iNET ZERO

your JNCIE training partner

iNET ZERO - JNCIE-ENT

Lab Preparation workbook

V2.1

Walktrough guide (2017)

For Juniper Networks ® - JNCIE-ENT Lab exam

https://t.me/learningnets

Copyright information

This walkthrough, iNET ZERO's JNCIE-ENT Lab Preparation Walkthrough v2.1 is developed by iNET ZERO.

All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means without the prior written permission of iNET ZERO.

This product cannot be used by or transferred to any other person. You are not allowed to rent, lease, loan or sell iNET ZERO training products including this workbook.

You are not allowed to modify, copy, upload, email or distribute this workbook in any way. This product may only be used and printed for your own personal use and may not be used in any commercial way.

Warning: Besides standard anti piracy techniques like document watermarks, this workbook also contains a **steganography ID** making this workbook unique and always traceable to the original buyer.

Juniper (c), Juniper Networks (c), JNCIE, JNCIP, JNCIS, JNCIA, Juniper Networks Certified Internet Expert, are registered trademarks of Juniper Networks, Inc.

ISBN/EAN

978-90-819227-1-5

https://t.me/learningnets

About iNET ZERO's content developers and authors:

Mohammad Reza Ebrahimi



Reza lives in Dusseldorf, Germany. He has 15 years of experience in the IT and Telecommunications industry. He has been working for large Telecom companies such Nokia, Nokia Siemens and Huawei as an IP planner and as an IP engineer. He started his Juniper journey using the Inetzero studying materials.He is a double JNCIE (JNCIE-ENT#412, JNCIE-SP#1945) and a JNCIP-SEC holder. He also has experience with Cisco systems and is a CCNP/CCDP certified.

Jörg Buesink



Jörg lives in the Netherlands near Amsterdam and has more than 13 years of experience in the IT and networking industry. He has worked for several large ISPs / service providers in the role of technical consultant, designer and network architect. He has extensive experience in network implementation, design and architecture and teached several networking classes. Jörg is triple JNCIE certified (JNCIE-ENT#21, JNCIE-SP#284 and JNCIE-SEC#30) as well as triple CCIE#15032 (Routing/ Switching, Service provider and Security), Cisco CCDE#20110002 and Huawei HCIE#2188 (Routing and Switching) certified.

lands. All rights reserved

https://t.me/learningnets

Table of Contents

General information 6
Exam strategy 6
Workbook updates and configuration files 8
iNET ZERO rack rental service 8
Chapter One: General System Features 9
Task 1: Initial System Configuration 17
Task 2: User Authentication and Authorization 25
Task 3: Syslog configuration 28
Task 4: SNMP configuration 30
Task 5: Firewall filters 31
Chapter 2: L2 Switching and Security 43
Task 1: L2 Switching Network Deployment 50
Task 2: Virtual Chassis 53
Task 3: VLAN Configuration 57
Task 4: Spanning Tree Configuration 75
Task 5: L2 Switching Security Features 89
Task 6: VRRP Configuration 93
Task 7: Provider bridging / Q-in-Q implementation 101
Chapter 3: IGP Routing 105
Task 1: Base Network and Virtual Router Deployment 115
Task 2: Multi Area OSPF Configuration 118
Task 3: External Configuration and Redistribution Policies 137
Task 4: Protocol Independent Routing and Policies 150
Chapter 4: BGP Routing 154
Task 1: Base Network Deployment 164
Task 2: Internal BGP Configuration 166
Task 3: External BGP Configuration 175
Task 4: BGP Routing policies 181
Task 5: Filter Based Forwarding 197
Chapter 5 : Multicast Routing and Switching 201
Task 1: Base configuration 208
Task 2: Multicast Configuration 210
Chapter 6: Class of Service 224
Task 1: Base Network Deployment 230
Task 2: SRX Class of Service 231
Task 3: EX class of Service 249
Chapter 7: A Full Day Lab Challenge 1 267

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 1: Initial System Configuration 267
Task 2: L2 switching and security 282
Task 3: IGP Routing 318
Task 4: BGP Configuration 347
Task 5: Multicast Routing and Switching 378
Task 6: Class of Service 388

Chapter 8: A Full Day Lab Challenge 2 401

Task 1: Initial System Configuration 401
Task 2: Layer2 Configuration 407
Task 3: Protocol Independent Routing 434
Task 4: IGP Routing 442
Task 5: BGP Routing 478
Task 6: Multicast Routing 511
Task 7: Class of Service 521

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

General information

Exam strategy

Lab preparation

To earn the prestigious JNCIE-ENT credential you must first pass the following Juniper Networks written exams: JNCIA-Junos, JNCIS-ENT and JNCIP-ENT and finally a very challenging 8 hour lab exam. You must have extensive hands-on experience with Junos Enterprise Routing and Switching on SRX and EX platforms. When preparing for the JNCIE-ENT lab exam ensure you fully understand the exam topics and objectives. The latest update for the JNCIE-ENT lab exam topics and objectives can be found on the Juniper Networks certification website.

Use this workbook to verify if you are ready for the JNCIE-ENT lab exam as it covers very challenging tasks specifically written for JNCIE-ENT lab preparation. The workbook tasks are written at the same difficulty level as the JNCIE-ENT lab exam and follow the same structure.

We recommend at least going twice through all the workbook tasks and the final lab challenge. The first time, work through all the tasks and use the solutions as often as you like. The second time try to solve all the tasks and final lab without looking at the solutions.

While using this workbook ensure you take study notes and record the time spend on each task. Time is of essence for the JNCIE-ENT lab exam. Try to complete the final lab exam within 8 hours!

As mentioned you need extensive hands-on experience on SRX and EX devices. If you do not have these devices available, you might also consider using iNET ZEROs rack rental service. When using iNETZERO's rack rental service you do no need to worry about patching cables, correct port numbering, software versions, etc as the iNET ZERO rack rental service will match the physical topologies used in this workbook.

Day before the exam

We advice you not to do extensive lab preparation anymore on this day as you should be ready by now, perhaps some review of your study notes, but that's it! On the day before your exam your goal is to stay relaxed and rest.

When travelling multiple hours to the JNCIE lab testing location, please ensure you book a hotel and arrive on the day before of your exam. If possible, visit the JNCIE lab-testing center and familiarize yourself with the location. Go to bed early, as tomorrow it will be a long day.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Exam day

When you wake up, ensure you dress comfortable and take a good breakfast. This will be a very long day! Ensure you arrive 30 minutes before the exam start time. Ensure you do not forget to take a valid ID with you, exam registration confirmation and if you prefer your own keyboard. If you forget your ID the JNCIE lab exam proctor will not allow you to take the exam. When it's time to enter the exam room, the JNCIE lab exam proctor will direct you to your workstation. Using the SecureCRT software you can access the JNCIE terminal server, which provides console and out of band access (fxp) to the exam devices.

When the proctor tells you that the exam begins, you are allowed to open the exam documents. Before immediately starting with the first question, we would like to advise you to read the entire exam and study the topology first to fully understand all requirements.

Be familiar with the notepad application as it is a real time saver. For example if you have to apply BGP configuration to multiple devices, with just some small changes then notepad is your friend.

Timing is important! Be quick, but pay attention to detail. A simple typo in a policy name can cause you to fail a complete task, even if the policy is working as expected. Ensure that you meet **all** requirements as mentioned in the exam documents. You will not be given partial points for tasks.

Do not assume any other requirement then mentioned in the tasks. If you are unsure about a task requirement then ask the proctor. While the JNCIE lab exam proctor cannot help you in solving the tasks, he or she can certainly help you in clarifying the task requirements.

Notify the JNCIE lab exam proctor immediately if you suspect a hardware issue. The proctor will verify if it is indeed a hardware issue. If it is, the proctor will fix the issue and might extend your time. If not, the proctor will tell you that the issue is caused by your own configuration and that no additional time will be given.

Make sure that you have at least ***one*** task correct in ***each*** section of the exam, if not you will fail the entire exam even if your total score is above the passing requirement!

Be sure to understand that a configuration mistake might result in point loss for multiple sections of the exam. As an example: if you apply a misconfigured firewall filter it might bring your BGP, IGP neighbors and perhaps also other protocols down, which results in point loss for multiple sections.

After the exam

That's it! There is nothing more that you can do for now. It's time to wait and hope you passed the JNCIE lab exam. It usually takes two to three weeks (15 business days) before you receive your results. If you pass great! if not, we recommend you to register again as soon as possible as all topics are still fresh in your memory! Some candidates wait a couple of months between attempts. This usually results in losing your configuration speed and it takes more time and energy to memorize all the topics again.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Workbook updates and configuration files

Together with this workbook you should have received configuration files that you will need while working through the workbook chapters. The workbook configurations .zip file consist of baseline and final configurations. If, for example, you would like to start with Chapter 4, you load the baseline configuration files for Chapter 4. We also include the final configurations for each chapter so you will be able to see all tasks completed and working.

Although we test our products and have an extensive review process, it still can happen that an error is found and needs to be corrected in either this workbook or the configuration files. Also iNET ZERO might make changes or additions to the workbook or the configuration files.

IMPORTANT: Please visit the following url regularly as it contains the procedure to obtain the latest version of our workbook or configuration files.

Workbook and configuration updates: **http://www.inetzero.com/pics/wb/ent/entv20.txt**

iNET ZERO rack rental service

Did you know that this workbook can be used in combination with our premium JNCIE rack rental service? Take a look at our website for the latest information www.inetzero.com

When using iNET ZERO rack rental service, please make sure to **reload** all devices after you have applied the chapter configurations. The reason for this is that the SRX devices might still be in "**flow**" mode, where "**packet mode**" is required for the JNCIE-ENT workbook tasks. It can also be that a reload is needed because of EX switch virtual-chassis or SRX clustering configurations reset.

Warning:

Please do **NOT change the root account password for any of our devices to prevent unnecessary password recovery. Thank you for your cooperation!**

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Chapter One: General System Features

This chapter focuses on initial system configuration and general system features. You will configure various tasks, such as configuring host names, management network access, user authentication and authorization, NTP, SNMP, syslog and RE protection firewall filters. You will be operating 8 devices, four SRXs (SRX1 to SRX4) and four EXs (EX1 to EX4). The topology for Chapter 1 is shown in figure 1.

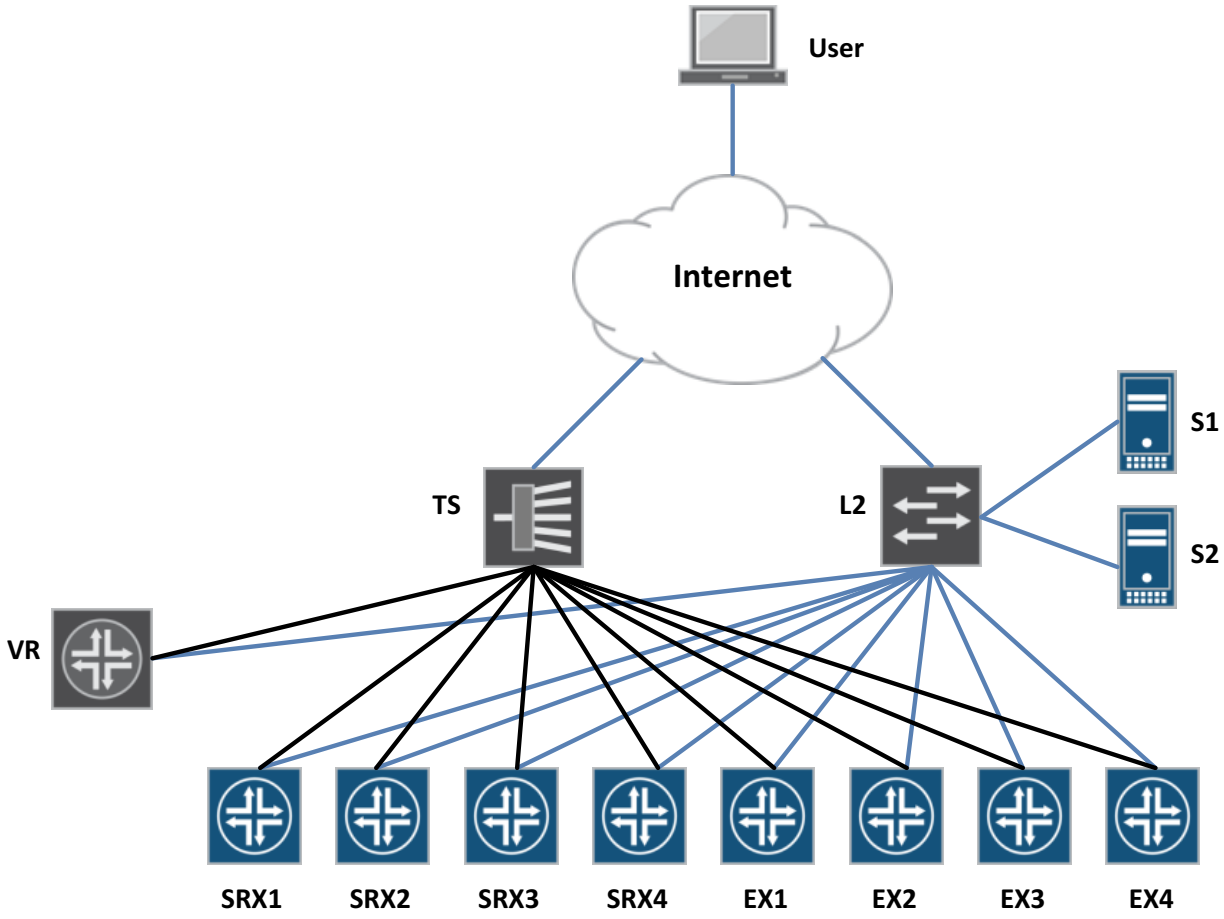


Figure 1

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

JNCIE-ENT candidates should be familiar with the following topics:

- NTP
- Syslog
- Authentication and Authorization
- Automation (scripting, archival)
- RPM
- Stateless firewall configuration
- VRRP
- IPv4 and IPv6 addressing
- Aggregated Ethernet
- GRE
- Logical tunnel interface
- BFD
- Time saver

Overview.

Junos OS, or JUNOS, is a FreeBSD-based operating system developed by Juniper. JUNOS is used across a whole portfolio of devices that includes routers, switches and security devices. JUNOS offers support for many of the most commonly used features within networks. The supported features are mostly consistent across the different device types that Juniper has to offer.

NTP

The Network Time Protocol (NTP) can be used to synchronize network devices to a single and accurate time source. NTP is supported up to version 4. Configuring the network with NTP will make troubleshooting a lot easier. This is because you will be able to correlate events from different devices by examining their timestamp.

The NTP configuration is done at the **[edit system ntp]** hierarchy level.

As soon as the time difference between the local device and the NTP server is more than 128 milliseconds, the NTP client will synchronize with the server.

If the time difference between the NTP client and the NTP server exceeds 1000 seconds, the device clock is not automatically synchronized. In this case, you can use the **set date ntp <*address*>** command to synchronize the device to the NTP server.

The configuration of an NTP-boot server will make sure that the NTP client synchronizes itself to the NTP server whenever the device boots. This will synchronize the NTP clients to the NTP server at boot time regardless of the time difference between the device and the NTP server.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

You can use the **show ntp status** and **show ntp associations** command to see the synchronization status**.** When multiple NTP servers are configured, the ***(asterisk) symbol displayed in **show ntp associations** displays the server that is selected for synchronization.

Syslog

JUNOS also offers a plethora of general logging features that can help track events and provide insights into what is happening on the network. The syslog configuration is located in [**edit system syslog]** hierarchy level. Here, you can specify both facility as well as severity. The facility defines what class of messages you want to log. By specifying the severity, you can control the level of detail that should be logged.

Apart from storing the information in a local file, the system can be configured to log to users, or remote syslog servers.

Local information is stored in the **/var/log** directory. To see the logs, you can use the **show log <*filename*>** command. The default system log messages are stored in the **messages** file.

Authentication and Authorization

Permissions and restrictions are managed through the configuration of classes. A class is a template that defines a set of privileges. These privileges can be granted to users by placing them in a class. On a device running Junos OS, every user has to belong to a class.

Inside the class configuration, the **permission** keyword is used to configure permission flags. Every permission flag represents a certain amount of actions/capabilities a user will be granted access to. Apart from the permission flags, you can also use the **deny-commands**, **deny-configuration, allow-commands** and **allow-configuration** statements. These statements allow you to use regular expressions, offering you a way to further increase the granularity in defining what users may or may not do.

To check the permissions that are assigned to the current user, you can issue the **show cli authorization** command.

Apart from local password authentication, JUNOS also supports remote authentication methods such as TACACs+ and RADIUS. Remote authentication methods can be combined with local password authentication. In such cases, local password authentication usually serves as a last resort for when IP connectivity with the remote authentication system is lost.

To make sure both Radius authenticated users and locally configured users can log in, use the **authentication-order [radius password]** configuration command.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

When the **authentication-order radius** command is used, local users can be authenticated only when radius server is unreachable.

Automation (scripting, archival)

Commit scripts should be either manually added to **/var/db/scripts/commit** or fetched from a remote server by specifying the URL in the configuration.

The syntax for this is as follows: **source "protocol ://username@host:/address/script_name"**

In this case, the **refresh** command is a mandatory command before committing to fetch the script from server.

The script configuration is located in **[edit system scripts commit]** hierarchy level.

The archival of the configuration can be specified in the **[edit system archival]** hierarchy level. The archival of the configuration can be either timed (**transfer-interval**) or per commit (**transfer-on-commit**). Transfer protocols that you can use for archival include both FTP and SCP.

RPM

Real-time Performance Monitoring, or RPM, can be used for service-level monitoring. Using RPM, the Junos device can measure network performance by examining round trip time, jitter and packet loss. The probe types include ICMP, HTTP, TCP and UDP.

For UDP and TCP probes, you have to configure the remote device as an RPM receiver in order to generate a response to the probe.

The RPM configuration is located in **[edit services rpm]** hierarchy level.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Stateless firewall configuration

The number one security measure to protect the control plane of a device running Junos OS is through the configuration and application of a firewall filter. A standard firewall filter can be applied to a JUNOS device's loopback interface. By doing so, all traffic to and from the routing-engine can be controlled. This helps to provide basic protection from potentially malicious traffic that is aimed at the routing-engine.

Firewall filters can also be used to manipulate, rate-limit, mark or control transit traffic that is passing through the device. When you are configuring a firewall filter, consider the following:

- Use the **log** option to log the packet header in a buffer within the PFE. The log can be shown using '**show firewall log**'
- Use the **syslog** option to log the packet header to a syslog file or to a remote syslog server
- Use **count** option to have a certain packet, or type of packet, increment a named counter
- Consider the use of a **prefix-list**, **apply-path** and a **source-prefix-filter** in a filter to simplify tasks
- Use **reject** to discard a packet and send and an ICMP destination unreachable message
- Use **discard** to silently drop traffic
- Use **policer** *policer-name* to subject the matched traffic to a policer
- Use port names instead of port numbers to simplify tasks and save time
- An implicit deny is applied at the end of a firewall filter
- Use the **insert** command to place a term before or after an existing term in the filter

VRRP

VRRP is a first hop router redundancy protocol, designed to eliminate the single point of failure in a scenario wherein a host relies on a default gateway. When configuring VRRP, take note of the following:

- The default VRRP priority is 100, the router with the highest priority becomes the active virtual-router
- By default, the Virtual IP address does not respond to ping. This can be changed using the **accept-data** option
- The virtual IP address must come from the same subnet that the interface address is configured in
- The cumulative priority cost of the tracked interfaces cannot exceed the current priority

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

IPv6 addressing

Take into account the following when dealing with IPv6 addressing configuration:

- An IPv6 address is 128 bits
- The first 48-bit prefix the routing-prefix, which is similar to the network ID in IPv4
- The next 16-bit is the subnet ID
- The last 64-bit make up the interface ID, which is similar to host ID in IPv4
- /127 mask is similar to /31 in IPv4 which can be used for point to point links
- There are three types of IPv6 addresses: unicast, multicast and anycast addresses
- There are three important IPv6 address scopes;
    - link-local: nodes connected to the same link can use this address to communicate with each other. The link-local addresses fall within the FE80::/10 prefix range
    - site-local: (deprecated) addresses used within a private IPv6 network. Site-local addresses fall within the FEC0::/10 prefix range
    - global: globally routable, used to connect to the Internet

Aggregated Ethernet

JUNOS also supports physical interface bundling, or link aggregation. This enables the grouping of multiple physical interfaces with the same properties (media, speed, duplex mode) into one logical interface called an aggregated interface, or AE. The use of link aggregation allows for increased redundancy and throughput. In general, AE interfaces can be used in the same way regular interfaces are used (as a Layer 3 interface, as a layer 2 trunk interface, etc.). The members of the AE interface can be a statically configured to participate in the LAG or they can use the LACP protocol . The LACP protocol allows for the automatic detection and deletion of member links. In addition to his, LACP also has certain link monitoring capabilities that can be used to check whether both ends of the link belong to the same AE bundle. The LACP protocol can be configured to operate in different modes and send messages at different intervals.

To configure an aggregated interface, ae3 for example, you should follow these three steps:

1. Define the device count under **[edit chassis aggregated-devices ethernet]**. Note that when you want to configure **ae3**, the **device-count** needs to be set to **4** (numbering starts at **ae0**).
2. Assign physical interface to the aggregated interface. Adding interfaces to **ae3** is done under the gigether-options of the physical interfaces, like this: **set interfaces ge-1/0/0 gigether-options 802.3ad ae3**
3. The last step is the configuration of the parameters for the logical AE interface itself. The parameters can include things like IPv4 address, IPv6 address, firewall filters, trunking, etc. The LACP and minimum-links configuration can be found under the **[edit interfaces ae3 aggregated-ether-options]** stanza.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

After configuring the interface, test the layer 3 connectivity by using ping to ensure that the interface is working.

GRE

Generic Routing Encapsulation (GRE) is an IP encapsulation protocol that can be used to transport packets. To configure a GRE interface on an SRX device, you should define the **source**, the **destination** and the **protocol family** under the **gre-0/0/0** interface. Note that:

1. The source and intermediate devices should have a route towards the destination.
2. You can, optionally, use **BFD** to monitor GRE interfaces.
3. Because 24 bytes are added to IP datagram, you should either enable fragmentation or increase the MTU along the path.

Logical tunnel interface

Logical tunnel (**lt**) interfaces can be used to connect two virtual routers or logical systems.

With the **peer-unit** command under the interface unit, you can specify what lt- subinterfaces are connected together. Depending on the networking need, the encapsulation can vary. On an lt interface, the protocol families **ethernet**, **frame-relay**, **inet**, **inet6**, **iso** and **mpls** are supported.

BFD

Bidirectional Forwarding Detection (BFD) is a light-weight hello mechanism that can be used to detect failures. BFD can be configured for static as well as dynamic routing protocols. Some example protocols for which you can use BFD include OSPF, BGP and LACP. You can also use BFD for static routes.

To configure BFD, you can use the **bfd-liveness-detection minimum-interval** command to specify transmit and receive intervals. The **multiplier** command specifies the number of missed hellos, after which the subscribed protocols will be notified that BFD is "down". The default hello multiplier is **3**.

BFD is configured in the hierarchy where you want to have it detect failures. So, for example, you can configure it under the [**protocols ospf area x interface**] configuration as well as under the [**routing-options static route x/x**] configuration.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Time saver

During the exam, many tasks will require you to apply identical configurations on multiple devices. The **load merge terminal relative** and **show | display set** are very helpful in helping you to quickly apply configuration on multiple devices.

When you forget a command, consider using either the **help apropos <*keyword*>** or the **help topic <*topic*>** commands.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 1: Initial System Configuration

(Please note that we explain the configuration details for one device only and we do not show any **commit** command in this walkthrough either.)

1) Load baseline configuration on all devices by issuing following command:

```
root@srx1# load override terminal
```

Use ctrl+D at the end then commit configuration.

You need to reboot SRX devices to change the forwarding mode to packet mode.

To verify the current device mode, use the following command:

```
root@srx1# run show security flow status
  Flow forwarding mode:
    Inet forwarding mode: packet based (reboot needed to change to
flow based)
    Inet6 forwarding mode: drop
    MPLS forwarding mode: packet based (reboot needed to change to
drop)
    ISO forwarding mode: drop
  Flow trace status
    Flow tracing status: off
```

2) Configure a hostname on each device.

```
root@srx1# set system host-name srx1
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

3) Configure the OoB management IPv4 and IPv6 addresses on each device.

```
root@srx1# set interfaces ge-0/0/0 unit 0 description "OoB management
connection"

root@srx1# set interfaces ge-0/0/0 unit 0 family inet address
10.10.1.1/24

root@srx1# set interfaces ge-0/0/0 unit 0 family inet6 address
fdb3:0309:419b:9ec4::1/64
```

To verify the configuration, issue a '**show interfaces terse**' or ping a host address that is already active on the OoB management network:

```
root@srx1# run ping 10.10.1.254
PING 10.10.1.254 (10.10.1.254): 56 data bytes
64 bytes from 10.10.1.254: icmp_seq=0 ttl=64 time=1.912 ms
64 bytes from 10.10.1.254: icmp_seq=1 ttl=64 time=7.307 ms
64 bytes from 10.10.1.254: icmp_seq=2 ttl=64 time=1.595 ms
64 bytes from 10.10.1.254: icmp_seq=3 ttl=64 time=1.468 ms
```

4) Enable each device to accept SSH and Telnet connections. Make sure that user **root** is not allowed to login using SSH. Do not allow more than 4 simultaneous SSH and Telnet connections and no more than 2 connections per minute.

All of the management protocols are configured under the [**system services**] stanza. The '**root-login deny**' will prevent remote SSH users from logging in as root.

There is no '**root-login**' configuration command for telnet. This is because by design, you cannot login as the root user via telnet. If you login via telnet, you can change the user to root by issuing the following command: '**start shell user root**'. You will login to the shell as the root user. From there, you can go to the cli using the command '**cli**'.

Furthermore, we specify a '**connection-limit**' (simultaneous users) and we use '**rate-limit**' to limit the number of connections to 2 per minute.

```
root@srx1# set system services ssh root-login deny

root@srx1# set system services ssh connection-limit 4

root@srx1# set system services ssh rate-limit 2

root@srx1# set system services telnet connection-limit 4

root@srx1# set system services telnet rate-limit 2
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

5) Configure static routes to the management networks. For IPv4, use 10.10.10/24 with next-hop 10.10.1.254. For IPv6, use fd19:2f8d:105a:69bd::/64 with next-hop fdb3:0309:419b:9ec4::fe. Make sure that these networks are never redistributed to any dynamic routing protocol. Ensure that the device is reachable while RPD is not started or is in the process of grateful restart:

Static routes are configured under the [**routing-options**] stanza. The options that we will need to configure for static routes are the following:

- **no-readvertise :** static routes configured with this option will never become eligible for redistribution into a dynamic routing protocol.
- **retain:** statically configured routes with this option are retained in the forwarding table when the routing protocol process shuts down.

First, we configure the IPv6 static route:

```
root@srx1# set routing-options rib inet6.0 static route
fd19:2f8d:105a:69bd::/64 next-hop fdb3:0309:419b:9ec4::fe

root@srx1# set routing-options rib inet6.0 static route
fd19:2f8d:105a:69bd::/64 retain

root@srx1# set routing-options rib inet6.0 static route
fd19:2f8d:105a:69bd::/64 no-readvertise
```

Then we move over to the IPv4 static route:

```
root@srx1# set routing-options static route 10.10.10.0/24 next-hop
10.10.1.254

root@srx1# set routing-options static route 10.10.10.0/24 retain

root@srx1# set routing-options static route 10.10.10.0/24 no-
readvertise
```

Through the configuration of a **backup-router,** we can make the device install a route to the management network when the RPD process is not running. This may happen when the RPD crashes or when the devices is booting. We configure the **backup-router** like this:

```
root@srx1# set system backup-router 10.10.1.254

root@srx1# set system backup-router destination 10.10.10.0/24

root@srx1# set system inet6-backup-router fdb3:0309:419b:9ec4::fe

root@srx1# set system inet6-backup-router destination
fd19:2f8d:105a:69bd::/64
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

If the **backup-router** would have been configured without a destination network, a default route would be installed towards the backup-router whenever the RPD process is not running.

Verification:

```
root@srx1# run show route table inet.0
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.10.1.0/24       *[Direct/0] 00:32:58
                    > via ge-0/0/0.0
10.10.1.1/32       *[Local/0] 00:32:58
                      Local via ge-0/0/0.0
10.10.10.0/24      *[Static/5] 00:32:58
                    > to 10.10.1.254 via ge-0/0/0.0

root@srx1# run show route table inet6.0
inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
fd19:2f8d:105a:69bd::/64
                   *[Static/5] 00:33:06
                    > to fdb3:309:419b:9ec4::fe via ge-0/0/0.0
fdb3:309:419b:9ec4::/64
                   *[Direct/0] 00:33:06
                    > via ge-0/0/0.0
fdb3:309:419b:9ec4::1/128
                   *[Local/0] 00:33:06
                      Local via ge-0/0/0.0
fe80::/64          *[Direct/0] 00:33:06
                    > via ge-0/0/0.0
fe80::b2a8:6e0f:fcb9:6100/128
                   *[Direct/0] 00:33:06
                    > via lo0.0
fe80::b2a8:6eff:feb9:6100/128
                   *[Local/0] 00:33:06
                      Local via ge-0/0/0.0
```

6) Configuring server S1 as the DNS server:

When we configure a device with a DNS server, the device will be able to perform a name lookup and resolve a hostnames into an addresses.

```
root@srx1# set system name-server 10.10.10.1
```

Note: if there would have been additional requirements requiring static host to IP mappings, we could have completed these by using the following configuration command **set system static-host-mapping *<name.com>* inet *<IP address>***.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

7) Set the time zone to *Europe/Amsterdam* on all of the devices:

```
root@srx1# set system time-zone Europe/Amsterdam
```

Verification:

```
root@srx1# run show system uptime
Current time: 2014-12-13 18:43:31 CET
System booted: 2014-12-13 16:39:24 CET (02:04:07 ago)
Protocols started: 2014-12-13 16:41:17 CET (02:02:14 ago)
Last configured: 2014-12-13 18:00:23 CET (00:43:08 ago) by root
 6:43PM  up 2:04, 1 user, load averages: 0.07, 0.10, 0.08
```

8) Make sure that srx1 synchronizes its time with NTP server S1. Ensure all NTP messages are authenticated using MD5 with password *workbook*. Configure NTP synchronization with server S1 at boot time:

NTP is configured under the [**system ntp**] stanza. As stated previously, NTP time synchronization is very important in that it helps you get a synchronized time stamp for log messages on the devices in your network.

In this case, we use the **boot-server** command to have the system synchronize with the NTP server during startup:

```
root@srx1# set system ntp boot-server 10.10.10.1
```

We also need to make sure that all NTP messages are authenticated. Because of this, we first need to configure an authentication-key:

```
root@srx1# set system ntp authentication-key 1 type md5
```

```
root@srx1# set system ntp authentication-key 1 value workbook
```

We then specify the NTP server and reference the same key. The following configuration will make sure that packets send by the local device will be authenticated:

```
root@srx1# set system ntp server 10.10.10.1 key 1
```

We also need to make sure that the router will only accept authenticated NTP messages. We can achieve this by specifying a '**trusted-key**' and referencing the previously configured key:

```
root@srx1# set system ntp trusted-key 1
```

Verification:

```
root@srx1# run show ntp associations
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
     remote           refid      st t when poll reach   delay   offset
jitter
==============================================================================
========
*10.10.10.1      LOCAL(0)        11 -   97  128  377    1.613   -3.350
1.118
 10.10.1.255     .BCST.          16 -    -   64    0    0.000    0.000
4000.00
root@srx1# run show ntp status
status=0664 leap_none, sync_ntp, 6 events, event_peer/strat_chg,
version="ntpd 4.2.0-a Sat Aug 25 08:03:34 UTC 2012 (1)",
processor="octeon", system="JUNOS11.4R5.5", leap=00, stratum=12,
precision=-17, rootdelay=1.613, rootdispersion=29.343, peer=39084,
refid=10.10.10.1,
reftime=d836fa48.1fae6061  Sat, Dec 13 2014 18:47:52.123, poll=7,
clock=d836fb81.d3e3f217  Sat, Dec 13 2014 18:53:05.827, state=4,
offset=-3.350, frequency=7.042, jitter=2.345, stability=0.027
```

9) Ensure that all other devices synchronize their clock with srx1 using broadcast NTP messages sourced from the OoB interface. MD5 authentication with password *workbook* must be used. Configure all your devices for NTP synchronization with srx1 at boot time.

First, we configure the srx1 with the following configuration command:

```
root@srx1# set system ntp broadcast 10.10.1.255 key 1
```

This will make srx1 operate in broadcast mode. The srx will start sending periodic broadcast messages to the configured address. By specifying the key, we configure srx1 to authenticate all NTP packets it sends.

After configuring srx1, we configure the rest of the devices to use srx1 as an NTP server. The other devices need to be configured to receive authenticated NTP packets:

```
root@srx2# set system ntp authentication-key 1 type md5

root@srx2# set system ntp authentication-key 1 value workbook

root@srx2# set system ntp trusted-key 1
```

And the devices need to be configured to listen for NTP broadcast messages:

```
root@srx2# set system ntp broadcast-client
```

Finally, we also need to configure the other devices to synchronize to srx1 when they are booting:

```
root@srx2# set system ntp boot-server 10.10.1.1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification:

```
root@srx2# run show ntp status
status=0664 leap_none, sync_ntp, 6 events, event_peer/strat_chg,
version="ntpd 4.2.0-a Sat Aug 25 08:03:34 UTC 2012 (1)",
processor="octeon", system="JUNOS11.4R5.5", leap=00, stratum=13,
precision=-17, rootdelay=4.023, rootdispersion=41.101, peer=42989,
refid=10.10.1.1,
reftime=d836fbcb.1eef83d7  Sat, Dec 13 2014 18:54:19.120, poll=6,
clock=d836fbf7.54d46083  Sat, Dec 13 2014 18:55:03.331, state=4,
offset=1.533, frequency=0.991, jitter=2.845, stability=0.008

[edit]
root@srx2# run show ntp associations
     remote           refid      st t when poll reach   delay   offset
jitter
==============================================================================
========
*10.10.1.1       10.10.10.1      12 -   49   64  377    2.421    1.533
2.122
```

10) Archival configuration:

The archival of the configuration is helpful when the internal storage of the device fails and when the device itself fails. It can also be convenient to use '**transfer-on-commit**' in addition to a daily backup. Suppose you have a daily backup policy and a device crashes at the end of the day. With the '**transfer-on-commit**' configuration command active, the changes from that day would not be lost.

```
root@srx1# set system archival configuration transfer-on-commit

root@srx1# set system archival configuration archive-sites
ftp://lab@10.10.10.1 password lab123
```

Note:

If you find yourself struggling to remember the format in which you need to specify the backup site, issue the following command: '**help topic system archive-sites**'.

Verification:

After issuing a commit, you can check the archival using the following command:

```
root@srx1# run show system configuration archival
/var/transfer/config/:
-rw-r-----  1 root  wheel       2787 Dec 13 15:51
Sun_juniper.conf.gz_20141213_145114
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

11) Make sure that the devices do not send ICMP redirects.

To disable the sending of ICMP redirects for the entire device, we configure the following:

```
root@srx1# set system no-redirects
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 2: User Authentication and Authorization

1) Configure the authentication method where the device first tries to authenticate users on the Radius server. If the authentication is not successful, have the device use a local password. Use S1 as a Radius server. Configure the device to retry contacting the Radius server 3 times and configure a timeout of 3 secons. Use *workbook* as the Radius shared secret.

JUNOS can authenticate users by checking a configured username/password that is stored in a local database on the device. Additionally, JUNOS can consult a centralized system such as RADIUS and TACACS+ to authenticate users that want to login.

The radius server is configured in the [**radius-server**] stanza. The configuration is pretty straightforward. Do mind the '**authentication-order**' command. This command specifies in what order the device will authenticate users.

```
root@srx1#set system radius-server 10.10.10.1 secret workbook

root@srx1#set system radius-server 10.10.10.1 timeout 3

root@srx1#set system radius-server 10.10.10.1 retry 3

root@srx1#set system authentication-order radius

root@srx1#set system authentication-order password
```

Note: A caveat could be to consult the local database only when either Radius or TACACS+ is unavailable. In these cases, you do not configure the password option. Without the "password" option, the local database is only checked when the radius or Tacacs server is not reachable.

2) Create a new user *lab*, with password *lab123*, on every device. Grant the user super user privileges.

Locally confgured users need a username, a password and a class. The class is a template that defines a group of privileges that users enjoy when they are logged into the device. The class can also specifiy what commands users are not allowed to use and, additionally, the class can determine how long a session can be idle before the user is automatically logged off.

There are four default login classes available in JUNOS. These are the **operator** (clear, network, reset, trace, view), **read-only** (view), **super-user**/**superuser** (all) and **unauthorized** (none) class. In this case, it is easy to use the preconfigured '**super-user**' class. We optionally configured the uid. When left unconfigured, the system assigns one automatically.

```
root@srx1#set system login user lab uid 2001

root@srx1#set system login user lab class super-user
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1#set system login user lab authentication plain-text-password
New password:lab123
Retype new password:lab123
```

Verification:

After issuing a commit, login again as user 'lab'.

3) Configure additionl users on all devices as defined in table 3.

```
root@srx1#set system login class limited permissions network
root@srx1#set system login class limited allow-commands "show
interfaces *"
root@srx1#set system login class limited deny-commands
"(telnet)|(ssh)"
root@srx1#set system login class privileged permissions all
root@srx1#set system login class privileged deny-commands
"(clear)|(configure)|(edit)|(start shell)"
root@srx1#set system login user noc uid 2000
root@srx1#set system login user noc class privileged
root@srx1#set system login user noc authentication plain-text-password
New password:noc123
Retype new password:noc123
root@srx1#set system login user ops uid 2002
root@srx1#set system login user ops class limited
root@srx1#set system login user ops authentication plain-text-password
New password:ops123
Retype new password:ops123
```

Verification:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

To verify your configuration, first login using the newly configured usernames and passwords. Then check the permissions with the **show cli authorization** command.

4) Make sure that for all users listed in table 3, sessions automatically logout after 5 minutes of idle time-out. Ensure that new user classes will automatically inherit this setting.

To configure an idle timeout for all current and future users, we can issue the following configuration commands:

```
root@srx1# set groups class-logout system login class <*> idle-timeout
5
```

This configuration commands creates a group configuration. What the group configuration could do is apply the '**idle-timeout 5**' to any class that is configured in the [**system login class**] stanza.

In order for the group configuration to have any effect, it needs to be applied. To apply the group, we issue a second configuration statement in [**system login**] stanza:

```
root@srx1# set system login apply-groups class-logout
```

The group does not have to be applied at the place where it affects the configuration. Another strategy could be to list all applied groups together at the top of the configuration. If we would have wanted to do so, we could have chosen to simply apply the group using the ' **set apply-groups class-logout**' configuration command.

To check whether or not the configured group has been applied correctly, we can issue the following command:

```
root@srx1#show | display inheritance
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 3: Syslog configuration

1) All critical messages, regardless of the facility, should be displayed on the user terminals. This should apply to all users that are logged in.

The JUNOS syslog facility can be configured to log to a file, to a remote location, but also to user CLI sessions. To log critical messages to any user, we configure the following:

```
root@srx1#set system syslog user * any critical
```

NOTE: the asterisk sign * represents all users.

2) All messages with severity info or higher need to be stored in the default log file. Instead of showing "last message repeated x times", the system should log the real message multiple times.

To start sending messages to a file, we configure this file under the [**system syslog file**] configuration stanza. When we configure this file, we need to specify what type of messages should be logged. In this case, we specify **'any info'**.

```
root@srx1#set system syslog file messages any info
```

By default, when a message repeats itself, you will not see lot of occurrences of that message. Instead, you will see the "last message repeated x times" statement appear in the log. To disable this default behavior, we issue the following configuration command:

```
root@srx1#set system syslog file messages allow-duplicates
```

3) A file named "interactive-commands" should record all the commands that users execute.

The facility we need to log to a file here is actually called 'interactive-commands'. We send all messages from this facility to the file by specifying 'any':

```
root@srx1#set system syslog file interactive-commands interactive-
commands any
```

4) All messages with severity info and higher should be stored on the S1 server in a structured format. The messages should be sourced by the OOB interface IP address.

Logging to a remote host is configured under the [**system syslog host**] configuration stanza. We specify 'any info', the same as what we configured for the file 'messages':

```
root@srx1#set system syslog host 10.10.10.1 any info
```

We then specify the OoB address as the source-address to use when sending these messages:

```
root@srx1#set system syslog host 10.10.10.1 source-address 10.10.1.1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Lastly, we configure the messages with the 'structured-data' configuration command. This format can make it easier for automated applications to read the information from a message.

```
root@srx1#set system syslog host 10.10.10.1 structured-data
```

5) Ensure that the maximum number of archived log files is 5. Each individual file should contain up to 1 Mb of log information.

We can configure characteristics that apply to all of the files configured in the [**system syslog file**] stanza in the [**system syslog archive**] stanza.

```
root@srx1#set system syslog archive size 1m

root@srx1#set system syslog archive files 5
```

Note: values configured specifically for a file override these statements.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 4: SNMP configuration

1) Configure the device with the read-only community string *workbook*. Make sure only server S1 can poll the device:

We first configure a clien-list. We only need to add 1 item to the list, the S1 server:

```
root@srx1#set snmp client-list S1 10.10.10.1/32
```

We then configure the snmp community and specify the client-list we just configured:

```
root@srx1#set snmp community workbook client-list-name S1
```

In this case, we chose to configure a client-list. We could have also chosen to configure a prefix-list and apply it using the same command ' **set snmp community workbook client-list-name *<prefix-list>***'. It is also possible to use both a prefix-list as well as a client-list at the same time. If you chose to do so, the prefix-list and the client-list cannot be configured with the same names.

Additionally, we configure the community to be read-only using the following command:

```
root@srx1#set snmp community workbook authorization read-only
```

2) Configure all devices to send SNMP traps to server S1 in case of link and/or chassis events. Make sure that traps are only send using the OoB management interface.

Traps are configured under the [**snmp trap-group**] stanza. We configure the device to send traps for both categories:

```
root@srx1#set snmp trap-group S1 categories chassis
root@srx1#set snmp trap-group S1 categories link
```

To have the device send the traps to S1, we configure the S1 server IP as a target:

```
root@srx1#set snmp trap-group S1 targets 10.10.10.1
```

To check SNMP statistics you can use the following command:

```
root@srx1# show snmp statistics
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 5: Firewall filters

In this task, we protect the Routing Engine using a firewall filter.

1) Configure an IPv4 firewall filter that allows the following protocols AH, ESP, BFD, VRRP, OSPF, BGP, PIM, IGMP and MSDP. Make sure that for the protocols VRRP, OSPF and PIM only messages from directly connected subnets are accepted.

To be able to only accept messages from directly connected subnets in certain terms, we will configure a prefix-list using 'apply-path':

```
root@srx1#set policy-options prefix-list direct-subnets apply-path
"interfaces <*> unit <*> family inet address <*>"
```

This prefix-list will automatically add all IP addresses configured under any of the interfaces on the device. Use the command '**show policy-options prefix-list direct-subnets | display inheritance**' to examine what subnets are added to the prefix list. We can reference this filter in the firewall terms where we need to permit only directly connected subnets.

We then move to the configuration of the firewall filter. We configure an individual term for each protocol. We use the protocol name instead of the port number to save time:

```
root@srx1#set firewall family inet filter protect-re-ipv4 term ah from
protocol ah

root@srx1#set firewall family inet filter protect-re-ipv4 term ah then
accept

root@srx1#set firewall family inet filter protect-re-ipv4 term esp
from protocol esp

root@srx1#set firewall family inet filter protect-re-ipv4 term esp
then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term bfd
from source-prefix-list direct-subnets

root@srx1#set firewall family inet filter protect-re-ipv4 term bfd
from protocol udp

root@srx1#set firewall family inet filter protect-re-ipv4 term bfd
from destination-port 3784

root@srx1#set firewall family inet filter protect-re-ipv4 term bfd
then accept
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1#set firewall family inet filter protect-re-ipv4 term vrrp
from source-prefix-list direct-subnets

root@srx1#set firewall family inet filter protect-re-ipv4 term vrrp
from protocol vrrp

root@srx1#set firewall family inet filter protect-re-ipv4 term vrrp
then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term ospf
from source-prefix-list direct-subnets

root@srx1#set firewall family inet filter protect-re-ipv4 term ospf
from protocol ospf

root@srx1#set firewall family inet filter protect-re-ipv4 term ospf
then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term bgp-
source from protocol tcp

root@srx1#set firewall family inet filter protect-re-ipv4 term bgp-
source from source-port bgp

root@srx1#set firewall family inet filter protect-re-ipv4 term bgp-
source then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term bgp-
destination from protocol tcp

root@srx1#set firewall family inet filter protect-re-ipv4 term bgp-
destination from destination-port bgp

root@srx1#set firewall family inet filter protect-re-ipv4 term bgp-
destination then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term pim
from source-prefix-list direct-subnets

root@srx1#set firewall family inet filter protect-re-ipv4 term pim
from protocol pim

root@srx1#set firewall family inet filter protect-re-ipv4 term pim
then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term igmp
from protocol igmp
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1#set firewall family inet filter protect-re-ipv4 term igmp
then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term msdp-
source from protocol tcp

root@srx1#set firewall family inet filter protect-re-ipv4 term msdp-
source from source-port msdp

root@srx1#set firewall family inet filter protect-re-ipv4 term msdp-
source then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term msdp-
destination from protocol tcp

root@srx1#set firewall family inet filter protect-re-ipv4 term msdp-
destination from destination-port msdp

root@srx1#set firewall family inet filter protect-re-ipv4 term msdp-
destination then accept
```

Note: since *from port* is not supported on EX switches, on that platform, we use source and destination port instead.

2) Configure the firewall filter to accept the RADIUS, DNS and SNMP management protocols only from the 10.10.10/24 network.

```
root@srx1#set firewall family inet filter protect-re-ipv4 term snmp
from source-address 10.10.10.0/24

root@srx1#set firewall family inet filter protect-re-ipv4 term snmp
from protocol udp

root@srx1#set firewall family inet filter protect-re-ipv4 term snmp
from destination-port snmp

root@srx1#set firewall family inet filter protect-re-ipv4 term snmp
then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term radius
from source-address 10.10.10.0/24

root@srx1#set firewall family inet filter protect-re-ipv4 term radius
from protocol udp

root@srx1#set firewall family inet filter protect-re-ipv4 term radius
from source-port radius
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1#set firewall family inet filter protect-re-ipv4 term radius
then accept

root@srx1#set firewall family inet filter protect-re-ipv4 term dns
from source-address 10.10.10.0/24

root@srx1#set firewall family inet filter protect-re-ipv4 term dns
from protocol udp

root@srx1#set firewall family inet filter protect-re-ipv4 term dns
from source-port domain

root@srx1#set firewall family inet filter protect-re-ipv4 term dns
then accept
```

3) Accept NTP messages only from the 10.10.10/24 and the 10.10.1/24 subnets.

We permit the source as well as the destination port for NTP. We first specify the source addresses:

```
root@srx1#set firewall family inet filter protect-re-ipv4 term ntp
from source-address 10.10.10.0/24

root@srx1#set firewall family inet filter protect-re-ipv4 term ntp
from source-prefix-list direct-subnets

root@srx1#set firewall family inet filter protect-re-ipv4 term ntp
from protocol udp

root@srx1#set firewall family inet filter protect-re-ipv4 term ntp
from source-port ntp

root@srx1#set firewall family inet filter protect-re-ipv4 term ntp
then accept
```

The RE filter will be applied in the input direction. Therefore, the destination port for NTP only needs to be allowed to the locally configured subnets:

```
root@srx1#set firewall family inet filter protect-re-ipv4 term ntp-
local from source-prefix-list direct-subnets

root@srx1#set firewall family inet filter protect-re-ipv4 term ntp-
local from protocol udp

root@srx1#set firewall family inet filter protect-re-ipv4 term ntp-
local from destination-port ntp

root@srx1#set firewall family inet filter protect-re-ipv4 term ntp-
local then accept
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

4) Accept SSH and Telnet from the 10.10.1/24 management network.

We start by allowing SSH messages sourced from the management network destined to the local device:

```
root@srx1#set firewall family inet filter protect-re-ipv4 term ssh
from source-address 10.10.1.0/24

root@srx1#set firewall family inet filter protect-re-ipv4 term ssh
from protocol tcp

root@srx1#set firewall family inet filter protect-re-ipv4 term ssh
from destination-port ssh

root@srx1#set firewall family inet filter protect-re-ipv4 term ssh
then accept
```

We configure an additional term to enable the local device to use SSH to log into other devices residing in the management network:

```
root@srx1#set firewall family inet filter protect-re-ipv4 term ssh-
local from source-address 10.10.1.0/24

root@srx1#set firewall family inet filter protect-re-ipv4 term ssh-
local from protocol tcp

root@srx1#set firewall family inet filter protect-re-ipv4 term ssh-
local from source-port ssh

root@srx1#set firewall family inet filter protect-re-ipv4 term ssh-
local then accept
```

When the local device logs into another device on the management network using SSH, the source address will be 10.10.1.x and the destination port will be SSH. The other device will reply with a source address from the management range and it will use the well-known SSH port as a source port in these messages. These replies need to be allowed by our inbound Routing Engine filter.

We apply the same logic to the telnet terms:

```
root@srx1#set firewall family inet filter protect-re-ipv4 term telnet
from source-address 10.10.1.0/24

root@srx1#set firewall family inet filter protect-re-ipv4 term telnet
from protocol tcp

root@srx1#set firewall family inet filter protect-re-ipv4 term telnet
from destination-port telnet

root@srx1#set firewall family inet filter protect-re-ipv4 term telnet
then accept
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1#set firewall family inet filter protect-re-ipv4 term telnet-local from source-address 10.10.1.0/24

root@srx1#set firewall family inet filter protect-re-ipv4 term telnet-local from protocol tcp

root@srx1#set firewall family inet filter protect-re-ipv4 term telnet-local from source-port telnet

root@srx1#set firewall family inet filter protect-re-ipv4 term telnet-local then accept
```

5) Configure the SRX devices to accept ICMP and traceroute messages while using a policer to limit the amount of traffic.

First, we configure the policer on our SRX devices:

```
root@srx1#set firewall policer re-policer if-exceeding bandwidth-limit 100k

root@srx1#set firewall policer re-policer if-exceeding burst-size-limit 25k

root@srx1#set firewall policer re-policer then discard
```

Then, we configure a term for ICMP and a term for traceroute and specify the previously configured policer:

```
root@srx1#set firewall family inet filter protect-re-ipv4 term icmp from protocol icmp

root@srx1#set firewall family inet filter protect-re-ipv4 term icmp then policer re-policer

root@srx1#set firewall family inet filter protect-re-ipv4 term icmp then accept


root@srx1#set firewall family inet filter protect-re-ipv4 term traceroute from protocol udp

root@srx1#set firewall family inet filter protect-re-ipv4 term traceroute from destination-port 33434-33534

root@srx1#set firewall family inet filter protect-re-ipv4 term traceroute then policer re-policer
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1#set firewall family inet filter protect-re-ipv4 term
traceroute then accept
```

Note: by default, a policer is 'term specific'. This means that an individual instance of the policer is generated for every term in which the policer is used.

6) Configure the ex devices to accept ICMP and traceroute messages.

This configuration is the similar to the configuration applied to the srx. Only the policer is missing, this is because the ex does not have the same capablities as the srx:

```
root@ex1#set firewall family inet filter protect-re-ipv4 term icmp
from protocol icmp

root@ex1#set firewall family inet filter protect-re-ipv4 term icmp
then accept

root@ex1#set firewall family inet filter protect-re-ipv4 term
traceroute from protocol udp

root@ex1#set firewall family inet filter protect-re-ipv4 term
traceroute from destination-port 33434-33534

root@ex1#set firewall family inet filter protect-re-ipv4 term
traceroute then accept
```

Note: you can rate limit ICMP for IPv4 and IPv6 on EX switches. The knob for this is found under the [**system internet-options icmpv4-rate-limit**] and the [**system internet-options icmpv6-rate-limit**] stanza. You can specify a 'bucket-size' or a 'packet-rate'.

7) Protect the srx devices from TCP syn/rst/fin attacks. Use a policer and drop excess traffic.

To protect the SRX devices from syn/rst/fin attacks, we start with the configuration of a policer:

```
root@srx1#set firewall policer tcp-flood-limit if-exceeding bandwidth-
limit 500k

root@srx1#set firewall policer tcp-flood-limit if-exceeding burst-
size-limit 125k

root@srx1#set firewall policer tcp-flood-limit then discard
```

We then configure a term that matches on '**tcp-flags**':

```
root@srx1#set firewall family inet filter protect-re-ipv4 term tcp-
flood from protocol tcp

root@srx1#set firewall family inet filter protect-re-ipv4 term tcp-
flood from tcp-flags "(syn & !ack) | fin | rst"
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1#set firewall family inet filter protect-re-ipv4 term tcp-flood then policer tcp-flood-limit

root@srx1#set firewall family inet filter protect-re-ipv4 term tcp-flood then accept
```

We have already configured a lot of terms that accept traffic. Since this is a terminating action, that traffic would never be subjected to this newly configured term (and it's policer). Through the use of the **insert** command, we can move this term to the top to make sure all traffic passes this term:

```
root@srx1#insert firewall family inet filter protect-re-ipv4 term tcp-flood before term ah
```

8) Configure the firewall filter to discard any other traffic and count the packets.

In the final term, we use the discard option to silently drop traffic. We also need to configure a named counter. This is done using the '**count**' option followed by a name. This can be any name:

```
root@srx1#set firewall family inet filter protect-re-ipv4 term last then count dropped-ipv4

root@srx1#set firewall family inet filter protect-re-ipv4 term last then discard
```

Note: after you have configured a named counter, the operational mode command 'show firewall' should display this counter and show you whether or not there have been any matches so far.

9) Apply the firewall fitler to ensure that it is used fot the protection of the Routing Engine.

We apply the filter to the loopback interface on the inbound direction. By doing so, we ensure that all inbound traffic destined to the Routing Engine will need to pass this filter:

```
root@srx1#set interfaces lo0 unit 0 family inet filter input protect-re-ipv4
```

10) Configure an IPv6 firewall filter allowing protocol messages from OSPF and BGP.

When configuring IPv6 filters, we use the 'next-header' in the from statements. Using this configuration statement, we can match the IPv6 extension headers that follow the IPv6 header. The most common extension headers are TCP and UDP.

Since OSPF uses neither of those, we specify 'next-header ospf':

```
root@srx1#set firewall family inet6 filter protect-re-ipv6 term ospf from next-header ospf

root@srx1#set firewall family inet6 filter protect-re-ipv6 term ospf then accept
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

For BGP, we first configure a term that allows BGP response messages to TCP destination-port 179:

```
root@srx1#set firewall family inet6 filter protect-re-ipv6 term bgp-source from next-header tcp

root@srx1#set firewall family inet6 filter protect-re-ipv6 term bgp-source from source-port bgp

root@srx1#set firewall family inet6 filter protect-re-ipv6 term bgp-source then accept
```

Then, we configure another term that allows inbound BGP messages to the local device on destination-port 179:

```
root@srx1#set firewall family inet6 filter protect-re-ipv6 term bgp-destination from next-header tcp

root@srx1#set firewall family inet6 filter protect-re-ipv6 term bgp-destination from destination-port bgp

root@srx1#set firewall family inet6 filter protect-re-ipv6 term bgp-destination then accept
```

11) Configure the SRX devices to accept IPv6 ICMP and traceroute messages. Ensure that the flow of messages is limited using a policer.

We can use the same policer we configured for IPv4. The configuration of that policer was as follows:

```
root@srx1#set firewall policer re-policer if-exceeding bandwidth-limit 100k

root@srx1#set firewall policer re-policer if-exceeding burst-size-limit 25k

root@srx1#set firewall policer re-policer then discard
```

```
root@srx1#set firewall family inet6 filter protect-re-ipv6 term icmp from next-header icmp6

root@srx1#set firewall family inet6 filter protect-re-ipv6 term icmp then policer re-policer

root@srx1#set firewall family inet6 filter protect-re-ipv6 term icmp then accept
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1#set firewall family inet6 filter protect-re-ipv6 term
traceroute from next-header udp

root@srx1#set firewall family inet6 filter protect-re-ipv6 term
traceroute from destination-port 33434-33534

root@srx1#set firewall family inet6 filter protect-re-ipv6 term
traceroute then policer re-policer

root@srx1#set firewall family inet6 filter protect-re-ipv6 term
traceroute then accept
```

12) Configure the EX devices to accept IPv6 ICMP and traceroute messages.

```
root@ex1#set firewall family inet6 filter protect-re-ipv6 term icmp
from next-header icmp6

root@ex1#set firewall family inet6 filter protect-re-ipv6 term icmp
then accept

root@ex1#set firewall family inet6 filter protect-re-ipv6 term
traceroute from next-header udp

root@ex1#set firewall family inet6 filter protect-re-ipv6 term
traceroute from destination-port 33434-33534

root@ex1#set firewall family inet6 filter protect-re-ipv6 term
traceroute then accept
```

13) Configure the IPv6 firewall filter to discard any other traffic and count the packets.

```
root@srx1#set firewall family inet6 filter protect-re-ipv6 term last
then count dropped-ipv6

root@srx1#set firewall family inet6 filter protect-re-ipv6 term last
then discard
```

14) Applying the IPv6 filter to the Loopback 0 interface, protect the control plane from malicious IPv6 traffic.

```
root@srx1#set interfaces lo0 unit 0 family inet6 filter input protect-
re-ipv6
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Note: you can display the named counters for the IPv4 as well as the IPv6 firewall filter you configured. You can do this by issuing the following commands:

```
root@srx1#run show firewall counter dropped-ipv4 filter protect-re-ipv4
```

```
root@srx1#run show firewall counter dropped-ipv6 filter protect-re-ipv6
```

15) Ensure that all the configuration for the firewall stanza is not displayed when users issue a 'show configuration' command.

We can use the 'apply-flags' configuration command for this. The command is useful to hide parts of the configuration, increasing readability.

On all devices:

```
root@srx1#set firewall family inet filter protect-re-ipv4 apply-flags omit
```

```
root@srx1#set firewall family inet6 filter protect-re-ipv6 apply-flags omit
```

On the srx devices, we also configured a policer for the ICMP and traceroute term.

On SRX devices only:

```
root@srx1#set firewall policer re-policer apply-flags omit
```

```
root@srx1#set firewall policer tcp-flood-limit apply-flags omit
```

Note: in case you want to what configuration is omitted when examining the configuration on the device, use the '**show firewall | display omit**' or '**show firewall | display set**'.

Task Verification:

A common mistake made in the JNCIE lab exam is an incorrect firewall filter that breaks something. This might cause you to lose many points.

The firewall filter terms should be checked one by one. After applying the firewall filters to the loopback interfaces, it is important to verify that all exam protocols and features are still working. For example the **show ospf neighbor**, **show bgp neighbor**, **show vrrp, show pim neighbors** and **show ntp association** can be used to verify the related terms. Use ping, telnet and SSH from other devices to check connectivity for those protocols.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

You can also contemplate activating the 'log' option for the last term, making the device log the packet headers of all discarded packets. By issuing a '**show firewall log**', you will be able to see if the filter is discarding any traffic that it is not supposed to. If you do so, and if the assignment did not specify that the last term should include the 'log' option, do not forget to remove it afterwards.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Chapter 2: L2 Switching and Security

A JNCIE candidate should be familiar with the following topics:

- Spanning-tree protocol
- VLANs
- Virtual chassis
- Layer 2 security features

Overview.

JUNOS runs on a whole variety of EX switches. These switches can be configured to suit todays Enterprise requirements. In addition to this, you will find that a lot of the layer 2 protocols and technologies will also be configurable on other platforms as well.

Spanning-tree protocol

RSTP is the default spanning-tree protocol that is enabled for EX Series Switches. RSTP is not VLAN aware and runs as a single instance. To be able to load balance traffic across multiple links, we should use either VSTP or MSTP.

In VSTP, a BPDU is send for each VLAN on every port. This can potentially overload the CPU. VSTP is limited to 253 different spanning tree topologies. VSTP and RSTP can be configured at the same time.

Using MSTP, you can map multiple VLANs to a spanning-tree instance and you can have multiple spanning-tree instances. MSTP is considered the most scalable of the different xSTP solutions.

To configure MSTP consider the following steps:

1. Delete **RSTP**
2. Define a **configuration name**
3. Define a **revision** number
4. Define a MSTI-to-VLAN mapping
5. Assign a **bridge-priority** (smaller priority is preferred)
    - To configure a port as edge you can put the **edge** command under the interface in the xSTP configuration.
    - To enable BPDU blocking use **bpdu-block-on-edge** command.
    - To change a switch port role you can use **priority** and **cost** commands.

To verify spanning-tree configuration the below commands are useful:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

show spanning-tree interface

show spanning-tree bridge

show spanning-tree mstp configuration

VLANs

- VLAN switching and trunking:

To configure an access port, you should enter **port-mode access** under the **family ethernet-switching**.

To configure a trunk port you should enter **port-mode trunk** under **family ethernet-switching**.

There are two different ways to assign a VLAN to an interface. You can create a vlan and assign it as a **member** at the interface level or you can assign an **interface** directly to a named VLAN.

To transmit untagged traffic on a trunk interface, configure the interface with the **native-vlan-id** statement.

- **Q-in-Q**:

To configure a Q-in-Q tunneling, consider the following steps:

1. on both ingress sides of the Q-in-Q tunnel:

set vlans *q-in-q* vlan-id 500

set vlans *q-in-q* dot1-tunneling customer-vlan 250

In the configuration above, a VLAN named 'q-in-q' is created. In this VLAN, VLAN 500 is the S-VLAN and VLAN 250 is C-VLAN.

set ethernet-switching-options dot1q-tunneling ether-type 0x8100

This configuration sets the Ether-type of tunneled VLANs to a value that equals 'normal', or non-tunneling, VLANS. This way, regular VLANs and tunneled VLANs can use the same trunk ports.

set interface ge-0/0/0 unit 0 family ethernet-switching vlan members *q-in-q*

2. On the intermediate devices:

set vlans *q-in-q* vlan-id 500

set vlans *q-in-q* dot1-tunneling

set ethernet-switching-options dot1q-tunneling ether-type 0x8100

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Voice VLAN

To send both tagged (voice) and untagged (data) traffic through a port while using separate VLANs, you can use the following command:

set ethernet-switching-options voip interface ge-0/0/0 vlan *voice*

You can issue a **show ethernet-switching interfaces** to verify the tagging and blocking status.

Private VLAN

PVLAN allows you to divide a broadcast domain to multiple subdomains.

A PVLAN consists of a primary VLAN and 1 or more secondary VLANs. The secondary VLANs can be either a community VLANs or an isolated VLANs. There can be multiple community VLANs but only 1 isolated VLAN.

Interfaces assigned to an isolated VLAN will only be able to forward packets to a promiscuous port or to the PVLAN trunk port.

Interfaces assigned to a community VLAN will be able to forward packets to a promiscuous port or the PVLAN trunk port. In addition to this, these interfaces will also be able to forward packets to other interfaces that are assigned the same community VLAN.

The Primary VLAN always includes a 802.1Q tag. Secondary vlans are not required to include a 802.1Q tag unless the PVLAN spans multiple switches.

To enable a VLAN for PVLAN functionality, you should use the **no-local-switching** command under the primary VLAN configuration.

To link a primary VLAN to its community VLANs, use the **primary-vlan** command.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Virtual Chassis

When a virtual chassis is configured as preprovisioned, new members are not able to join the virtual chassis automatically. The configuration of a preprovisioned virtual chassis requires the following steps:

1) **run show chassis hardware | match "routing engine"**
   - Issued on both switches**,** used to find the RE serial number.
2) **edit virtual-chassis**
3) **set preprovisioned**
4) **set member 0 role routing-engine**
5) **set member 0 serial-number *serial number of switch # 1***
6) **set member 1 role line-card**
7) **set member 1 serial-number *serial number of switch # 2***
8) **rename me0 to vme**
9) **setting the VC ports:**

On Switch #1:

request virtual-chassis vc-port set interface vcp-0

request virtual-chassis vc-port set interface vcp-1

On Switch #2:

request virtual-chassis vc-port set interface vcp-0

request virtual-chassis vc-port set interface vcp-1 member 1 (on the master switch)

If a switch had any configuration prior to the formation of the VC, and the member number changed during this formation, do not forget to change the interface configuration as well.

Verification commands to check the status of a Virtual-Chassis:

show virtual-chassis status

show virtual-chassis vc-port

show interface terse

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

The steps to deconstruct a VC are as follows:

- Disable the VC ports with the following commands on both switches:

request virtual-chassis vc-port set interface vcp-0 disable

request virtual-chassis vc-port set interface vcp-1 disable

- On both devices, recycle the opposite member-ids, because they should be their own master.

request virtual-chassis recycle member-id 0

request virtual-chassis recycle member-id 1

- Verify if that there are no other members.

show virtual-chassis

To ensure that the second switch is renumbered to the correct member-id, issue the following command:

request virtual-chassis renumber member-id 1 new-member-id 0

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Layer 2 security features

Port security (802.1X and MAC RADIUS)

To configure 802.1X port security, take the following into consideration:

- There are three roles in 802.1X:
  - Supplicant (user PC)
  - Authenticator (switch)
  - Authentication server (Radius server)
- Supplicant modes are:
  - **single**: only the first user which is connected to port is authenticated and the other users do not need to authenticate
  - **single-secure**: only one supplicant can connect to the port
  - **multiple**: multiple supplicants can connect to the port and the authentication is done individually
- MAC RADIUS is not another way of authentication a supplicant. It can be used when a supplicant does not have an 802.1X client. MAC RADIUS has a similar configuration to 802.1X.

To configure 802.1X, follow these steps:

- Define the radius server, configure a secret password and specify the source IP address **[edit access].**
- Define a profile and specify the authentication order for this profile.
- Configure the authenticator with the profile and define the port **[edit protocols dot1x]**

MAC RADIUS is configured in the same hierarchy. Note that the **mac-radius restrict** option can be used to enforce the use of MAC RADIUS only. This option will have the switch drop all 802.1X packets and the delay that occurs when a switch has to determine whether a device is a non-802.1X is also skipped.

The **run show dot1x interface** and **run show dot1x interface *interface-name* detail** commands are useful to verify the configuration.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Firewall Filters

Please keep the following in mind when you are configuring L2 firewall filters:

- Firewall rules are stateless
- The default action is to discard traffic
- Adding an entry to an existing firewall filter someplace other than at the end can be done using the **insert** command
- Egress firewall have no effect on traffic that is generated by the local route engine
- Configure port- or VLAN-based L2 firewall filters on EX switches under the **family ethernet-switching** hierarchy

MAC table Filtering:

- Configured under the **[edit ethernet-switching-options]** hierarchy
- The following MAC limiting violations can be specified:
  - none
  - drop
  - log
  - shutdown
- Verification can be done with the **run show log messages** command.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 1: L2 Switching Network Deployment

1) Load the baseline configuration for all devices using the following command:

```
root@srx1# load override terminal
```

Use ctrl+D at the end then commit configuration.

You need to reboot SRX devices to change it to packet mode.

2) Load baseline configuration of the VR device using the same command (access the VR with username ***root*** and password ***root123***).

3) Configure the following interfaces:

```
root@srx1# set interfaces ge-0/0/3 description "R1 connection"
root@srx1# set interfaces ge-0/0/3 unit 0 family inet address
192.168.0.2/30
root@srx1# set interfaces ge-0/0/6 description "EX1 connection No.1"
root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
port-mode trunk
root@srx1# set interfaces ge-0/0/9 description "EX1 connection No.2"
root@srx1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
port-mode trunk
root@srx1# set interfaces ge-0/0/10 description "EX2 connection"
root@srx1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
port-mode trunk
root@srx1# set chassis aggregated-devices ethernet device-count 1
root@srx1# set interfaces ge-0/0/1 description "SRX2 connection No.1"
root@srx1# set interfaces ge-0/0/1 gigether-options 802.3ad ae0
root@srx1# set interfaces ge-0/0/2 description "SRX2 connection No.2"
root@srx1# set interfaces ge-0/0/2 gigether-options 802.3ad ae0
root@srx1# set interfaces ae0 description "SRX2 aggregated connection"
root@srx1# set interfaces ae0 aggregated-ether-options lacp active
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set chassis aggregated-devices ethernet device-count 1
set interfaces ge-0/0/1 description "SRX1 connection No.1"
root@srx2# set interfaces ge-0/0/1 gigether-options 802.3ad ae0
root@srx2# set interfaces ge-0/0/2 description "SRX1 connection No.2"
root@srx2# set interfaces ge-0/0/2 gigether-options 802.3ad ae0
root@srx2# set interfaces ae0 description "SRX1 aggregated connection"
root@srx2# set interfaces ae0 aggregated-ether-options lacp active
```

Verification:

```
root@srx1# run show lacp interfaces
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout
Activity
      ge-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast
Active
      ge-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast
Active
      ge-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast
Active
      ge-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast
Active
    LACP protocol:        Receive State  Transmit State          Mux
State
      ge-0/0/1                  Current   Fast periodic Collecting
distributing
      ge-0/0/2                  Current   Fast periodic Collecting
distributing
```

Note: other network interfaces are configured in a similar way.

4) Configure OAM LFM on the links between ex1 and ex3, ex2 and ex4. Make sure ex1 and ex2 initiate the discovery process. Ensure OAM PDUs are sent every 600ms. If 5 PDUs are missed, the peer should be declared lost/down/

When configuring Ethernet OAM, use the following steps:

1) Enabling LFM (Link Fault Management) support on the interface
2) Configure the link discovery mode (active or passive)
3) Specify the PDU-sending interval

```
root@ex1# set protocols oam ethernet link-fault-management interface
ge-0/0/18.0 pdu-interval 600
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set protocols oam ethernet link-fault-management interface
ge-0/0/18.0 link-discovery active

root@ex1# set protocols oam ethernet link-fault-management interface
ge-0/0/18.0 pdu-threshold 5
```

Verification:

```
root@ex1# run show oam ethernet link-fault-management brief
  Interface: ge-0/0/18.0
    Status: Running, Discovery state: Send Any
    Peer address: 2c:6b:f5:35:1b:95
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: passive, Unidirectional mode: unsupported
      Remote loopback mode: unsupported, Link events: supported
      Variable requests: unsupported
```

Note: The other link between ex2 and ex4 is configured in similar way.

5) Enable BUM storm control on ex1 through ex4:

Storm control is configured under the [**ethernet-switching-options**] stanza:

```
root@ex1# set ethernet-switching-options storm-control interface all bandwidth 10000

root@ex1# set ethernet-switching-options storm-control interface all
multicast
```

To ensure that the interface will be temporarily disabled as soon as storm-control is triggered, we configure the following:

```
root@ex1# set ethernet-switching-options port-error-disable disable-
timeout 30
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 2: Virtual Chassis

1) Merge ex3 and ex4 into a Virtual Chassis. Ensure that ex3 becomes the master RE with member ID 0:

To meet the requirements of this task, we configure EX3 as member 0 and EX4 as member 1. The EX3 switch is configured with a higher priority and after applying the configuration, we issue a reboot on both switches:

```
root@ex3# set virtual-chassis member 0 mastership-priority 255

root@ex3# run request system reboot

root@ex4# set virtual-chassis member 1 mastership-priority 200

root@ex4# run request system reboot
```

2) Restore the interface configuration on the ex4 switch:

Reconfigure the hostname and restore the interfaces on the ex4 switch. Do not forget the protocol OAM configuration either:

```
root@ex3_ex4# set system host-name ex3_ex4

root@ex3_ex4# set interfaces ge-1/0/3 unit 0 family ethernet-switching

root@ex3_ex4# set interfaces ge-1/0/6 description "H12 connection"

root@ex3_ex4# set interfaces ge-1/0/6 unit 0 family ethernet-switching
port-mode access

root@ex3_ex4# set interfaces ge-1/0/7 description "H11 connection"

root@ex3_ex4# set interfaces ge-1/0/7 unit 0 family ethernet-switching
port-mode access

root@ex3_ex4# set interfaces ge-1/0/8 description "H10 connection"

root@ex3_ex4 set interfaces ge-1/0/8 unit 0 family ethernet-switching
port-mode access

root@ex3_ex4# set interfaces ge-1/0/19 description "EX2 connection"

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching port-mode trunk
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set protocols oam ethernet link-fault-management
interface ge-1/0/19.0 pdu-interval 600

root@ex3_ex4# set protocols oam ethernet link-fault-management
interface ge-1/0/19.0 link-discovery passive

root@ex3_ex4# set protocols oam ethernet link-fault-management
interface ge-1/0/19.0 pdu-threshold 5
```

3) Configure the vme.0 interface as management interface for the VC:

We simply rename the management interface from **me** to **vme**:

```
root@ex3_ex4#rename interfaces me0 to vme
```

Verification:

```
root@ex3_ex4# run show virtual-chassis

Virtual Chassis ID: aa76.8420.f19a
Virtual Chassis Mode: Enabled
                                          Mstr           Mixed
Neighbor List
Member ID  Status   Serial No    Model     prio  Role      Mode ID
Interface
0 (FPC 0)  Prsnt    BM0210045284 ex4200-24t 255  Master*     N  1
vcp-0
1 (FPC 1)  Prsnt    BM0210359941 ex4200-24t 200  Backup      N  0
vcp-1

Member ID for next new member: 2 (FPC 2)

{master:0}[edit]
root@ex3_ex4# run show interfaces terse
Interface               Admin Link Proto    Local
Remote
ge-0/0/0                up    down
ge-0/0/0.0              up    down eth-switch
ge-0/0/1                up    down
ge-0/0/1.0              up    down eth-switch
ge-0/0/2                up    down
ge-0/0/2.0              up    down eth-switch
ge-0/0/3                up    down
ge-0/0/3.0              up    down eth-switch
ge-0/0/4                up    down
ge-0/0/4.0              up    down eth-switch
ge-0/0/5                up    down
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
ge-0/0/6                up    up
ge-0/0/6.0              up    up   eth-switch
ge-0/0/7                up    up
ge-0/0/7.0              up    up   eth-switch
ge-0/0/8                up    up
ge-0/0/8.0              up    up   eth-switch
ge-0/0/9                up    up
ge-0/0/10               up    up
ge-0/0/11               up    up
ge-0/0/12               up    up
ge-0/0/13               up    up
ge-0/0/14               up    up
ge-0/0/15               up    up
ge-0/0/16               up    down
ge-0/0/17               up    down
ge-0/0/18               up    up
ge-0/0/18.0             up    up   eth-switch
ge-0/0/19               up    down
ge-0/0/20               up    down
ge-0/0/21               up    down
ge-0/0/22               up    up
ge-0/0/23               up    up
ge-1/0/0                up    down
ge-1/0/0.0              up    down eth-switch
ge-1/0/1                up    down
ge-1/0/1.0              up    down eth-switch
ge-1/0/2                up    down
ge-1/0/2.0              up    down eth-switch
ge-1/0/3                up    down
ge-1/0/3.0              up    down eth-switch
ge-1/0/4                up    down
ge-1/0/4.0              up    down eth-switch
ge-1/0/5                up    down
ge-1/0/6                up    up
ge-1/0/6.0              up    up   eth-switch
ge-1/0/7                up    up
ge-1/0/7.0              up    up   eth-switch
ge-1/0/8                up    up
ge-1/0/8.0              up    up   eth-switch
ge-1/0/9                up    up
ge-1/0/10               up    up
ge-1/0/11               up    up
ge-1/0/12               up    up
ge-1/0/13               up    up
ge-1/0/14               up    up
ge-1/0/15               up    up
ge-1/0/16               up    down
ge-1/0/17               up    down
ge-1/0/18               up    down
ge-1/0/19               up    up
ge-1/0/19.0             up    up   eth-switch
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
ge-1/0/20                up    down
ge-1/0/21                up    down
ge-1/0/22                up    up
ge-1/0/23                up    up
vcp-0                    up    up
vcp-0.32768              up    up
vcp-1                    up    down
vcp-1.32768              up    down
bme0                     up    up
bme0.32768               up    up   inet     128.0.0.1/2
                                             128.0.0.16/2
                                             128.0.0.32/2
                                    tnp      0x10
bme0.32770               down  up   eth-switch
dsc                      up    up
gre                      up    up
ipip                     up    up
jsrv                     up    up
jsrv.1                   up    up   inet     128.0.0.127/2
lo0                      up    up
lo0.0                    up    up   inet
                                    inet6    fe80::2e6b:f50f:fc35:1b80
lsi                      up    up
me0                      up    up
me0.0                    up    up   eth-switch
mtun                     up    up
pimd                     up    up
pime                     up    up
tap                      up    up
vlan                     up    up
vme                      up    up
vme.0                    up    up   inet     10.10.1.13/24
                                    inet6    fdb3:309:419b:9ec4::13/64

fe80::2e6b:f5ff:fe35:1b82/64


{master:0}[edit]
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 3: VLAN Configuration

1) The configuration of the VLANs:

```
root@srx1# set vlans Kappa vlan-id 600

root@srx1# set vlans Omega vlan-id 400

root@srx1# set vlans Secured vlan-id 500

root@srx1# set vlans Sigma vlan-id 700

root@srx1# set vlans Theta vlan-id 200

root@srx1# set vlans Voice vlan-id 601

root@srx1# set vlans Zeta vlan-id 30

root@srx1# set vlans default vlan-id 1

root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 200

root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 300

root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 400

root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 500

root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 600

root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 601

root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 701

root@srx1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 200

root@srx1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 400

root@srx1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 500
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 600

root@srx1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 601

root@srx1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 701

root@srx1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 200

root@srx1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 400

root@srx1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 500

root@srx1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 600

root@srx1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 601

root@srx1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 200

root@srx1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 400

root@srx1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 500

root@srx1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 600

root@srx1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 601

root@srx1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 701
```

```
root@srx1# run show vlans
Name            Tag     Interfaces
Kappa           600
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
Omega           400
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
Secured         500
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Sigma           700
                        None
Sigma-transport 701
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*
Theta           200
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
Voice           601
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
Zeta            300
                        ge-0/0/6.0*
default         1
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
```

```
root@srx2# set vlans Kappa vlan-id 600

root@srx2# set vlans Omega vlan-id 400

root@srx2# set vlans Secured vlan-id 500

root@srx2# set vlans Sigma vlan-id 700

root@srx2# set vlans Theta vlan-id 200

root@srx2# set vlans Voice vlan-id 601

root@srx2# set vlans Zeta vlan-id 300

root@srx2# set vlans default vlan-id 1
```

```
root@srx2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 200

root@srx2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 300

root@srx2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 400

root@srx2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 500

root@srx2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 600

root@srx2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 601
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 701

root@srx2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 200

root@srx2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 400

root@srx2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 500

root@srx2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 600

root@srx2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 601

root@srx2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 701

root@srx2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 200

root@srx2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 400

root@srx2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 500

root@srx2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 600

root@srx2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 601


root@srx2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 200

root@srx2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 400

root@srx2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 500

root@srx2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 600
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 601

root@srx2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 701
```

```
root@srx2# run show vlans
Name            Tag     Interfaces
Kappa           600
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
Omega           400
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
Secured         500
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
Sigma           700
                        None
Sigma-transport 701
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*
Theta           200
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
Voice           601
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
Zeta            300
                        ge-0/0/6.0*
default         1
                        ae0.0*, ge-0/0/6.0*, ge-0/0/9.0*, ge-0/0/10.0*
```

```
root@ex1# set vlans Kappa vlan-id 600

root@ex1# set vlans Omega vlan-id 400

root@ex1# set vlans Secured vlan-id 500

root@ex1# set vlans Sigma-transport vlan-id 701

root@ex1# set vlans Sigma-transport interface ge-0/0/8.0

root@ex1# set vlans Theta vlan-id 200

root@ex1# set vlans Theta interface ge-0/0/7.0

root@ex1# set vlans Voice vlan-id 601

root@ex1# set vlans Zeta vlan-id 300

root@ex1# set vlans default vlan-id 1
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 200

root@ex1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 300

root@ex1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 400

root@ex1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 500

root@ex1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 600

root@ex1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 601

root@ex1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 701

root@ex1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 200

root@ex1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 400

root@ex1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 500

root@ex1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 600

root@ex1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 601

root@ex1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 701

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 200

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 400

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 500
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 600

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 601

root@ex1# set interfaces ge-0/0/16 unit 0 family ethernet-switching
vlan members 300

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members 200

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members 400

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members 500

root@ex1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 200

root@ex1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 300

root@ex1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 400

root@ex1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 500

root@ex1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 600

root@ex1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 601
```

```
root@ex1# run show vlans brief
Name            Tag
Kappa           600
Omega           400
Secured         500
Sigma-transport 701
Theta           200
Theta-community 2002
Voice           601
Zeta            300
__pvlan_Theta_ge-0/0/7.0__
__pvlan_Theta_isiv__ 2001
default         1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set vlans Kappa vlan-id 600

root@ex2# set vlans Omega vlan-id 400

root@ex2# set vlans Omega interface ge-0/0/7.0

root@ex2# set vlans Secured vlan-id 500

root@ex2# set vlans Sigma-transport vlan-id 701

root@ex2# set vlans Theta vlan-id 200

root@ex2# set vlans Voice vlan-id 601

root@ex2# set vlans Zeta vlan-id 300

root@ex2# set vlans default vlan-id 1


root@ex2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 200

root@ex2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 300

root@ex2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 400

root@ex2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 500

root@ex2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 600

root@ex2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 601

root@ex2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 701

root@ex2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 200

root@ex2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 400

root@ex2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 500
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 600

root@ex2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 601

root@ex2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 701

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 200

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 400

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 500

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 600

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 601

root@ex2# set interfaces ge-0/0/17 unit 0 family ethernet-switching
vlan members 300

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
vlan members 200

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
vlan members 400

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
vlan members 500


root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 200

root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 300

root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 400

root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 500
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 600

root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 601
```

```
root@ex2# run show vlans brief

Name            Tag
Kappa           600
Omega           400
Secured         500
Sigma-transport 701
Theta           200
Theta-community 2002
Voice           601
Zeta            300
__pvlan_Theta_isiv__ 2001
default         1
```

```
root@srx3# set vlans Kappa vlan-id 600

root@srx3# set vlans Omega vlan-id 400

root@srx3# set vlans Secured vlan-id 500

root@srx3# set vlans Sigma vlan-id 700

root@srx3# set vlans Theta vlan-id 200

root@srx3# set vlans Voice vlan-id 601

root@srx3# set vlans Zeta vlan-id 300

root@srx3# set vlans Zeta interface ge-0/0/14.0

root@srx3# set vlans default vlan-id 1


root@srx3# set interfaces ge-0/0/1 unit 0 family ethernet-switching
vlan members all

root@srx3# set interfaces ge-0/0/2 unit 0 family ethernet-switching
vlan members all

root@srx3# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 300
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# run show vlans
Name            Tag     Interfaces
Kappa           600
                        ge-0/0/1.0*, ge-0/0/2.0*
Omega           400
                        ge-0/0/1.0*, ge-0/0/2.0*
Secured         500
                        ge-0/0/1.0*, ge-0/0/2.0*
Sigma           700
                        ge-0/0/1.0*, ge-0/0/2.0*
Theta           200
                        ge-0/0/1.0*, ge-0/0/2.0*
Voice           601
                        ge-0/0/1.0*, ge-0/0/2.0*
Zeta            300
                        ge-0/0/1.0*, ge-0/0/2.0*, ge-0/0/12.0*
default         1
                        ge-0/0/1.0*, ge-0/0/2.0*
```

```
root@srx4# set vlans Kappa vlan-id 600

root@srx4# set vlans Omega vlan-id 400

root@srx4# set vlans Secured vlan-id 500

root@srx4# set vlans Sigma vlan-id 700

root@srx4# set vlans Theta vlan-id 200

root@srx4# set vlans Voice vlan-id 601

root@srx4# set vlans Zeta vlan-id 300

root@srx4# set vlans Zeta interface ge-0/0/14.0

root@srx4# set vlans default vlan-id 1


root@srx4# set interfaces ge-0/0/1 unit 0 family ethernet-switching
vlan members all

root@srx4# set interfaces ge-0/0/2 unit 0 family ethernet-switching
vlan members all

root@srx4# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 300
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx4# run show vlans
Name            Tag     Interfaces
Kappa           600
                        ge-0/0/1.0*, ge-0/0/2.0*
Omega           400
                        ge-0/0/1.0*, ge-0/0/2.0*
Secured         500
                        ge-0/0/1.0*, ge-0/0/2.0*
Sigma           700
                        ge-0/0/1.0*, ge-0/0/2.0*
Theta           200
                        ge-0/0/1.0*, ge-0/0/2.0*
Voice           601
                        ge-0/0/1.0*, ge-0/0/2.0*
Zeta            300
                        ge-0/0/1.0*, ge-0/0/2.0*, ge-0/0/13.0*
default         1
                        ge-0/0/1.0*, ge-0/0/2.0*
```

```
root@ex3_ex4# set vlans Kappa vlan-id 600

root@ex3_ex4# set vlans Omega vlan-id 400

root@ex3_ex4# set vlans Omega interface ge-1/0/8.0

root@ex3_ex4# set vlans Omega interface ge-1/0/6.0

root@ex3_ex4# set vlans Secured vlan-id 500

root@ex3_ex4# set vlans Sigma vlan-id 700

root@ex3_ex4# set vlans Theta vlan-id 200

root@ex3_ex4# set vlans Theta interface ge-0/0/6.0

root@ex3_ex4# set vlans Theta interface ge-0/0/8.0

root@ex3_ex4# set vlans Voice vlan-id 601

root@ex3_ex4# set vlans Zeta vlan-id 300

root@ex3_ex4# set vlans default vlan-id 1


root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching vlan members 200

root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching vlan members 400
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching vlan members 500

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching vlan members 200

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching vlan members 400

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching vlan members 500
```

2) Ensure that H3, H7 and H9 cannot communicate with the other hosts in VLAN 200. You can use VLAN2001 to accomplish this requirement

This task is completed using the private VLAN feature. We configure the private VLAN using the **no-local-switching** command. Because the private VLAN spans multiple switches, additional configuration is needed. In the primary private VLAN configuration, the links that connect the switches together need to be specified using the **pvlan-trunk** command.

We need to specify an isolation-ID and we also need to configure a secondary isolated-VLAN that will require its own 802.1Q tag. This isolated VLAN also needs to be configured on all of the switches.

On ex1:

```
root@ex1# set vlans Theta interface ge-0/0/18.0 pvlan-trunk

root@ex1# set vlans Theta interface ae0.0 pvlan-trunk

root@ex1# set vlans Theta no-local-switching

root@ex1# set vlans Theta isolation-id 2001
```

```
root@ex1# run show vlans extensive
VLAN: Theta, Created at: Sat Dec 20 16:30:57 2014
802.1Q Tag: 200, Internal index: 7, Admin State: Enabled, Origin:
Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 5 (Active = 5), Untagged  1 (Active = 1)
      ae0.0*, tagged, trunk, pvlan-trunk
      ge-0/0/6.0*, tagged, trunk
      ge-0/0/9.0*, tagged, trunk
      ge-0/0/10.0*, tagged, trunk
      ge-0/0/18.0*, tagged, trunk, pvlan-trunk
      ge-0/0/7.0*, untagged, access
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Secondary VLANs: Isolated 1, Community  0, Inter-switch-isolated  1
  Isolated VLANs :
      __pvlan_Theta_ge-0/0/7.0__
  Inter-switch-isolated VLAN :
      __pvlan_Theta_isiv__
```

On ex2:

```
root@ex2# set vlans Theta interface ge-0/0/19.0 pvlan-trunk

root@ex2# set vlans Theta interface ae0.0 pvlan-trunk

root@ex2# set vlans Theta isolation-id 2001
```

```
root@ex2# run show vlans extensive
VLAN: Theta, Created at: Sat Dec 20 16:33:38 2014
802.1Q Tag: 200, Internal index: 7, Admin State: Enabled, Origin:
Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 5 (Active = 5), Untagged  0 (Active = 0)
      ae0.0*, tagged, trunk, pvlan-trunk
      ge-0/0/6.0*, tagged, trunk
      ge-0/0/9.0*, tagged, trunk
      ge-0/0/10.0*, tagged, trunk
      ge-0/0/19.0*, tagged, trunk, pvlan-trunk
Secondary VLANs: Isolated 0, Community  0, Inter-switch-isolated  1
```

On ex3_ex4:

```
root@ex3_ex4# set vlans Theta interface ge-0/0/18.0 pvlan-trunk

root@ex3_ex4# set vlans Theta interface ge-1/0/19.0 pvlan-trunk

root@ex3_ex4# set vlans Theta no-local-switching

root@ex3_ex4# set vlans Theta isolation-id 2001
```

```
root@ex3_ex4# run show vlans extensive
VLAN: Theta, Created at: Sat Dec 20 16:42:43 2014
802.1Q Tag: 200, Internal index: 7, Admin State: Enabled, Origin:
Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged  4 (Active = 4)
      ge-0/0/18.0*, tagged, trunk, pvlan-trunk
      ge-1/0/19.0*, tagged, trunk, pvlan-trunk
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
      ge-0/0/6.0*, untagged, access
      ge-0/0/7.0*, untagged, access
      ge-0/0/8.0*, untagged, access
      ge-1/0/7.0*, untagged, access
Secondary VLANs: Isolated 2, Community  0, Inter-switch-isolated  1
  Isolated VLANs :
      __pvlan_Theta_ge-0/0/6.0__
      __pvlan_Theta_ge-0/0/8.0__
    Inter-switch-isolated VLAN :
      __pvlan_Theta_isiv__
```

3) Communication between the two hosts is accomplished through the configuration of the community VLAN, using VLAN-id 2002.

On ex3_ex4:

```
root@ex3_ex4# set vlans Theta-community vlan-id 2002

root@ex3_ex4# set vlans Theta-community interface ge-0/0/7.0

root@ex3_ex4# set vlans Theta-community interface ge-1/0/7.0

root@ex3_ex4# set vlans Theta-community primary-vlan Theta
```

```
root@ex3_ex4# run show vlans extensive
VLAN: Theta, Created at: Sat Dec 20 16:42:43 2014
802.1Q Tag: 200, Internal index: 7, Admin State: Enabled, Origin:
Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged  4 (Active = 4)
      ge-0/0/18.0*, tagged, trunk, pvlan-trunk
      ge-1/0/19.0*, tagged, trunk, pvlan-trunk
      ge-0/0/6.0*, untagged, access
      ge-0/0/7.0*, untagged, access
      ge-0/0/8.0*, untagged, access
      ge-1/0/7.0*, untagged, access
Secondary VLANs: Isolated 2, Community  1, Inter-switch-isolated  1
  Isolated VLANs :
      __pvlan_Theta_ge-0/0/6.0__
      __pvlan_Theta_ge-0/0/8.0__
  Community VLANs :
      Theta-community
  Inter-switch-isolated VLAN :
      __pvlan_Theta_isiv__
VLAN: Theta-community, Created at: Sat Dec 20 16:42:43 2014
802.1Q Tag: 2002, Internal index: 8, Admin State: Enabled, Origin:
Static
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Private VLAN Mode: Community, Primary VLAN: Theta
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged  2 (Active = 2)
      ge-0/0/18.0*, tagged, trunk, pvlan-trunk
      ge-1/0/19.0*, tagged, trunk, pvlan-trunk
      ge-0/0/7.0*, untagged, access
      ge-1/0/7.0*, untagged, access
```

Note: The community VLAN configuration is also applied to ex1 and ex2

4) A native VLAN should be configured on all devices. It needs to be transported untagged on all trunk interfaces:

```
root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
native-vlan-id 1

root@srx1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
native-vlan-id 1

root@srx1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
native-vlan-id 1

root@srx1# set interfaces ae0 unit 0 family ethernet-switching native-
vlan-id 1
```

```
root@srx1# run show vlans extensive
VLAN: default, Created at: Sat Dec 20 16:31:38 2014
802.1Q Tag: 1, Internal index: 10, Admin State: Enabled, Origin:
Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged  4 (Active = 4)
      ae0.0*, untagged, trunk
      ge-0/0/6.0*, untagged, trunk
      ge-0/0/9.0*, untagged, trunk
      ge-0/0/10.0*, untagged, trunk
```

5) The voice VLAN feature needs to be configured on the ex1 and the ex2 switch. We need to specify that the tagged (voice) and untagged (data) traffic can pass through the following ports:

```
root@ex1# set vlans Kappa interface ge-0/0/0.0

root@ex1# set vlans Kappa interface ge-0/0/1.0

root@ex1# set vlans Kappa interface ge-0/0/2.0

root@ex1# set vlans Kappa interface ge-0/0/3.0

root@ex1# set vlans Kappa interface ge-0/0/4.0

root@ex1# set ethernet-switching-options voip interface ge-0/0/0.0
vlan Voice
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set ethernet-switching-options voip interface ge-0/0/1.0
vlan Voice

root@ex1# set ethernet-switching-options voip interface ge-0/0/2.0
vlan Voice

root@ex1# set ethernet-switching-options voip interface ge-0/0/3.0
vlan Voice

root@ex1# set ethernet-switching-options voip interface ge-0/0/4.0
vlan Voice
```

```
root@ex1# run show vlans extensive
VLAN: Voice, Created at: Sat Dec 20 16:30:57 2014
802.1Q Tag: 601, Internal index: 9, Admin State: Enabled, Origin:
Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 9 (Active = 4), Untagged  0 (Active = 0)
      ae0.0*, tagged, trunk
      ge-0/0/0.0, tagged, access
      ge-0/0/1.0, tagged, access
      ge-0/0/2.0, tagged, access
      ge-0/0/3.0, tagged, access
      ge-0/0/4.0, tagged, access
      ge-0/0/6.0*, tagged, trunk
      ge-0/0/9.0*, tagged, trunk
      ge-0/0/10.0*, tagged, trunk
```

6) To support power over Ethernet for IP phones and to enable the Link Layer Discovery Protocol for VOIP deployments, we configure the following on both the ex1 and the ex2 switch:

```
root@ex1# set poe interface ge-0/0/0

root@ex1# set poe interface ge-0/0/1

root@ex1# set poe interface ge-0/0/2

root@ex1# set poe interface ge-0/0/3

root@ex1# set poe interface ge-0/0/4

root@ex1# set protocols lldp-med interface ge-0/0/4.0

root@ex1# set protocols lldp-med interface ge-0/0/3.0

root@ex1# set protocols lldp-med interface ge-0/0/2.0

root@ex1# set protocols lldp-med interface ge-0/0/1.0

root@ex1# set protocols lldp-med interface ge-0/0/0.0
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# run show poe interface
Interface Admin status Oper status Max power Priority Power
consumption Class
 ge-0/0/0 Enabled      OFF          15.4W     Low      0.0W
not-applicable
 ge-0/0/1 Enabled      OFF          15.4W     Low      0.0W
not-applicable
 ge-0/0/2 Enabled      OFF          15.4W     Low      0.0W
not-applicable
 ge-0/0/3 Enabled      OFF          15.4W     Low      0.0W
not-applicable
 ge-0/0/4 Enabled      OFF          15.4W     Low      0.0W
not-applicable
 ge-0/0/5 Disabled     Disabled     0.0W      Low      0.0W
not-applicable
 ge-0/0/6 Disabled     Disabled     0.0W      Low      0.0W
not-applicable
 ge-0/0/7 Disabled     Disabled     0.0W      Low      0.0W
not-applicable
root@ex1# run show lldp

LLDP                     : Enabled
Advertisement interval   : 30 seconds
Transmit delay           : 2 seconds
Hold timer               : 4 seconds
Notification interval    : 0 Second(s)
Config Trap Interval     : 0 seconds
Connection Hold timer    : 300 seconds

LLDP MED                 : Enabled
MED fast start count     : 3 Packets


Interface       Parent Interface    LLDP        LLDP-MED
ge-0/0/0.0      -                   Enabled     Enabled
ge-0/0/1.0      -                   Enabled     Enabled
ge-0/0/2.0      -                   Enabled     Enabled
ge-0/0/3.0      -                   Enabled     Enabled
ge-0/0/4.0      -                   Enabled     Enabled
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 4: Spanning Tree Configuration

1) The configuration is done based on the description in the theory section of this chapter.

Keep in mind that the default bridge priority is 32k. To have a switch elected as the CST root, a lower priority should be configured under the **[edit protocols mstp]** hierarchy.

```
root@srx1# set protocols mstp configuration-name north

root@srx1# set protocols mstp revision-level 1

root@srx1# set protocols mstp bridge-priority 8k

root@srx1# set protocols mstp msti 1 bridge-priority 8k

root@srx1# set protocols mstp msti 1 vlan 200

root@srx1# set protocols mstp msti 1 vlan 600

root@srx1# set protocols mstp msti 1 vlan 601

root@srx1# set protocols mstp msti 2 vlan 400

root@srx1# set protocols mstp msti 2 vlan 500
```

```
root@srx1# run show spanning-tree mstp configuration
MSTP information
Context identifier      : 0
Region name             : north
Revision                : 1
Configuration digest    : 0xf3d7ba47b1fbb2ea735860b1e839f961


MSTI     Member VLANs
   0 0-199,201-299,301-399,401-499,501-599,602-4094
   1 200,600-601
   2 400,500
```

```
root@srx2# set protocols mstp configuration-name north

root@srx2# set protocols mstp revision-level 1

root@srx2# set protocols mstp msti 1 vlan 200

root@srx2# set protocols mstp msti 1 vlan 600

root@srx2# set protocols mstp msti 1 vlan 601
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set protocols mstp msti 2 bridge-priority 8k

root@srx2# set protocols mstp msti 2 vlan 400

root@srx2# set protocols mstp msti 2 vlan 500

root@srx2# set protocols mstp msti 2 interface ge-0/0/9.0 priority 16
```

```
root@srx2# run show spanning-tree mstp configuration
MSTP information
Context identifier      : 0
Region name             : north
Revision                : 1
Configuration digest    : 0xf3d7ba47b1fbb2ea735860b1e839f961


MSTI     Member VLANs
   0 0-199,201-299,301-399,401-499,501-599,602-4094
   1 200,600-601
   2 400,500
```

```
root@ex1# set protocols mstp configuration-name north

root@ex1# set protocols mstp revision-level 1

root@ex1# set protocols mstp msti 1 vlan 200

root@ex1# set protocols mstp msti 1 vlan 600

root@ex1# set protocols mstp msti 1 vlan 601

root@ex1# set protocols mstp msti 2 vlan 400

root@ex1# set protocols mstp msti 2 vlan 500
```

```
root@ex2# run show spanning-tree mstp configuration
MSTP information
Context identifier      : 0
Region name             : north
Revision                : 1
Configuration digest    : 0xf3d7ba47b1fbb2ea735860b1e839f961


MSTI     Member VLANs
   0 0-199,201-299,301-399,401-499,501-599,602-4094
   1 200,600-601
   2 400,500
```

```
root@ex2# set protocols mstp configuration-name north
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set protocols mstp revision-level 1

root@ex2# set protocols mstp msti 1 vlan 200

root@ex2# set protocols mstp msti 1 vlan 600

root@ex2# set protocols mstp msti 1 vlan 601

root@ex2# set protocols mstp msti 2 vlan 400

root@ex2# set protocols mstp msti 2 vlan 500
```

```
root@ex2# run show spanning-tree mstp configuration
MSTP information
Context identifier      : 0
Region name             : north
Revision                : 1
Configuration digest    : 0xf3d7ba47b1fbb2ea735860b1e839f961


MSTI     Member VLANs
   0 0-199,201-299,301-399,401-499,501-599,602-4094
   1 200,600-601
   2 400,500
```

Verification:

```
root@srx1# run show spanning-tree bridge

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 8192.b0:a8:6e:b9:61:10
  CIST regional root                : 8192.b0:a8:6e:b9:61:10
  CIST internal root cost           : 0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 5
  Time since last topology change   : 7245 seconds
  Topology change initiator         : ge-0/0/6.0
  Topology change last recvd. from  : 3c:94:d5:46:73:09
  Local parameters
    Bridge ID                       : 8192.b0:a8:6e:b9:61:10
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  MSTI regional root                 : 8193.b0:a8:6e:b9:61:10
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Number of topology changes         : 4
  Topology change initiator          : ge-0/0/9.0
  Topology change last recvd. from   : 3c:94:d5:46:73:0c
  Local parameters
    Bridge ID                        : 8193.b0:a8:6e:b9:61:10
    Extended system ID               : 0
    Internal instance ID             : 1

STP bridge parameters for MSTI 2
  MSTI regional root                 : 8194.78:19:f7:18:3d:90
  Root cost                          : 10000
  Root port                          : ae0.0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 19
  Number of topology changes         : 4
  Topology change initiator          : ae0.0
  Topology change last recvd. from   : 78:19:f7:18:3e:00
  Local parameters
    Bridge ID                        : 32770.b0:a8:6e:b9:61:10
    Extended system ID               : 0
    Internal instance ID             : 2

root@srx2# run show spanning-tree bridge

STP bridge parameters
Context ID                           : 0
Enabled protocol                     : MSTP

STP bridge parameters for CIST
  Root ID                            : 8192.b0:a8:6e:b9:61:10
  Root cost                          : 0
  Root port                          : ae0.0
  CIST regional root                 : 8192.b0:a8:6e:b9:61:10
  CIST internal root cost            : 10000
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 19
  Message age                        : 0
  Number of topology changes         : 5
  Time since last topology change    : 7256 seconds
  Topology change initiator          : ae0.0
  Topology change last recvd. from   : b0:a8:6e:b9:61:80
  Local parameters
    Bridge ID                        : 32768.78:19:f7:18:3d:90
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 8193.b0:a8:6e:b9:61:10
  Root cost                         : 10000
  Root port                         : ae0.0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Number of topology changes        : 4
  Topology change initiator         : ae0.0
  Topology change last recvd. from  : b0:a8:6e:b9:61:80
  Local parameters
    Bridge ID                       : 32769.78:19:f7:18:3d:90
    Extended system ID              : 0
    Internal instance ID            : 1

STP bridge parameters for MSTI 2
  MSTI regional root                : 8194.78:19:f7:18:3d:90
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 4
  Topology change initiator         : ge-0/0/9.0
  Topology change last recvd. from  : 2c:6b:f5:8f:87:cc
  Local parameters
    Bridge ID                       : 8194.78:19:f7:18:3d:90
    Extended system ID              : 0
    Internal instance ID            : 2

root@ex1# run show spanning-tree bridge

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 8192.b0:a8:6e:b9:61:10
  Root cost                         : 0
  Root port                         : ge-0/0/6.0
  CIST regional root                : 8192.b0:a8:6e:b9:61:10
  CIST internal root cost           : 20000
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Message age                       : 0
  Number of topology changes        : 4
  Time since last topology change   : 7273 seconds
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Topology change initiator          : ge-0/0/6.0
  Topology change last recvd. from   : 2e:6b:f5:8f:87:c3
  Local parameters
    Bridge ID                        : 32768.3c:94:d5:46:73:01
    Extended system ID               : 0
    Internal instance ID             : 0

STP bridge parameters for MSTI 1
  MSTI regional root                 : 8193.b0:a8:6e:b9:61:10
  Root cost                          : 20000
  Root port                          : ge-0/0/9.0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 19
  Number of topology changes         : 3
  Topology change initiator          : ge-0/0/9.0
  Topology change last recvd. from   : 2e:6b:f5:8f:87:c3
  Local parameters
    Bridge ID                        : 32769.3c:94:d5:46:73:01
    Extended system ID               : 0
    Internal instance ID             : 1

STP bridge parameters for MSTI 2
  MSTI regional root                 : 8194.78:19:f7:18:3d:90
  Root cost                          : 25000
  Root port                          : ae0.0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Number of topology changes         : 3
  Topology change initiator          : ge-0/0/18.0
  Topology change last recvd. from   : b0:a8:6e:b9:61:09
  Local parameters
    Bridge ID                        : 32770.3c:94:d5:46:73:01
    Extended system ID               : 0
    Internal instance ID             : 2

root@ex2# run show spanning-tree bridge

STP bridge parameters
Context ID                           : 0
Enabled protocol                     : MSTP

STP bridge parameters for CIST
  Root ID                            : 8192.b0:a8:6e:b9:61:10
  Root cost                          : 0
  Root port                          : ge-0/0/6.0
  CIST regional root                 : 8192.b0:a8:6e:b9:61:10
  CIST internal root cost            : 15000
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Hello time                          : 2 seconds
  Maximum age                         : 20 seconds
  Forward delay                       : 15 seconds
  Hop count                           : 18
  Message age                         : 0
  Number of topology changes          : 5
  Time since last topology change     : 7273 seconds
  Topology change initiator           : ge-0/0/6.0
  Topology change last recvd. from    : 78:19:f7:18:3d:86
  Local parameters
    Bridge ID                         : 32768.2c:6b:f5:8f:87:c1
    Extended system ID                : 0
    Internal instance ID              : 0

STP bridge parameters for MSTI 1
  MSTI regional root                  : 8193.b0:a8:6e:b9:61:10
  Root cost                           : 15000
  Root port                           : ge-0/0/6.0
  Hello time                          : 2 seconds
```

```
Maximum age                         : 20 seconds
  Forward delay                       : 15 seconds
  Hop count                           : 18
  Number of topology changes          : 4
  Topology change initiator           : ge-0/0/6.0
  Topology change last recvd. from    : 78:19:f7:18:3d:86
  Local parameters
    Bridge ID                         : 32769.2c:6b:f5:8f:87:c1
    Extended system ID                : 0
    Internal instance ID              : 1

STP bridge parameters for MSTI 2
  MSTI regional root                  : 8194.78:19:f7:18:3d:90
  Root cost                           : 20000
  Root port                           : ge-0/0/9.0
  Hello time                          : 2 seconds
  Maximum age                         : 20 seconds
  Forward delay                       : 15 seconds
  Hop count                           : 19
  Number of topology changes          : 5
  Topology change initiator           : ae0.0
  Topology change last recvd. from    : b0:a8:6e:b9:61:0a
  Local parameters
    Bridge ID                         : 32770.2c:6b:f5:8f:87:c1
    Extended system ID                : 0
    Internal instance ID              : 2
```

2) In order to achieve our objective, we configure a lower cost for interface ge-6/0/0 and a higher cost for ge-10/0/0 on ex2:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set protocols mstp interface ge-0/0/6.0 cost 5000

root@ex2# set protocols mstp interface ge-0/0/10.0 cost 200000
```

```
root@ex2# run show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port
State  Role
                         port ID         bridge ID          Cost
ae0.0            128:1        128:1  32768.2c6bf58f87c1     10000  FWD
DESG
ge-0/0/6.0       128:519      128:519  32768.7819f7183d90      5000  FWD
ROOT
ge-0/0/7.0       128:520      128:520  32768.2c6bf58f87c1     20000  FWD
DESG
ge-0/0/8.0       128:521      128:521  32768.2c6bf58f87c1     20000  FWD
DESG
ge-0/0/9.0       128:522      128:522  32768.7819f7183d90     20000  BLK
ALT
ge-0/0/10.0      128:523      128:523   8192.b0a86eb96110    200000  BLK
ALT
ge-0/0/17.0      128:530      128:530  32768.2c6bf58f87c1     20000  FWD
DESG
ge-0/0/19.0      128:532      128:532  32768.2c6bf58f87c1     20000  FWD
DESG
```

3) To meet the task requirement, the interface cost and port priority were changed. Keep in mind that the port priority is used in MSTP calculation when two ports have same cost to the root bridge:

```
root@ex1# set protocols mstp msti 1 interface ge-0/0/9.0 priority 16

root@ex1# set protocols mstp msti 2 interface ae0.0 cost 5000
```

```
root@ex1# run show spanning-tree interface
Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port
State  Role
                         port ID         bridge ID          Cost
ae0.0            128:1        128:1  32769.2c6bf58f87c1     10000  BLK
ALT
ge-0/0/6.0       128:519      128:519   8193.b0a86eb96110     20000  BLK
ALT
ge-0/0/7.0       128:520      128:520  32769.3c94d5467301     20000  FWD
DESG
ge-0/0/9.0        16:522       16:522   8193.b0a86eb96110     20000  FWD
ROOT
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
ge-0/0/10.0     128:523       128:523   32769.7819f7183d90      20000   BLK
ALT
ge-0/0/18.0     128:531       128:531   32769.3c94d5467301      20000   FWD
DESG



Spanning tree interface parameters for instance 2

Interface     Port ID     Designated       Designated          Port
State   Role
                           port ID          bridge ID           Cost
ae0.0             128:1         128:1   32770.2c6bf58f87c1       5000   FWD
ROOT
ge-0/0/6.0      128:519       128:519   32770.b0a86eb96110      20000   BLK
ALT
ge-0/0/9.0      128:522       128:522   32770.b0a86eb96110      20000   BLK
ALT
ge-0/0/10.0     128:523       128:523    8194.7819f7183d90     200000   BLK
ALT
ge-0/0/18.0     128:531       128:531   32770.3c94d5467301      20000   FWD
DES
```

4) By changing cost and port priority, we can ensure that port ge-0/0/9.0 is the root port for instance 2 and that port ge-0/0/6.0 is the root port for instance 1 on ex2:

```
root@ex2# set protocols mstp msti 1 interface ge-0/0/6.0 cost 5000

root@ex2# set protocols mstp msti 1 interface ge-0/0/10.0 cost 200000

root@ex2# set protocols mstp msti 2 interface ge-0/0/9.0 priority 16
```

```
root@ex2# run show spanning-tree interface
Spanning tree interface parameters for instance 1

Interface     Port ID     Designated       Designated          Port
State   Role
                           port ID          bridge ID           Cost
ae0.0             128:1         128:1   32769.2c6bf58f87c1      10000   FWD
DESG
ge-0/0/6.0      128:519       128:519   32769.7819f7183d90       5000   FWD
ROOT
ge-0/0/9.0      128:522       128:522   32769.7819f7183d90      20000   BLK
ALT
ge-0/0/10.0     128:523       128:523    8193.b0a86eb96110     200000   BLK
ALT
ge-0/0/19.0     128:532       128:532   32769.2c6bf58f87c1      20000   FWD
DESG

Spanning tree interface parameters for instance 2
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Interface     Port ID    Designated      Designated          Port
State  Role
                         port ID         bridge ID           Cost
ae0.0            128:1        128:1  32770.2c6bf58f87c1      10000  FWD
DESG
ge-0/0/6.0     128:519      128:519   8194.7819f7183d90      20000  BLK
ALT
ge-0/0/7.0     128:520      128:520  32770.2c6bf58f87c1      20000  FWD
DESG
ge-0/0/9.0      16:522       16:522   8194.7819f7183d90      20000  FWD
ROOT
ge-0/0/10.0    128:523      128:523  32770.b0a86eb96110      20000  BLK
ALT
ge-0/0/19.0    128:532      128:532  32770.2c6bf58f87c1      20000  FWD
DESG
```

5) We can complete this task by lowering cost of the ex3_ex4 interface that is facing ex2:

```
root@ex3_ex4# set protocols rstp interface ge-1/0/19.0 cost 5000
```

```
root@ex3_ex4# run show spanning-tree interface | no-more

Spanning tree interface parameters for instance 0

Interface     Port ID    Designated      Designated          Port
State  Role
                         port ID         bridge ID           Cost
ge-0/0/6.0     128:519      128:519  32768.2c6bf5351b81      20000  FWD
DESG
ge-0/0/7.0     128:520      128:520  32768.2c6bf5351b81      20000  FWD
DESG
ge-0/0/8.0     128:521      128:521  32768.2c6bf5351b81      20000  FWD
DESG
ge-0/0/18.0    128:531      128:531   8192.b0a86eb96110      20000  BLK
ALT
ge-1/0/6.0     128:631      128:631  32768.2c6bf5351b81      20000  FWD
DESG
ge-1/0/7.0     128:632      128:632  32768.2c6bf5351b81      20000  FWD
DESG
ge-1/0/8.0     128:633      128:633  32768.2c6bf5351b81      20000  FWD
DESG
ge-1/0/19.0    128:644      128:532   8192.b0a86eb96110       5000  FWD
ROOT
```

6) The MSTP configuration on srx3 and srx4 is similar to the configuration previously applied to srx1 and srx2. The difference lies in the MSTP region name. We also specify different bridge-priority to achieve the objectives specified in this task:

```
root@srx3# set protocols mstp configuration-name south
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set protocols mstp revision-level 1

root@srx3# set protocols mstp bridge-priority 32k

root@srx3# set protocols mstp msti 3 bridge-priority 8k

root@srx3# set protocols mstp msti 3 vlan 300
```

```
root@srx3# run show spanning-tree mstp configuration
MSTP information
Context identifier      : 0
Region name             : south
Revision                : 1
Configuration digest    : 0xf3fd759eb02331250fda30f7fce939b4


MSTI     Member VLANs
   0 0-299,301-4094
   3 300
```

```
root@srx4# set protocols mstp configuration-name south

root@srx4# set protocols mstp revision-level 1

root@srx4# set protocols mstp bridge-priority 60k

root@srx4# set protocols mstp msti 3 vlan 300
```

```
root@srx4# run show spanning-tree mstp configuration
MSTP information
Context identifier      : 0
Region name             : south
Revision                : 1
Configuration digest    : 0xf3fd759eb02331250fda30f7fce939b4


MSTI     Member VLANs
   0 0-299,301-4094
   3 300
root@srx3# run show spanning-tree bridge

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 8192.b0:a8:6e:b9:61:10
  Root cost                         : 20000
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Root port                          : ge-0/0/12.0
  CIST regional root                 : 32768.3c:8a:b0:5c:7a:b0
  CIST internal root cost            : 0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 20
  Message age                        : 1
  Number of topology changes         : 4
  Time since last topology change    : 7260 seconds
  Topology change initiator          : ge-0/0/12.0
  Topology change last recvd. from   : 3c:94:d5:46:73:13
  Local parameters
    Bridge ID                        : 32768.3c:8a:b0:5c:7a:b0
    Extended system ID               : 0
    Internal instance ID             : 0

STP bridge parameters for MSTI 3
  MSTI regional root                 : 8195.3c:8a:b0:5c:7a:b0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Number of topology changes         : 3
  Topology change initiator          : ge-0/0/12.0
  Topology change last recvd. from   : 3c:61:04:8d:be:82
  Local parameters
    Bridge ID                        : 8195.3c:8a:b0:5c:7a:b0
    Extended system ID               : 0
    Internal instance ID             : 1
root@srx4# run show spanning-tree bridge

STP bridge parameters
Context ID                           : 0
Enabled protocol                     : MSTP

STP bridge parameters for CIST
  Root ID                            : 8192.b0:a8:6e:b9:61:10
  Root cost                          : 20000
  Root port                          : ge-0/0/1.0
  CIST regional root                 : 32768.3c:8a:b0:5c:7a:b0
  CIST internal root cost            : 20000
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 19
  Message age                        : 1
  Number of topology changes         : 1
  Time since last topology change    : 7782 seconds
  Topology change initiator          : ge-0/0/1.0
  Topology change last recvd. from   : 2c:6b:f5:8f:87:d4
  Local parameters
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Bridge ID                          : 61440.3c:61:04:8d:be:90
    Extended system ID                 : 0
    Internal instance ID               : 0

STP bridge parameters for MSTI 3
  MSTI regional root                   : 8195.3c:8a:b0:5c:7a:b0
  Root cost                            : 20000
  Root port                            : ge-0/0/2.0
  Hello time                           : 2 seconds
  Maximum age                          : 20 seconds
  Forward delay                        : 15 seconds
  Hop count                            : 19
  Number of topology changes           : 2
  Topology change initiator            : ge-0/0/2.0
  Topology change last recvd. from     : 3c:8a:b0:5c:7a:a2
  Local parameters
    Bridge ID                          : 32771.3c:61:04:8d:be:90
    Extended system ID                 : 0
    Internal instance ID               : 1
```

7) Creating a new instance in region north and ensure that ex1 is the root-bridge for this instance:

```
root@ex1# set protocols mstp msti 3 bridge-priority 8k

root@ex1# set protocols mstp msti 3 vlan 300
```

```
root@ex1# run show spanning-tree bridge
STP bridge parameters for MSTI 3
  MSTI regional root                   : 8195.3c:94:d5:46:73:01
  Hello time                           : 2 seconds
  Maximum age                          : 20 seconds
  Forward delay                        : 15 seconds
  Number of topology changes           : 3
  Topology change initiator            : ge-0/0/16.0
  Topology change last recvd. from     : 2e:6b:f5:8f:87:c3
  Local parameters
    Bridge ID                          : 8195.3c:94:d5:46:73:01
    Extended system ID                 : 0
    Internal instance ID               : 3
```

On srx1, srx2 and ex2:

```
root@srx1# set protocols mstp msti 3 vlan 300
```

8) Make sure that srx4 uses ge-0/0/2 to forward traffic to srx3 for VLAN 300 by making sure that the ge-0/0/2 port is the root port for instance 3:

On srx3 and srx4:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3#set protocols mstp msti 3 interface ge-0/0/2.0 priority 16
```

```
root@srx4# run show spanning-tree interface
Spanning tree interface parameters for instance 3

Interface    Port ID    Designated      Designated         Port
State  Role
                        port ID         bridge ID          Cost
ge-0/0/1.0     128:514      128:514   8195.3c8ab05c7ab0     20000  BLK
ALT
ge-0/0/2.0      16:515       16:515   8195.3c8ab05c7ab0     20000  FWD
ROOT
ge-0/0/13.0    128:526      128:526  32771.3c61048dbe90     20000  BLK
ALT
ge-0/0/14.0    128:527      128:527  32771.3c61048dbe90     20000  FWD
DESG
```

9) In order to make sure that the ports can transition to the forwarding state immediately, we configure the ports as **edge** ports. In order to simplify the configuration, we this options by using the **interface-range** command. To ensure that the ports are disabled as soon as BPDUs are received, we use **bpdu-block-on-edge** option:

```
root@ex3_ex4# set interfaces interface-range access-ports member-range
ge-0/0/0 to ge-0/0/4

root@ex3_ex4# set interfaces interface-range access-ports member-range
ge-1/0/0 to ge-1/0/4

root@ex3_ex4# set interfaces interface-range access-ports unit 0
family ethernet-switching

root@ex3_ex4# set protocols rstp interface access-ports edge

root@ex3_ex4# set protocols rstp bpdu-block-on-edge
```

10) Make sure that the ex1 and ex2 uplinks to ex3_ex3 will be put into the blocking state as soon as superior BPDUs are received:

```
root@ex1# set protocols mstp interface ge-0/0/18.0 no-root-port

root@ex2# set protocols mstp interface ge-0/0/19.0 no-root-port
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 5: L2 Switching Security Features

1) The 802.1X configuration is based on the explanation in the beginning of this chapter. The interface range feature is used for VLAN assignment:

```
root@ex3_ex4# set vlans Secured interface access-ports
root@ex3_ex4# set access radius-server 10.10.10.1 secret workbook
root@ex3_ex4# set access profile dot1x authentication-order radius
root@ex3_ex4# set access profile dot1x radius authentication-server
10.10.10.1
root@ex3_ex4# set protocols dot1x authenticator authentication-
profile-name dot1x
root@ex3_ex4# set protocols dot1x authenticator interface ge-1/0/3.0
mac-radius restrict
```

2) Enabling MAC RADIUS for non-802.1X enabled devices:

```
root@ex3_ex4# set protocols dot1x authenticator interface access-ports
mac-radius
```

3) Configure ge-1/0/3 to bypass 802.1X authentication and perform MAC RADIUS immediately.

```
root@ex3_ex4# set protocols dot1x authenticator interface ge-1/0/3.0
mac-radius restrict
```

4) Enable captive portal authentication as a last resort solution:

NOTE: A captive portal is configured as a service. A captive portal redirects Web requests to a login page. A captive portal profile is configured in a way that is similar to an 802.1X profile. For the configuration, we need to enable the http service on the switch and define a term in the control-plane firewall filter that accepts incoming web requests.

```
root@ex3_ex4# set system services web-management http
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set services captive-portal authentication-profile-name
dot1x

root@ex3_ex4# set services captive-portal interface access-ports

root@ex3_ex4# set services captive-portal secure-authentication http

root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term
http from protocol tcp

root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term
http from destination-port http

root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term
http then accept

root@ex3_ex4# insert firewall family inet filter protect-re-ipv4 term
http before term icmp
```

5) Enabling the MAC address security feature that limits MAC movements.

```
root@ex1# set ethernet-switching-options secure-access-port vlan all
mac-move-limit 5

root@ex1# set ethernet-switching-options secure-access-port vlan all
mac-move-limit action drop
```

6) Enabling MAC limitation, configured with the drop action when violated:

```
root@ex1# set ethernet-switching-options secure-access-port interface
secured-ports mac-limit 2

root@ex1# set ethernet-switching-options secure-access-port interface
secured-ports mac-limit action drop
```

7) Enabling MAC limitation for port ge-0/0/7.0, configured to shutdown the port when violated:

```
root@ex1# set ethernet-switching-options secure-access-port interface
ge-0/0/7.0 mac-limit 1

root@ex1# set ethernet-switching-options secure-access-port interface
ge-0/0/7.0 mac-limit action shutdown
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

8) Configuring ARP inspection, DHCP snooping and trusted DHCP server:

Note: above features are security methods to prevent rogue access to the network.

To complete this task, we also need to add a term to the control-plane firewall filter accepting DHCP.

```
root@ex3_ex4# set ethernet-switching-options secure-access-port
interface ge-0/0/18.0 dhcp-trusted

root@ex3_ex4# set ethernet-switching-options secure-access-port
interface ge-1/0/19.0 dhcp-trusted

root@ex3_ex4# set ethernet-switching-options secure-access-port vlan
400 arp-inspection

root@ex3_ex4# set ethernet-switching-options secure-access-port vlan
400 examine-dhcp

root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term
dhcp-1 from protocol udp

root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term
dhcp-1 from source-port dhcp

root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term
dhcp-1 then accept

root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term
dhcp-2 from protocol udp

root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term
dhcp-2 from destination-port dhcp

root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term
dhcp-2 then accept

root@ex3_ex4# insert firewall family inet filter protect-re-ipv4 term
dhcp-2 before term http

root@ex3_ex4# insert firewall family inet filter protect-re-ipv4 term
dhcp-1 before term dhcp-2
```

```
root@ex3_ex4# run show dhcp snooping binding
DHCP Snooping Information:
MAC address        IP address        Lease (seconds)  Type     VLAN
Interface
3C:61:04:8D:BE:86  192.168.31.21              74955  dynamic  Omega
ge-1/0/6.0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# run show arp inspection statistics
Interface      Packets received     ARP inspection pass  ARP inspection
failed
  ge-1/0/6                  21                      21
0
  ge-1/0/7                   0                       0
0
  ge-1/0/8                   0                       0
0
  ge-1/0/9                   0                       0
0
 ge-1/0/10                   0                       0
0
 ge-1/0/11                   0                       0
0
 ge-1/0/12                   0                       0
0
 ge-1/0/13                   0                       0
0
 ge-1/0/14                   0                       0
0
 ge-1/0/15                   0                       0
0
 ge-1/0/16                   0                       0
0
 ge-1/0/17                   0                       0
0
 ge-1/0/18                   0                       0
0
 ge-1/0/19                  38                      38
0
```

Note: you will an output similar to the above after DHCP configuration.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 6: VRRP Configuration

1) Configure the Layer 3 interface for VLAN200, VLAN300 and VLAN400.

To configure L3 interfaces for a VLAN, we first need to define a VLAN interface with an IP address:

```
root@srx1# set interfaces vlan unit 200 family inet address
192.168.30.252/24

root@srx1# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::2/64

root@srx1# set interfaces vlan unit 300 family inet6 address
fe80::2/64

root@srx1# set interfaces vlan unit 400 family inet address
192.168.31.252/24
```

Afterwards, we bind the newly created interface to the VLAN:

```
root@srx1# set vlans Omega l3-interface vlan.400

root@srx1# set vlans Theta l3-interface vlan.200

root@srx1# set vlans Zeta l3-interface vlan.300
```

```
root@srx1# run show vlans brief
                                                    Ports
Name            Tag       Primary Address         Active/Total
Kappa           600                                 4/4
Omega           400       192.168.31.252/24         4/4
Secured         500                                 4/4
Sigma           700
Sigma-transport 701                                 3/3
Theta           200       192.168.30.252/24         4/4
Voice           601                                 4/4
Zeta            300                                 1/1
                          fd08:a131:17d8:82d6::2/64
```

On srx2:

```
root@srx2# set interfaces vlan unit 200 family inet address
192.168.30.253/24

root@srx2# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::3/64
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set interfaces vlan unit 300 family inet6 address
fe80::3/64

root@srx2# set interfaces vlan unit 400 family inet address
192.168.31.253/24

root@srx2# set vlans Omega l3-interface vlan.400

root@srx2# set vlans Theta l3-interface vlan.200

root@srx2# set vlans Zeta l3-interface vlan.300
```

```
root@srx2# run show vlans brief
                                                    Ports
Name            Tag       Primary Address         Active/Total
Kappa           600                                  4/4
Omega           400       192.168.31.253/24          4/4
Secured         500                                  4/4
Sigma           700
Sigma-transport 701                                  3/3
Theta           200       192.168.30.253/24          4/4
Voice           601                                  4/4
Zeta            300                                  1/1
                          fd08:a131:17d8:82d6::3/64
                                                    Ports
Name            Tag       Primary Address         Active/Total
default         1                                    4/4
```

2) The configuration of the DHCP pool and the necessary changes in the firewall filters:

```
root@srx1# set system services dhcp pool 192.168.31.0/24 address-range
low 192.168.31.20

root@srx1# set system services dhcp pool 192.168.31.0/24 address-range
high 192.168.31.39

root@srx1# set system services dhcp pool 192.168.31.0/24 router
192.168.31.254
```

Firewall Filter Configuration:

```
root@srx1# set firewall family inet filter protect-re-ipv4 term dhcp-1
from protocol udp

root@srx1# set firewall family inet filter protect-re-ipv4 term dhcp-1
from source-port dhcp

root@srx1# set firewall family inet filter protect-re-ipv4 term dhcp-1
then accept
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set firewall family inet filter protect-re-ipv4 term dhcp-2
from protocol udp

root@srx1# set firewall family inet filter protect-re-ipv4 term dhcp-2
from destination-port dhcp

root@srx1# set firewall family inet filter protect-re-ipv4 term dhcp-2
then accept

root@srx1# insert firewall family inet filter protect-re-ipv4 term
dhcp-2 before term telnet-local

root@srx1# insert firewall family inet filter protect-re-ipv4 term
dhcp-1 before term dhcp-2
```

On srx2:

```
root@srx2# set system services dhcp pool 192.168.31.0/24 address-range
low 192.168.31.40

root@srx2# set system services dhcp pool 192.168.31.0/24 address-range
high 192.168.31.59

root@srx2# set system services dhcp pool 192.168.31.0/24 router
192.168.31.254
```

Firewall Filter Configuration:

```
root@srx2# set firewall family inet filter protect-re-ipv4 term dhcp-1
from protocol udp

root@srx2# set firewall family inet filter protect-re-ipv4 term dhcp-1
from source-port dhcp

root@srx2# set firewall family inet filter protect-re-ipv4 term dhcp-1
then accept

root@srx2# set firewall family inet filter protect-re-ipv4 term dhcp-2
from protocol udp

root@srx2# set firewall family inet filter protect-re-ipv4 term dhcp-2
from destination-port dhcp

root@srx2# set firewall family inet filter protect-re-ipv4 term dhcp-2
then accept
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# insert firewall family inet filter protect-re-ipv4 term
dhcp-2 before term telnet-local

root@srx2# insert firewall family inet filter protect-re-ipv4 term
dhcp-1 before term dhcp-2
```

3) The IPv4 VRRP configuration specifying srx 1 as the master for VLAN 200 and srx2 as the master for VLAN 400:

On srx1:

```
root@srx1# set interfaces vlan unit 200 family inet address
192.168.30.252/24 vrrp-group 1 virtual-address 192.168.30.254

root@srx1# set interfaces vlan unit 200 family inet address
192.168.30.252/24 vrrp-group 1 priority 200

root@srx1# set interfaces vlan unit 200 family inet address
192.168.30.252/24 vrrp-group 1 accept-data

root@srx1# set interfaces vlan unit 400 family inet address
192.168.31.252/24 vrrp-group 2 virtual-address 192.168.31.254

root@srx1# set interfaces vlan unit 400 family inet address
192.168.31.252/24 vrrp-group 2 accept-data
```

On srx2:

```
root@srx2# set interfaces vlan unit 200 family inet address
192.168.30.253/24 vrrp-group 1 virtual-address 192.168.30.254

root@srx2# set interfaces vlan unit 200 family inet address
192.168.30.253/24 vrrp-group 1 accept-data

root@srx2# set interfaces vlan unit 400 family inet address
192.168.31.253/24 vrrp-group 2 virtual-address 192.168.31.254

root@srx2# set interfaces vlan unit 400 family inet address
192.168.31.253/24 vrrp-group 2 priority 200

root@srx2# set interfaces vlan unit 400 family inet address
192.168.31.253/24 vrrp-group 2 accept-data
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification:

```
root@srx1# run show vrrp brief
Interface     State       Group   VR state VR Mode   Timer    Type
Address
vlan.200      up              1   master   Active      A  0.157 lcl
192.168.30.252
                                                             vip
192.168.30.254
vlan.400      up              2   backup   Active      D  3.281 lcl
192.168.31.252
                                                             vip
192.168.31.254
                                                             mas
192.168.31.253
root@srx2# run show vrrp
Interface     State       Group   VR state VR Mode   Timer    Type
Address
vlan.200      up              1   backup   Active      D  2.921 lcl
192.168.30.253
                                                             vip
192.168.30.254
                                                             mas
192.168.30.252
vlan.400      up              2   master   Active      A  0.155 lcl
192.168.31.253
                                                             vip
192.168.31.254
```

4) The IPv6 VRRP configuration specifying srx 1 as the master for VLAN 300:

On srx1:

```
root@srx1# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::2/64 vrrp-inet6-group 3 virtual-inet6-address
fd08:a131:17d8:82d6::1

root@srx1# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::2/64 vrrp-inet6-group 3 virtual-link-local-
address fe80::300

root@srx1# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::2/64 vrrp-inet6-group 3 priority 200

root@srx1# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::2/64 vrrp-inet6-group 3 accept-data
```

On srx2:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::3/64 vrrp-inet6-group 3 virtual-inet6-address
fd08:a131:17d8:82d6::1

root@srx2# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::3/64 vrrp-inet6-group 3 virtual-link-local-
address fe80::300

root@srx2# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::3/64 vrrp-inet6-group 3 priority 100

root@srx2# set interfaces vlan unit 300 family inet6 address
fd08:a131:17d8:82d6::3/64 vrrp-inet6-group 3 accept-data
```

Verification:

```
root@srx1# run show vrrp interface vlan.300
vlan.300       up              3   master          Active    lcl
fd08:a131:17d8:82d6::2
                                                             vip
fe80::300
                                                             vip
fd08:a131:17d8:82d6::1
root@srx2# run show vrrp interface vlan.300
vlan.300       up              3   backup   Active      D  3.549 lcl
fd08:a131:17d8:82d6::3
                                                                 vip
fe80::300
                                                                 vip
fd08:a131:17d8:82d6::1
                                                                 mas
fe80::2
```

5) Authenticating IPv4 VRRP messages with MD5:

On srx1:

```
root@srx1# set interfaces vlan unit 200 family inet address
192.168.30.252/24 vrrp-group 1 authentication-type md5

root@srx1# set interfaces vlan unit 200 family inet address
192.168.30.252/24 vrrp-group 1 authentication-key workbook

root@srx1# set interfaces vlan unit 400 family inet address
192.168.31.252/24 vrrp-group 2 authentication-type md5

root@srx1# set interfaces vlan unit 400 family inet address
192.168.31.252/24 vrrp-group 2 authentication-key workbook
```

On srx1:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set interfaces vlan unit 200 family inet address
192.168.30.253/24 vrrp-group 1 authentication-type md5

root@srx2# set interfaces vlan unit 200 family inet address
192.168.30.253/24 vrrp-group 1 authentication-key workbook

root@srx2# set interfaces vlan unit 400 family inet address
192.168.31.253/24 vrrp-group 2 authentication-key workbook

root@srx2# set interfaces vlan unit 400 family inet address
192.168.31.253/24 vrrp-group 2 track interface ge-0/0/3.0 priority-
cost 150
```

Verification:

```
root@srx1# run show vrrp brief
Interface     State       Group   VR state VR Mode   Timer    Type
Address
vlan.200      up              1   master   Active      A  0.157 lcl
192.168.30.252
                                                           vip
192.168.30.254
vlan.300      up              3   master   Active      A  0.527 lcl
fd08:a131:17d8:82d6::2
                                                           vip
fe80::300
                                                           vip
fd08:a131:17d8:82d6::1
vlan.400      up              2   backup   Active      D  3.281 lcl
192.168.31.252
                                                           vip
192.168.31.254
                                                           mas
192.168.31.253
root@srx2# run show vrrp brief
Interface     State       Group   VR state VR Mode   Timer    Type
Address
vlan.200      up              1   backup   Active      D  3.341 lcl
192.168.30.253
                                                           vip
192.168.30.254
                                                           mas
192.168.30.252
vlan.300      up              3   backup   Active      D  3.549 lcl
fd08:a131:17d8:82d6::3
                                                           vip
fe80::300
                                                           vip
fd08:a131:17d8:82d6::1
                                                           mas
fe80::2
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
vlan.400        up                    2   master   Active      A  0.238 lcl
192.168.31.253
                                                                 vip
192.168.31.254
```

6) Configure srx1 and srx2 to track their ge-0/0/3.0 interface:

```
root@srx1# set interfaces vlan unit 200 family inet address
192.168.30.252/24 vrrp-group 1 track interface ge-0/0/3.0 priority-
cost 150

root@srx2# set interfaces vlan unit 400 family inet address
192.168.31.253/24 vrrp-group 2 track interface ge-0/0/3.0 priority-
cost 150
```

7) Configuring the required static routes on srx1 and srx2:

```
root@srx1# set routing-options static route 0.0.0.0/0 next-hop
192.168.0.1

root@srx2# set routing-options static route 0.0.0.0/0 next-hop
192.168.0.5
```

8) PING verification:

```
root@srx3# run ping routing-instance H7 192.168.168.168
PING 192.168.168.168 (192.168.168.168): 56 data bytes
64 bytes from 192.168.168.168: icmp_seq=0 ttl=63 time=1.637 ms
64 bytes from 192.168.168.168: icmp_seq=1 ttl=63 time=1.545 ms
64 bytes from 192.168.168.168: icmp_seq=2 ttl=63 time=1.419 ms
64 bytes from 192.168.168.168: icmp_seq=3 ttl=63 time=1.439 ms
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 7: Provider bridging / Q-in-Q implementation

1) To familiarize yourself with the setup, check the H1 and H2 routing-instances on srx1 and srx2:

```
root@srx1# show routing-instances H1
H1 {
        instance-type virtual-router;
        interface ge-0/0/8.700;
        routing-options {
            static {
                route 0.0.0.0/0 next-hop 192.168.1.254;
            }
        }
    }
root@srx1# show interfaces ge-0/0/8
description "H1 connection to EX1";
        vlan-tagging;
        unit 700 {
            vlan-id 700;
            family inet {
                address 192.168.1.1/24;
            }
        }
root@srx2# show routing-instances H2
H2 {
        instance-type virtual-router;
        interface ge-0/0/8.700;
        routing-options {
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
            static {
                route 0.0.0.0/0 next-hop 192.168.1.254;
            }
        }
    }
root@srx2# show interfaces ge-0/0/8
       description "H2 connection to EX2";
        vlan-tagging;
        unit 700 {
            vlan-id 700;
            family inet {
                address 192.168.1.2/24;
            }
        }
    }
```

We can see that both instances have an interface that connects the routing-instance to a switch. Both interfaces are configured with 'vlan-tagging' and the IP address configuration resides on a sub-interface that is configured for VLAN 700. The H1 instance is connected to ex1 interface ge-0/0/8.0 and the H2 instance is connected to ex2 interface ge-0/0/8.0. Both hosts reside in the same subnet and both hosts are using the same default gateway.

2) Configuring the transport VLAN, VLAN 701, on srx1, srx2, ex1 and ex2:

```
root@srx1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 701

root@srx1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 701

root@srx1# set interfaces ae0 unit 0 family ethernet-switching vlan
members 701

root@srx2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 701
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 701

root@srx2# set interfaces ae0 unit 0 family ethernet-switching vlan
members 701

root@ex1# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 701

root@ex1# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 701

root@ex2# set interfaces ge-0/0/6 unit 0 family ethernet-switching
vlan members 701

root@ex2# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 701
```

3) Apply the Many-to-One bundling configuration on ex1 and apply the Q-in-Q ether-type configuration on both ex1 and ex2:

```
root@ex1# set ethernet-switching-options dot1q-tunneling ether-type
0x8100

root@ex1# set vlans Sigma-transport dot1q-tunneling customer-vlans 700

root@ex2# set ethernet-switching-options dot1q-tunneling ether-type
0x8100
```

Note: there are three different ways you can map a C-VLANs to an S-VLAN:

Mapping a specific interface: most specific configuration. Here, you can specify a single C-VLAN on a single interfaces and map it to an S-VLAN. Here, we use the **mapping** option specifically for the interfaces in the VLAN configuration.

Many-to-one bundling : Here, you specify a single C-VLAN that is to be mapped to the S-VLAN. This option is a little less granular when we compare it to the previous method. Here, the **customer-vlans** option is used.

All-in-one bundling: Here, all packets from access interfaces are mapped to the S-VLAN. This is the least specific option. Here, the **dot1q-tunneling** option is used without specifying any C-VLAN.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

4) We need to map the interface on ex2, tunneling only on that specific interface into VLAN 701:

```
root@ex2# set vlans Sigma-transport interface ge-0/0/8.0 mapping 700
push
```

When mapping a specific interface, the configuration applies to the referenced interface only. With this configuration, we can choose to use either swap or push. In this case, we need the packet to keep the original VLAN tag, so we choose push. We can use the swap option when VLAN translation is required.

5) By issuing a ping from H1 to H2, we can verify that the Q-in-Q tunnel working properly:

```
root@srx1# run ping routing-instance H1 192.168.1.2
PING 192.168.1.2 (192.168.1.2): 56 data bytes
64 bytes from 192.168.1.2: icmp_seq=0 ttl=64 time=23.846 ms
64 bytes from 192.168.1.2: icmp_seq=1 ttl=64 time=6.585 ms
64 bytes from 192.168.1.2: icmp_seq=2 ttl=64 time=9.223 ms
64 bytes from 192.168.1.2: icmp_seq=3 ttl=64 time=2.208 ms
64 bytes from 192.168.1.2: icmp_seq=4 ttl=64 time=1.211 ms
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Chapter 3: IGP Routing

A JNCIE-ENT candidate should be familiar with OSPF. This includes the following topics:

- Configuring OSPFv2 and OSPFv3
- Multi-area OSPF implementation
- Summarization and route filtering
- Redistributing routes to OSPF
- Virtual link and Multi area Adjacencies
- OSPF consideration
- Monitoring and troubleshooting

Configuring OSPFv2 and OSPFv3

OSPFv2 is a link-state protocol that can be used to distribute IPv4 routing-information. Link State Advertisements (LSAs) are used to reliable flood routing information across the AS. These LSAs are stored in an LSDB. The SPF algorithm is run against the records in this database and the outcome is placed into the routing-table.

OSPFv2 uses the following packets:

1. **Hello** used to dynamically discover, establish and maintain neighbor relationships. Send to 224.0.0.5 (or all DRs 224.0.0.6) on a regular interval.
2. **Database Description** used during neighbor relationship establishment on multiaccess links. Packets transmit LSA headers between OSPF routers.
3. **Link-State Request** used to request missing LSAs or LSAs that are not up-to-date.
4. **Link-State Update** send in response to LSR or link state change. Carries a collection of LSAs.
5. **Link- State Acknowledgement** every flooded LSA is acknowledged with an LSAck, providing for a reliable flooding.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Inside the LSDB, we find different types of LSAs. These types are the following:

Type 1: Router-LSA.

Description of the state and costs of the routers' link inside an area. Originated for each area a routers belongs to and flooded throughout the entire area.

Type 2: Network-LSA.

Originated by the DR, describing the networks connected to a multi-access segment. Flooded throughout the area.

Type 3: Network-Summary LSA.

Describes inter-area routes to networks in other areas, originated by the ABR. These LSAs are flooded throughout the area.

Type 4: ASBR-Summary LSA.

Originated by an ABR attached to an area that contains an ASBR. Describes inter-area routes to AS boundary routers. Basically the ABR telling other areas, or the rest of the OSPF domain, where they can find a particular ASBR.

Type 5: External LSA.

Originated by ASBRs, describes the routes towards external destinations that have been redistributed into OSPF.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

During the formation of an OSPF neighbor relationship, the router will go through various states. The description of the different neighbor states (or conversation states) is as follows;

Down: nothing happened. OSPF is awaiting a start event. When there is a start event, a Hello is sent towards the neighbor.

Init: Hello is seen but it did not contain the router's own router-id.

2Way: Bidirectional communication is established. Both routers see their own router-id listed in the 'neighbor' field in the received hello packets.

ExStart: The routers determine who is slave and who is master during the following exchange state.

Exchange: The master and slave exchange LSA headers, describing their databases to each other. LSR's and LSU's may also be send during this state in case a router is missing some LSA's.

Loading: Link state information is exchanged in the form of LSR's and LSU's.

Full: When the LSR list is empty, the state is changed to Full.

OSPFv3 was developed to offer support for both IPv6 and IPv4. To quote RFC 5340: 'The fundamental mechanisms of OSPF (flooding, Designated Router (DR) election, area support, (Shortest Path First) SPF calculations, etc.) remain unchanged.'

In the [**protocols ospf3**] stanza, you can configure the IPv6 realm. Under the [**protocols ospf3 realm ipv4-unicast**], you can configure the IPv4 realm. In the IPv4 realm, you will find that some features (like plain or MD5 authentication, virtual link and traffic engineering) are not supported.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Apart from the address family (or realm) support, there are some additional differences between OSPFv2 and OSPFv3:

- In OSPFv3, authentication can only be done using IPsec.
- OSPF3 has a different packet format and it introduces new LSA types (some of which have a new flooding scope):

OSPFv3 LSA type - function code:	OSPFv3 LSA name:	OSPFv2 LSA type:	OSPFv2 LSa name:
1 (Area scope)	Router-LSA	1	Router-LSA
2 (Area scope)	Network-LSA	2	Network-LSA
3 (Area scope)	Inter-Area-Prefix-LSA	3	Summary-LSA
4 (Area scope)	Inter-Area-Router-LSA	4	ASBR-Summary-LSA
5 (AS scope)	AS-External-LSA	5	AS-External-LSA
6 (Area scope)	Deprecated		
7 (Area scope)	NSSA-LSA	7	NSSA External LSA
8 (link-local)	Link-LSA		
9 (Area scope)	Intra-Area-Prefix-LSA		

- OSPFv3 does protocol processing per-link instead of per subnet
- There is explicit support for multiple instances of OSPF3 on a single link
- The router-ID is still a 32-bit number but it must be configured manually

Very often, you'll find OSPFv2 configured for IPv4 and OSPFv3 configured for IPv6.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Multi-area OSPF implementation

To be able to optimize and scale a network running OSPF, it can be divided into smaller groups called areas. An area is a collection of networks that are administratively grouped together. A separate LSDB is maintained for every single area and the LSDB describing an area should be identical on all the routers inside that area.

Through the configuration of these different areas, LSA flooding is restrained. Furthermore, by dividing the network into different areas, you will be able to summarize prefixes. Summarization can have two positive effects on the network. It can reduce the size of the LSDB and it provide for extra stability. The extra stability is achieved by prevent routing-instability in one area, or a network external to OSPF, from affecting another area's LSDB.

A multi-area OSPF network consists of several different areas that must be directly connected to the backbone area. Routers connected to multiple areas are called Area Border Routers. These ABRs can transmit network information between different OSPF areas. Routers injecting information from outside of the OSPF domain (i.e. static, RIP, etc.) are called Autonomous System Boundary Routers. By default, an area is considered a 'regular' or 'normal' area and supports the flooding of LSA type 1, 2, 3, 4 and 5.

Apart from the 'regular' areas, there are a number of other stub area types to consider:

Stub Area: the regular stub area is void of any LSA type 4 and 5. Instead of having the external routing-information, stub routers rely on a default route from the ABR(s) inside the stub area. The stub area cannot contain an ASBR and only LSA type 1, 2 and 3 are flooded throughout the area.

Totally Stubby Area: totally stubby areas have their LSDBs further reduced by relying solely on the default route generated by the ABR.

Not So Stubby Area (NSSA): it is possible for routers inside the NSSA to be an ASBR and inject external routing-information into the OSPF domain. To redistribute external routing information into OSPF, the NSSA ASBR injects type 7 LSAs. These type 7 LSAs are translated into type 5 LSAs by the ABR that borders the non-NSSA area(s). When there are multiple ABRs, the ABR with the highest router-id does the translation. External routing information from other areas is not send into the NSSA. Inside the NSSA area, we can find type 1, 2 3 and 7 LSAs.

Totally NSSA: the ABR for a totally NSSA will not inject type 3 LSAs, making the routers inside the area rely on the default route for all internal and external routes that are not generated from within the NSSA.

Converting a stub or NSSA area into a totally stub or NSSA area can be done by using the **no-summaries** command. When you inject a route into a totally stubby area or totally NSSA, you can have its metric and its metric type adjusted. When you play around with the metric values, always keep in mind that a type 1 route will always be preferred over a type 2 route.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Also note that because the different stub areas do not contain all routing information, a virtual-link cannot transit these type of areas.

Summarization and route filtering

We can hide network instability and increase the network scalability through summarization and filtering. OSPF imposes several restriction on where you can perform summarization and route filtering. Because the LSDB has to be consistent on all routers within an area, you cannot summarize within an area. The general rule is that you can summarize where LSA and route conversion takes place (ABR/ASBR).

Within Junos, performing OSPF summarization and route-filtering can be done using the following options:

- Inter-area internal LSA summarization and filtering on the ABR, using the **area-range** configuration command. The **restrict** option can be used to filter out the routes contained within the summary.
- Inter-area internal LSA filtering on the ABR using the **network-summary-import** and **network-summary-export** configuration command.
- External route summarization and filtering on the ASBR through the use of aggregate routes and export policies.
- Inter-area NSSA external route summarization and filtering on the NSSA ABR using the **nssa area-range** command.
- Route-table filtering of external routes using import policies.

Keep in mind that only type 1, 2 and 7 LSAs can be summarized or filtered. Type 5 LSAs cannot be summarized or filtered.

Redistributing routes into OSPF

When you configure route-redistribution on an OSPF router, you turn the router into an ASBR. By default, the external routes redistributed into OSPF are external type 2 routes. The metric for these routes are never incremented. Through the use of a routing policy, you are able to change the route type. This is done using the **'then external type 1'** action. This will redistribute the matched routing information as external type 1. The external type 1 routes will have their metrics incremented when they are flooded throughout the network.

By using external type 1 and external type 2 routes when redistributing routing-information on multiple routers, you can easily create a primary and a backup path.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Virtual link and Multi area Adjacencies

All of the OSPF areas have to be connected to area 0. When an OSPF router does not have a physical connection to the backbone area, we can still connect the router to area 0 by means of a virtual link.

The virtual-link is configured in area 0 using the following configuration command:

[edit protocols ospf area 0.0.0.0]

virtual-link neighbor-id *router-id* transit-area *area-id*

The verification of the virtual-link can be done by issuing the **show ospf neighbor** and **show ospf interface** commands. The virtual-link can be configured with additional parameters such as authentication, altered hello interval, etc.

As stated previously, keep in mind that a virtual link cannot be established through a stub area.

To configure an multi-area adjacency, we need to add the **secondary** option to the interface in the OSPF configuration. The state of the secondary interface is always point-to-point. For more information about multi-area adjacency , please study RFC 5185.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

OSPF options and considerations.

- To divert traffic away from a router, we can use the **overload** command. This will set the metric for all of the routers' transit links (except links into stub areas) to 0xFFFF (the highest value).
- Type 1 and 2 LSAs are summarized using the **area-range** command under **[edit protocols ospf area X].** The type 7 LSAs summarization is configured under the **[edit protocols ospf area X nssa]** stanza.
- By default, every link has a cost of 1. We can change this manually, link by link, using the **metric** command. We can also change the way OSPF calculates the metric by issuing the **reference-bandwidth** command. Using this command, we change the metric calculation to *cost=Ref bandwidth/interface bandwidth*.
- To avoid a DR and BDR election, we can configure the link to be point-to-point using the '**set protocols ospf area x interface x interface-type p2p'** configuration command. Avoiding a DR and BDR election speeds up the neighbor relationship formation. It also avoids the network from advertising type 2 LSAs.
- To make sure a router is ineligible to become either the DR or the BDR on a link, you can configure it with a priority 0.
- To inject a prefix that is configured on an interface without forming an OSPF neighbor relationship on that interface, configure it using the **passive** option.
- The **no-nssa-abr** command disables exporting type 7 LSAs on ABR and ASBR alike.

Monitoring and troubleshooting OSPF

There are many reasons that can prevent an OSPF neighbor relationship to form between two routers:

1. Timer mismatch

2. Area number mismatch

3. Authentication mismatch (type of authentication, used passphrase, key-id in case of MD5)

4. Area type mismatch (stub area on one end versus 'regular' area on the other end)

5. Subnet mask mismatch

6. Duplicate router-id

7. MTU mismatch

8. A DR priority of 0 configured on both routers

9. Duplicate interface IP address

10. Interface type mismatch (p2p and p2mp for example)

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

To investigate the problems related to an OSPF neighbor relationship establishment, consider one of the following methods or commands:

- **monitor traffic interface *interface-name***

The monitor traffic interface command is used to display packets send to and from the routing-engine. Adding the keyword **detail** or **extensive** will give you additional information from the fields contained within the packet header. You can also narrow down the packets that you want to display by specifying a matching condition. For example, to see detailed OSPF information send and received across interface ge-0/0/4.123, issue the following command: **monitor traffic interface ge-0/0/4.123 matching "proto ospf" extensive**

- activating **traceoptions**

Traceoptions for OSPF are activated under the [**protocols ospf traceoptions**] or [**protocols ospf3 traceoptions**] stanza. To start a trace for general OSPF information, and to have the output stored in a file called 'ospf.log', you can configure the following:

set protocols ospf traceoptions file ospf.log
set protocols ospf traceoptions flag normal

You can use the '**show log ospf.log | last**' command to see the latest messages logged to the file. To continuously follow any messages added to the log file, you can issue the '**monitor start ospf.log**' command.

- **show ospf(3) neighbor**

This command displays what neighbor relationships are formed across the interfaces running OSPF. It will also display the priorities for the interfaces on the other side, the router-ID of the neighboring router and the countdown timer until the neighbor is declared dead. Using the extensive option will also tell you how long the adjacency has been up.

- **show ospf(3) interface**

This command will output the options that the interface is configured for. It will list the interface type, the area that the interface is placed in, the DR and BDR for the link (0.0.0.0 or none for p2p links) and the number of neighbors active on the link.

Using the 'extensive' option will also inform you of the configured Hello and Dead interval, the retransmit interval, area number, area-type, authentication type and subnet-mask.

- **show interface terse**

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Used to check the subnet that is configured on the interface and to verify whether or not the interface is up.

- **show interfaces ge-0/0/x.x | match MTU**

To verify the MTU that is applicable to the address family active on the interface.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 1: Base Network and Virtual Router Deployment

1) The ex3 and ex4 switches are configured as a virtual chassis. The topology in this chapter requires you to split the virtual chassis. To split the virtual-chassis, we deactivate the VC ports on both switches and we renumber the ex4 switch:

```
root@ex3_ex4> request virtual-chassis vc-port set interface  vcp-0
disable

root@ex3_ex4> request virtual-chassis vc-port set interface  vcp-1
disable

root@ex4> request virtual-chassis renumber member-id 1 new-member-id 0
```

Note: for more information about splitting of a virtual chassis please refer to chapter 2.

Verification:

```
root@ex4# run show virtual-chassis
Virtual Chassis ID: d57d.0a9c.2594
Virtual Chassis Mode: Enabled
                                          Mstr           Mixed
Neighbor List
Member ID  Status   Serial No    Model     prio  Role      Mode ID
Interface
0 (FPC 0)  Prsnt    BM0210359941 ex4200-24t 128  Master*      N

Member ID for next new member: 1 (FPC 1)
```

2) Load the base configuration on the devices:

```
root@srx1# load override terminal
```

One Ctrl+D is needed to load configuration then enter and commit.

Do not forget to reboot SRX devices to use packet mode.

3) Same procedures are used to load base configuration on the virtual device and AS switch.

Use **username *lab* password *lab123*** to connect these devices.

4) Configure interfaces, For example for srx1:

```
root@srx1# set interfaces ge-0/0/4 vlan-tagging
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set interfaces ge-0/0/4 unit 12 description "L3 connection
to SRX2"

root@srx1# set interfaces ge-0/0/4 unit 12 vlan-id 12

root@srx1# set interfaces ge-0/0/4 unit 12 family inet address
172.30.0.1/30

root@srx1# set interfaces ge-0/0/4 unit 12 family inet6 address
fd17:f0f4:f691::1/126

root@srx1# set interfaces ge-0/0/4 unit 13 description "L3 connection
to EX1"

root@srx1# set interfaces ge-0/0/4 unit 13 vlan-id 13

root@srx1# set interfaces ge-0/0/4 unit 13 family inet address
172.30.0.5/30

root@srx1# set interfaces ge-0/0/4 unit 13 family inet6 address
fd17:f0f4:f691::5/126

root@srx1# set interfaces ge-0/0/4 unit 14 description "L3 connection
to ISP1"

root@srx1# set interfaces ge-0/0/4 unit 14 vlan-id 14

root@srx1# set interfaces ge-0/0/4 unit 14 family inet address
172.30.80.2/30

root@srx1# set interfaces ge-0/0/4 unit 14 family inet6 address
fd17:f0f4:f691:80::2/126

root@srx1# set interfaces lo0 unit 0 family inet address
172.30.15.1/32

root@srx1# set interfaces lo0 unit 0 family inet6 address
fd17:f0f4:f691:15::1/128
```

Proceed and configure the other interfaces on the rest of the devices according to the table.

5) Configure the virtual routers on srx4:

```
root@srx1# set routing-instances SRX4a instance-type virtual-router

root@srx1# set routing-instances SRX4a interface ge-0/0/5.39

root@srx1# set routing-instances SRX4a interface ge-0/0/5.59
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set routing-instances SRX4a interface lo0.1
root@srx1# set routing-instances SRX4b instance-type virtual-router
root@srx1# set routing-instances SRX4b interface ge-0/0/5.31
root@srx1# set routing-instances SRX4b interface lo0.2
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 2: Multi Area OSPF Configuration

1) Configure multi-area OSPF according to the diagram. Ensure that no OSPF LSA type 2 are present in any of the LSDBs. Make sure that routers are explicitly configured with a router ID equal to the loopback IP address:

We can explicitly specify the router ID on all routers under the [**routing-options router-id**] stanza.

The other requirement is to ensure that no type 2 LSA is present in any of the LSDBs. A type 2 LSA, or Network LSA, is originated by the DR and describes the networks connected to a multi-access segment. By configuring all non-loopback interfaces as point-to-point interfaces, we ensure that no type 2 LSA is originated anywhere.

srx1:

```
root@srx1# set routing-options router-id 172.30.15.1

root@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface lo0.0 passive

root@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.12 interface-type p2p

root@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.13 interface-type p2p

root@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.14 passive

root@srx1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.13
interface-type p2p

root@srx1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.12
interface-type p2p

root@srx1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.14
passive

root@srx1# set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
```

srx2:

```
root@srx2# set routing-options router-id 172.30.15.2

root@srx2# set protocols ospf area 0.0.0.10 interface ge-0/0/4.25
interface-type p2p
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface lo0.0 passive

root@srx2# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.12 interface-type p2p

root@srx2# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.23 interface-type p2p

root@srx2# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.24 interface-type p2p

root@srx2# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.22 passive

root@srx2# set protocols ospf3 area 0.0.0.0 interface lo0.0 passive

root@srx2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.12
interface-type p2p

root@srx2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.23
interface-type p2p

root@srx2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.24
interface-type p2p

root@srx2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.22
passive
```

srx3:

```
root@srx3# set routing-options router-id 172.30.15.3

root@srx3# set protocols ospf3 realm ipv4-unicast area 0.0.0.2
interface ge-0/0/4.34 interface-type p2p

root@srx3# set protocols ospf3 realm ipv4-unicast area 0.0.0.2
interface ge-0/0/4.39 interface-type p2p

root@srx3# set protocols ospf3 realm ipv4-unicast area 0.0.0.2
interface lo0.0 passive

root@srx3# set protocols ospf3 realm ipv4-unicast area 0.0.0.3
interface ge-0/0/4.31 interface-type p2p

root@srx3# set protocols ospf3 area 0.0.0.2 interface lo0.0 passive
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set protocols ospf3 area 0.0.0.2 interface ge-0/0/4.39
interface-type p2p

root@srx3# set protocols ospf3 area 0.0.0.2 interface ge-0/0/4.34
interface-type p2p

root@srx3# set protocols ospf3 area 0.0.0.3 interface ge-0/0/4.31
interface-type p2p
```

srx4:

```
root@srx4# set routing-options router-id 172.30.15.4

root@srx4# set protocols ospf3 realm ipv4-unicast area 0.0.0.2
interface ge-0/0/4.64 interface-type p2p

root@srx4# set protocols ospf3 realm ipv4-unicast area 0.0.0.2
interface ge-0/0/4.34 interface-type p2p

root@srx4# set protocols ospf3 realm ipv4-unicast area 0.0.0.2
interface lo0.0 passive

root@srx4# set protocols ospf3 area 0.0.0.2 interface ge-0/0/4.64
interface-type p2p

root@srx4# set protocols ospf3 area 0.0.0.2 interface ge-0/0/4.34
interface-type p2p

root@srx4# set protocols ospf3 area 0.0.0.2 interface lo0.0 passive

root@srx4# set protocols ospf3 area 0.0.0.2 interface ge-0/0/4.41
passive
```

srx4a:

```
root@srx4# set routing-instances SRX4a routing-options router-id
172.30.15.9

root@srx4# set routing-instances SRX4a protocols ospf3 realm ipv4-
unicast area 0.0.0.0 interface lo0.1 passive

root@srx4# set routing-instances SRX4a protocols ospf3 realm ipv4-
unicast area 0.0.0.0 interface ge-0/0/5.59 interface-type p2p

root@srx4# set routing-instances SRX4a protocols ospf3 realm ipv4-
unicast area 0.0.0.2 interface ge-0/0/5.39 interface-type p2p
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx4# set routing-instances SRX4a protocols ospf3 area 0.0.0.0
interface ge-0/0/5.59 interface-type p2p

root@srx4# set routing-instances SRX4a protocols ospf3 area 0.0.0.0
interface lo0.1 passive

root@srx4# set routing-instances SRX4a protocols ospf3 area 0.0.0.2
interface ge-0/0/5.39 interface-type p2p
```

srx4b:

```
root@srx4# set routing-instances SRX4b routing-options router-id
172.30.15.10

root@srx4# set routing-instances SRX4b protocols ospf3 realm ipv4-
unicast area 0.0.0.3 interface ge-0/0/5.31 interface-type p2p

root@srx4# set routing-instances SRX4b protocols ospf3 realm ipv4-
unicast area 0.0.0.3 interface lo0.2 passive

root@srx4# set routing-instances SRX4b protocols ospf3 area 0.0.0.3
interface lo0.2 passive

root@srx4# set routing-instances SRX4b protocols ospf3 area 0.0.0.3
interface ge-0/0/5.31 interface-type p2p
```

ex1:

```
root@ex1# set routing-options router-id 172.30.15.5

root@ex1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface lo0.0 passive

root@ex1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/14.13 interface-type p2p

root@ex1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/14.23 interface-type p2p

root@ex1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/14.56 interface-type p2p

root@ex1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/14.59 interface-type p2p
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set protocols ospf3 realm ipv4-unicast area 0.0.0.1 area-
range 172.30.0.32/2

root@ex1# set protocols ospf3 realm ipv4-unicast area 0.0.0.1
interface ge-0/0/14.57 interface-type p2p

root@ex1# set protocols ospf3 area 0.0.0.0 interface lo0.0 passive

root@ex1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/14.13
interface-type p2p

root@ex1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/14.23
interface-type p2p

root@ex1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/14.56
interface-type p2p

root@ex1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/14.59
interface-type p2p
```

ex2:

```
root@ex2# set routing-options router-id 172.30.15.6

root@ex2# set protocols ospf area 0.0.0.10 interface ge-0/0/14.61
interface-type p2p

root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface lo0.0 passive

root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/14.56 interface-type p2p

root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/14.24 interface-type p2p

root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.1
interface ge-0/0/14.68 interface-type p2p

root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.2
interface ge-0/0/14.64 interface-type p2p

root@ex2# set protocols ospf3 area 0.0.0.0 interface lo0.0 passive

root@ex2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/14.56
interface-type p2p
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/14.24
interface-type p2p

root@ex2# set protocols ospf3 area 0.0.0.2 interface ge-0/0/14.64
interface-type p2p
```

ex3:

```
root@ex3# set routing-options router-id 172.30.15.7

root@ex3# set protocols ospf3 realm ipv4-unicast area 0.0.0.1
interface ge-0/0/14.57 interface-type p2p

root@ex3# set protocols ospf3 realm ipv4-unicast area 0.0.0.1
interface ge-0/0/14.78 interface-type p2p

root@ex3# set protocols ospf3 realm ipv4-unicast area 0.0.0.1
interface lo0.0 passive
```

ex4:

```
root@ex4# set routing-options router-id 172.30.15.8

root@ex4# set protocols ospf3 realm ipv4-unicast export to.ospf

root@ex4# set protocols ospf3 realm ipv4-unicast area 0.0.0.1
interface ge-0/0/14.68 interface-type p2p

root@ex4# set protocols ospf3 realm ipv4-unicast area 0.0.0.1
interface ge-0/0/14.78 interface-type p2p

root@ex4# set protocols ospf3 realm ipv4-unicast area 0.0.0.1
interface lo0.0 passive
```

Verification:

```
root@srx1# run show ospf3 neighbor instance all
Instance: master
Realm: ipv6-unicast
ID              Interface             State     Pri   Dead
172.30.15.2     ge-0/0/4.12           Full      128     38
  Neighbor-address fe80::7a19:f700:c18:3d84
172.30.15.5     ge-0/0/4.13           Full      128     37
  Neighbor-address fe80::3e94:d500:d46:7311
Realm: ipv4-unicast
ID              Interface             State     Pri   Dead
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
172.30.15.2      ge-0/0/4.12             Full       128     33
  Neighbor-address fe80::7a19:f700:c18:3d84
172.30.15.5      ge-0/0/4.13             Full       128     36
  Neighbor-address fe80::3e94:d500:d46:7311

root@srx2# run show ospf3 neighbor instance all
Instance: master
Realm: ipv6-unicast
ID               Interface              State     Pri   Dead
172.30.15.1      ge-0/0/4.12             Full       128     37
  Neighbor-address fe80::b2a8:6e00:cb9:6104
172.30.15.5      ge-0/0/4.23             Full       128     35
  Neighbor-address fe80::3e94:d500:1746:7311
172.30.15.6      ge-0/0/4.24             Full       128     31
  Neighbor-address fe80::2e6b:f500:188f:87d1
Realm: ipv4-unicast
ID               Interface              State     Pri   Dead
172.30.15.1      ge-0/0/4.12             Full       128     36
  Neighbor-address fe80::b2a8:6e00:cb9:6104
172.30.15.5      ge-0/0/4.23             Full       128     36
  Neighbor-address fe80::3e94:d500:1746:7311
172.30.15.6      ge-0/0/4.24             Full       128     32
  Neighbor-address fe80::2e6b:f500:188f:87d1

root@srx3# run show ospf3 neighbor instance all
Instance: master
Realm: ipv6-unicast
ID               Interface              State     Pri   Dead
172.30.15.4      ge-0/0/4.34             Full       128     30
  Neighbor-address fe80::3e61:400:228d:be84
172.30.15.9      ge-0/0/4.39             Full       128     30
  Neighbor-address fe80::3e61:400:278d:be85
172.30.15.10     ge-0/0/4.31             Full       128     33
  Neighbor-address fe80::3e61:400:1f8d:be85
Realm: ipv4-unicast
ID               Interface              State     Pri   Dead
172.30.15.4      ge-0/0/4.34             Full       128     36
  Neighbor-address fe80::3e61:400:228d:be84
172.30.15.9      ge-0/0/4.39             Full       128     32
  Neighbor-address fe80::3e61:400:278d:be85
172.30.15.10     ge-0/0/4.31             Full       128     34
  Neighbor-address fe80::3e61:400:1f8d:be85

root@srx4# run show ospf3 neighbor instance all
Instance: master
Realm: ipv6-unicast
ID               Interface              State     Pri   Dead
172.30.15.3      ge-0/0/4.34             Full       128     35
  Neighbor-address fe80::3e8a:b000:225c:7aa4
172.30.15.6      ge-0/0/4.64             Full       128     35
  Neighbor-address fe80::2e6b:f500:408f:87d1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Realm: ipv4-unicast
ID              Interface              State     Pri   Dead
172.30.15.3     ge-0/0/4.34            Full      128     38
  Neighbor-address fe80::3e8a:b000:225c:7aa4
172.30.15.6     ge-0/0/4.64            Full      128     33
  Neighbor-address fe80::2e6b:f500:408f:87d1

Instance: SRX4a

Realm: ipv6-unicast
ID              Interface              State     Pri   Dead
172.30.15.5     ge-0/0/5.59            Full      128     38
  Neighbor-address fe80::3e94:d500:3b46:7311
172.30.15.3     ge-0/0/5.39            Full      128     39
  Neighbor-address fe80::3e8a:b000:275c:7aa4

Realm: ipv4-unicast
ID              Interface              State     Pri   Dead
172.30.15.5     ge-0/0/5.59            Full      128     37
  Neighbor-address fe80::3e94:d500:3b46:7311
172.30.15.3     ge-0/0/5.39            Full      128     37
  Neighbor-address fe80::3e8a:b000:275c:7aa4

Instance: SRX4b

Realm: ipv6-unicast
ID              Interface              State     Pri   Dead
172.30.15.3     ge-0/0/5.31            Full      128     37
  Neighbor-address fe80::3e8a:b000:1f5c:7aa4

Realm: ipv4-unicast
ID              Interface              State     Pri   Dead
172.30.15.3     ge-0/0/5.31            Full      128     36
  Neighbor-address fe80::3e8a:b000:1f5c:7aa4
```

2) Configure OSPF to automatically calculate a metric value of 100 for all 1Gbps links:

The assignment specifies that the calculated metric for a 1Gbps link should be 100. When we configure a reference bandwidth, the costs are calculated using the following formula: **cost = ref-bandwidth/bandwidth**.

In order for a 1Gbps link to produce a metric of 100, we need to set the reference bandwidth to 100g (100/1=100).

```
root@srx1#set protocols ospf3 realm ipv4-unicast reference-bandwidth
100g

root@srx1#set protocols ospf3 reference-bandwidth 100g
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

3) This task has already been completed in the first task. We used the OSPFv3 protocol and configured a realm for IPv4. This gave us a single instance of OSPF that supports IPv4 and IPv6.

4) Area 1 routers ex3 and ex4 are not configured with IPv6. Use an OSPF process that would not require redistribution between Area 0 and Area 1. Area 1 routes should be present as Inter Area in Area 0:

To complete this task, we change Area 1 into an NSSA Area. Since the area type has to match on all devices, we configure this on ex1, ex2, ex3 and ex4:

```
root@ex1#set protocols ospf3 realm ipv4-unicast area 0.0.0.1 nssa
root@ex2#set protocols ospf3 realm ipv4-unicast area 0.0.0.1 nssa
root@ex3#set protocols ospf3 realm ipv4-unicast area 0.0.0.1 nssa
root@ex4#set protocols ospf3 realm ipv4-unicast area 0.0.0.1 nssa
```

5) The following is the IPsec configuration required for Area 1:

```
root@ex1#set security ipsec security-association ospfv3-sha-1 mode
transport
root@ex1#set security ipsec security-association ospfv3-sha-1 manual
direction bidirectional protocol ah
root@ex1#set security ipsec security-association ospfv3-sha-1 manual
direction bidirectional spi 256
root@ex1#set security ipsec security-association ospfv3-sha-1 manual
direction bidirectional authentication algorithm hmac-sha1-96
root@ex1#set security ipsec security-association ospfv3-sha-1 manual
direction bidirectional authentication key ascii-text
12345678901234567890
root@ex1#set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface
ge-0/0/14.57 interface-type p2p
root@ex1#set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface
ge-0/0/14.57 ipsec-sa ospfv3-sha-1
```

We configured the security association under the [**security ipsec**] stanza and we applied it under the interfaces configured for ospf3. The next thing we need to do is add a term in the IPv6 firewall filter accepting the Authentication Header (AH):

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1#set firewall family inet6 filter protect-re-ipv6 term ah from
next-header ah

root@ex1#set firewall family inet6 filter protect-re-ipv6 term ah then
accept

root@ex1#insert firewall family inet6 filter protect-re-ipv6 term ah
before term icmp
```

We then configure the opposite end of the link.

On ex3:

```
root@ex3#set security ipsec security-association ospfv3-sha-1 mode
transport

root@ex3#set security ipsec security-association ospfv3-sha-1 manual
direction bidirectional protocol ah

root@ex3#set security ipsec security-association ospfv3-sha-1 manual
direction bidirectional spi 256

root@ex3#set security ipsec security-association ospfv3-sha-1 manual
direction bidirectional authentication algorithm hmac-sha1-96

root@ex3#set security ipsec security-association ospfv3-sha-1 manual
direction bidirectional authentication key ascii-text
12345678901234567890

root@ex3#set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface
ge-0/0/14.57 interface-type p2p

root@ex3#set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface
ge-0/0/14.57 ipsec-sa ospfv3-sha-1

root@ex3#set firewall family inet6 filter protect-re-ipv6 term ah from
next-header ah

root@ex3#set firewall family inet6 filter protect-re-ipv6 term ah then
accept

root@ex3#insert firewall family inet6 filter protect-re-ipv6 term ah
before term icmp
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

The above configuration should be repeated for all links in Area 1, not forgetting the ex4 and ex2 switches. After applying the configuration, we can issue the following command to verify whether the OSPF adjacencies are still UP:

```
root@ex2# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.30      ge-0/0/14.61           Full      172.30.47.1      128
36
```

By issuing the command 'show ospf3 interface detail' we can verify what IPSec SA name is used on the interface, if any.

6) Area 2 has already been configured for both IPv4 and IPv6 earlier on.

7) Configuring a virtual link between srx4a and ex2.

When the link between ex1 and srx4a is disconnected, srx4a loses the one and only link it has in Area 0. To make sure that Area 0 does not split when this happens, we configure a virtual-link towards ex2 from srx4a. The virtual-link is always configured inside area 0. The neighbor-id is the router-id of the OSPF router at the other end of the virtual-link, and the transit Area is the area that the virtual-link traverses.

In this case, the virtual-link traverses Area 2:

```
root@ex2#set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id
172.30.15.9 transit-area 0.0.0.2
```

```
root@srx4#set routing-instances SRX4a protocols ospf3 area 0.0.0.0
virtual-link neighbor-id 172.30.15.6 transit-area 0.0.0.2
```

```
root@srx4# run show ospf3 neighbor instance SRX4a
Instance: SRX4a

Realm: ipv6-unicast
ID               Interface              State     Pri   Dead
172.30.15.5      ge-0/0/5.59            Full      128     34
  Neighbor-address fe80::3e94:d500:3b46:7311
172.30.15.6      vl-172.30.15.6         Full        0     30
  Neighbor-address fd17:f0f4:f691:2::1
172.30.15.3      ge-0/0/5.39            Full      128     31
  Neighbor-address fe80::3e8a:b000:275c:7aa4
```

8) We complete this task by configuring gre interfaces to tunnel between srx3 and srx1/srx2.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Both srx1 and srx2 are active inside Area 0. Through the establishment of an OSPF neighbor relationship across a gre-interface, we can connect srx3 to the backbone area. By placing the gre interface inside Area 3, we also turn srx1 and srx2 into ABR routers.

We start out with the configuration of Area 3 on srx3 and srx4b:

```
root@srx3#set protocols ospf3 area 0.0.0.3 interface ge-0/0/4.31
interface-type p2p
```

```
root@srx4#set routing-instances SRX4b protocols ospf3 realm ipv4-
unicast reference-bandwidth 100g
```

```
root@srx4#set routing-instances SRX4b protocols ospf3 realm ipv4-
unicast area 0.0.0.3 interface ge-0/0/5.31 interface-type p2p
```

```
root@srx4#set routing-instances SRX4b protocols ospf3 realm ipv4-
unicast area 0.0.0.3 interface lo0.2 passive
```

```
root@srx4#set routing-instances SRX4b protocols ospf3 reference-
bandwidth 100g
```

```
root@srx4#set routing-instances SRX4b protocols ospf3 area 0.0.0.3
interface lo0.2 passive
```

```
root@srx4#set routing-instances SRX4b protocols ospf3 area 0.0.0.3
interface ge-0/0/5.31 interface-type p2p
```

Verification:

```
root@srx4# run show ospf3 neighbor instance SRX4b
Instance: SRX4b
Realm: ipv6-unicast
ID              Interface             State     Pri   Dead
172.30.15.3     ge-0/0/5.31           Full      128     34
  Neighbor-address fe80::3e8a:b000:1f5c:7aa4
```

Even though Area 3 is not connected to the backbone area yet, the neighbor relationship between srx3 and srx4b will form. We then continue and configure the gre interfaces on srx1, srx2 and srx3. We also need to include these interfaces in our OSPF configuration.

Note that the gre tunnel source addresses should not be advertised into Area 2, otherwise suboptimal routing could occur. We can prevent this from happening by using the 'area-range x.x.x.x restrict' configuration command.

On srx3:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set interfaces gr-0/0/0 unit 0 description "GRE tunnel to
SRX1"

root@srx3# set interfaces gr-0/0/0 unit 0 tunnel source 172.30.15.3

root@srx3# set interfaces gr-0/0/0 unit 0 tunnel destination
172.30.15.1

root@srx3# set interfaces gr-0/0/0 unit 0 family inet address
172.30.0.130/30

root@srx3# set interfaces gr-0/0/0 unit 0 family inet6 address
fd17:f0f4:f691:3::6/126

root@srx3# set interfaces gr-0/0/0 unit 1 description "GRE tunnel to
SRX2"

root@srx3# set interfaces gr-0/0/0 unit 1 tunnel source 172.30.15.3

root@srx3# set interfaces gr-0/0/0 unit 1 tunnel destination
172.30.15.2

root@srx3# set interfaces gr-0/0/0 unit 1 family inet address
172.30.0.134/30

root@srx3# set interfaces gr-0/0/0 unit 1 family inet6 address
fd17:f0f4:f691:3::a/126

root@srx3# set protocols ospf3 realm ipv4-unicast area 0.0.0.3
interface gr-0/0/0.0 interface-type p2p

root@srx3# set protocols ospf3 realm ipv4-unicast area 0.0.0.3
interface gr-0/0/0.0 metric 1000

root@srx3# set protocols ospf3 realm ipv4-unicast area 0.0.0.3
interface gr-0/0/0.1 interface-type p2p

root@srx3# set protocols ospf3 realm ipv4-unicast area 0.0.0.3
interface gr-0/0/0.1 metric 1000

root@srx3# set protocols ospf3 area 0.0.0.3 interface gr-0/0/0.0
interface-type p2p

root@srx3# set protocols ospf3 area 0.0.0.3 interface gr-0/0/0.0
metric 1000

root@srx3# set protocols ospf3 area 0.0.0.3 interface gr-0/0/0.1
interface-type p2p
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set protocols ospf3 area 0.0.0.3 interface gr-0/0/0.1
metric 1000

root@srx3# set protocols ospf3 realm ipv4-unicast area 0.0.0.2 area-
range 172.30.15.3/32 restrict

root@srx3# set protocols ospf3 area 0.0.0.2 area-range
fd17:f0f4:f691:15::3/128 restrict
```

On srx1 and srx2:

Here, we also configure gre interfaces and include them in the OSPF configuration. To the prevent tunnel addresses from being advertised into Area 0, we again use the '**area-range x.x.x.x restrict** ' configuration command:

```
root@srx1# set interfaces gr-0/0/0 unit 0 description "GRE tunnel to
SRX3"

root@srx1# set interfaces gr-0/0/0 unit 0 tunnel source 172.30.15.1

root@srx1# set interfaces gr-0/0/0 unit 0 tunnel destination
172.30.15.3

root@srx1# set interfaces gr-0/0/0 unit 0 family inet address
172.30.0.129/30

root@srx1# set interfaces gr-0/0/0 unit 0 family inet6 address
fd17:f0f4:f691:3::5/126

root@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.3
interface gr-0/0/0.0 interface-type p2p

root@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.3
interface gr-0/0/0.0 metric 1000

root@srx1# set protocols ospf3 area 0.0.0.3 interface gr-0/0/0.0
interface-type p2p

root@srx1# set protocols ospf3 area 0.0.0.3 interface gr-0/0/0.0
metric 1000

root@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0 area-
range 172.30.15.1/32 restrict

root@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0 area-
range 172.30.15.2/32 restrict
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set protocols ospf3 area 0.0.0.0 area-range
fd17:f0f4:f691:15::1/128 restrict

root@srx1# set protocols ospf3 area 0.0.0.0 area-range
fd17:f0f4:f691:15::2/128 restrict
```

Verification:

```
root@srx1# run show ospf3 neighbor instance all
Instance: master
Realm: ipv6-unicast
ID              Interface             State     Pri   Dead
172.30.15.2     ge-0/0/4.12           Full      128     38
  Neighbor-address fe80::7a19:f700:c18:3d84
172.30.15.5     ge-0/0/4.13           Full      128     37
  Neighbor-address fe80::3e94:d500:d46:7311
172.30.15.3     gr-0/0/0.0            Full      128     34
  Neighbor-address fe80::3e8a:b000:5c:7aa0
Realm: ipv4-unicast
ID              Interface             State     Pri   Dead
172.30.15.2     ge-0/0/4.12           Full      128     33
  Neighbor-address fe80::7a19:f700:c18:3d84
172.30.15.5     ge-0/0/4.13           Full      128     36
  Neighbor-address fe80::3e94:d500:d46:7311
172.30.15.3     gr-0/0/0.0            Full      128     34
  Neighbor-address fe80::3e8a:b000:5c:7aa0
root@srx1# run show route protocol ospf 172.30.15.10

inet.0: 58 destinations, 60 routes (58 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.15.10/32    *[OSPF3/10] 01:03:12, metric 1100
                    > via gr-0/0/0.0
```

9) BFD continuity checking must be configured for all links in Area 2.

The **minimum-interval** is used to specify the interval between the BFD packets. The **multiplier** determines how many packets may be lost before the neighbor is declared down. We need to configure BFD for both realms separately:

On ex2:

```
root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.2
interface ge-0/0/14.64 bfd-liveness-detection minimum-interval 400

root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.2
interface ge-0/0/14.64 bfd-liveness-detection multiplier 3
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set protocols ospf3 area 0.0.0.2 interface ge-0/0/14.64 bfd-
liveness-detection minimum-interval 400

root@ex2# set protocols ospf3 area 0.0.0.2 interface ge-0/0/14.64 bfd-
liveness-detection multiplier 3
```

The above configuration will lead to a 1.2 second detect time. To complete this task, we also need to add a new term in the firewall filter to accept BFD packets:

```
root@ex2# set firewall family inet filter protect-re-ipv4 term bfd
from source-prefix-list direct-subnets

root@ex2# set firewall family inet filter protect-re-ipv4 term bfd
from protocol udp

root@ex2# set firewall family inet filter protect-re-ipv4 term bfd
from destination-port 3784

root@ex2# set firewall family inet filter protect-re-ipv4 term bfd
then accept

root@ex2# insert firewall family inet6 filter protect-re-ipv6 term
bfd before term icmp
```

We should apply a similar configuration to every device in Area 2.

Verification:

```
root@srx4# run show bfd session brief
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval
Multiplier
fe80::2e6b:f500:408f:87d1 Up        ge-0/0/4.64    1.200     0.400
3
fe80::3e8a:b000:225c:7aa4 Up        ge-0/0/4.34    1.200     0.400
3
fe80::3e8a:b000:275c:7aa4 Up        ge-0/0/5.39    1.200     0.400
3

3 sessions, 6 clients
Cumulative transmit rate 7.5 pps, cumulative receive rate 7.5 pps
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

10) Area 1 should contain LSA Type 3, 4 or 5.

The Area type we are looking for here is an NSSA with the '**no-summary**' option. The area was already configured as an NSSA earlier on, so in this case we can focus on the new objectives.

We configure ex2 and ex3, the ABR routers, not to send any summaries into the area using the '**nssa no-summaries**' configuration command. The ex3 and ex4 switch will still need to be able to reach the rest of the network, so a default route is injected through the '**default-lsa default-metric 10**' configuration command. The option '**type-7**' was configured to make sure no type 3 LSA is injected into the Area.

On ex2 and ex3:

```
root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.1 nssa
default-lsa default-metric 10

root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.1 nssa
default-lsa type-7

root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.1 nssa no-
summaries
```

Verification:

```
root@ex3# run show route protocol ospf
inet.0: 32 destinations, 48 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0          *[OSPF3/150] 01:06:48, metric 110, tag 0
                    > to 172.30.0.33 via ge-0/0/14.57
```

11) To complete this task, we need to write a route policy that exports direct links and RIP routes to OSPF.

Since this is an NSSA area, ASBR routers will automatically inject the external route information using type 7 LSAs. The ABRs of the NSSA, in this case ex1 and ex2, will translate the type 7 LSAs into type 5 LSAs.

On ex3 and ex4:

```
root@ex3# set policy-options policy-statement to.ospf term 1 from
protocol rip

root@ex3# set policy-options policy-statement to.ospf term 1 then
accept

root@ex3# set policy-options policy-statement to.ospf term 2 from
protocol direct
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3# set policy-options policy-statement to.ospf term 2 from
route-filter 172.30.0.48/30 exact

root@ex3# set policy-options policy-statement to.ospf term 2 then
accept
```

Since the routes we aim to export to OSPF are IPv4 only, we configure the export policy for the '**ipv4-unicast**' realm:

```
root@ex3# set protocols ospf3 realm ipv4-unicast export to.ospf
```

Verification:

```
root@ex2# run show route 172.30.0.48/30
inet.0: 76 destinations, 92 routes (76 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
172.30.0.48/30     *[OSPF3/150] 01:22:46, metric 0, tag 0
                    > to 172.30.0.38 via ge-0/0/14.68
```

The above command shows us that the route is present on the ex2 device. A deeper analysis on what routes are being advertised by either ex3 and ex4 could be done by issuing the '**show ospf3 database advertising-router self realm ipv4-unicast extensive**' command. Here, we can see exactly what route is advertised and what the LSA type for each route is.

Instead of specifying '**self**', we can also specify a router-id to look at what another router in the area has advertised.

12) To summarize routes from Area 1 into Area 0, we should configure the ABRs (ex1 and ex2) using the **area-range** command.

Additionally, we need more specific routes to be suppressed. We can achieve this by adding the **restrict** option to the '**nssa area-range**' configuration.

On both ABRs (ex1 and ex2):

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 area-range
172.30.0.32/27

set protocols ospf3 realm ipv4-unicast area 0.0.0.1 nssa area-range
172.30.0.48/29 restrict
```

Verification:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# run show route protocol ospf 172.30.0.32/27
inet.0: 58 destinations, 60 routes (58 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
172.30.0.32/27     *[OSPF3/10] 01:32:31, metric 400
                    > to 172.30.0.6 via ge-0/0/4.13
```

13) We can check all OSPF adjacencies by issuing a **show ospf neighbor** command on every device. All adjacencies should list the state 'Full'.

```
root@srx1# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
172.30.15.2     ge-0/0/4.12           Full      128     31
  Neighbor-address fe80::7a19:f700:c18:3d84
172.30.15.5     ge-0/0/4.13           Full      128     32
  Neighbor-address fe80::3e94:d500:d46:7311
172.30.15.3     gr-0/0/0.0            Full      128     37
  Neighbor-address fe80::3e8a:b000:5c:7aa0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 3: External Configuration and Redistribution Policies

1) RIPng configuration to DC2:

The configuration of RIPng is pretty simple and straightforward. We configure a group name and add the interface on which we want to run RIPng to that group. In this example, we name the group 'DC2':

```
root@srx4# set protocols ripng group DC2 neighbor ge-0/0/4.41
```

```
root@srx4# run show ripng neighbor instance all
Instance: master
                    Source                                Dest
In
Neighbor      State  Address                               Address  Send
Recv Met
--------      -----  -------                               -------  ---- ---
- ---
ge-0/0/4.41      Up fe80::3e61:400:298d:be84              ff02::9  yes yes
1

root@srx4# run show route protocol ripng
inet6.0: 58 destinations, 61 routes (58 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
fd17:f0f4:f691:f::/80
                   *[RIPng/100] 01:33:58, metric 2, tag 0
                    > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:1::/80
                   *[RIPng/100] 01:33:58, metric 2, tag 0
                    > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:2::/80
                   *[RIPng/100] 01:33:58, metric 2, tag 0
                    > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:3::/80
                   *[RIPng/100] 01:33:58, metric 2, tag 0
                    > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:4::/80
                   *[RIPng/100] 01:33:58, metric 2, tag 0
                    > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:5::/80
                   *[RIPng/100] 01:33:58, metric 2, tag 0
                    > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:6::/80
                   *[RIPng/100] 01:33:58, metric 2, tag 0
                    > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:7::/80
                   *[RIPng/100] 01:33:58, metric 2, tag 0
                    > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:8::/80
                   *[RIPng/100] 01:33:58, metric 2, tag 0
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                     > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:9::/80
                    *[RIPng/100] 01:33:58, metric 2, tag 0
                     > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:a::/80
                    *[RIPng/100] 01:33:58, metric 2, tag 0
                     > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:b::/80
                    *[RIPng/100] 01:33:58, metric 2, tag 0
                     > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:c::/80
                    *[RIPng/100] 01:33:58, metric 2, tag 0
                     > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:d::/80
                    *[RIPng/100] 01:33:58, metric 2, tag 0
                     > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:e::/80
                    *[RIPng/100] 01:33:58, metric 2, tag 0
                     > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
fd17:f0f4:f691:f:f::/80
                    *[RIPng/100] 01:33:58, metric 2, tag 0
                     > to fe80::21f:1200:292e:6f81 via ge-0/0/4.41
ff02::9/128         *[RIPng/100] 01:56:41, metric 1
                       MultiRecv
```

2) Ensure that srx4 advertises a single type 5 LSA that summarizes all RIPng prefixes received from DC2:

If we were to simply export RIPng into OSPF, an external type 5 LSA would have been created per RIPng route. For this assignment, we advertise a single summary route of the RIPng routes to OSPF by redistributing an aggregate route. The RIPng routes will act as contributing routes to the aggregate. By exporting the aggregate to OSPF, we manage to advertise DC2 as a single type 5 LSA:

```
root@srx4# set routing-options rib inet6.0 aggregate route
fd17:f0f4:f691:f::/64

root@srx4# set policy-options policy-statement to.ospfv3 term 1 from
protocol aggregate

root@srx4# set policy-options policy-statement to.ospfv3 term 1 from
route-filter fd17:f0f4:f691:f::/64 exact

root@srx4# set policy-options policy-statement to.ospfv3 term 1 then
accept

root@srx4# set policy-options policy-statement to.ospfv3 term 2 then
reject

root@srx4# set protocols ospf3 export to.ospfv3
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification:

```
root@srx1# run show route table inet6.0 fd17:f0f4:f691:f::/64 exact
fd17:f0f4:f691:f::/64
                   *[OSPF3/150] 05:39:27, metric 0, tag 0
                    > to fe80::7a19:f700:c18:3d84 via ge-0/0/4.12
                      to fe80::3e94:d500:d46:7311 via ge-0/0/4.13
```

3) Have srx4 advertise a single IPv6 route that summarizes all internal prefixes to DC2:

Here we basically perform the same steps that we executed in the previous task. We summarize the internal prefixes using an aggregate route and then we send a single summary into RIPng.

Even though JUNOS imports all RIPng routes by default, it does not export any. For this assignment, we configure an aggregate IPv6 route that represents a summary of the internal prefixes. Then, we write an export policy to have the route exported to RIPng:

```
root@srx4# set routing-options rib inet6.0 aggregate route
fd17:f0f4:f691::/48

root@srx4# set policy-options policy-statement ripng.out term 1 from
protocol aggregate

root@srx4# set policy-options policy-statement ripng.out term 1 from
route-filter fd17:f0f4:f691::/48 exact

root@srx4# set policy-options policy-statement ripng.out term 1 then
accept

root@srx4# set protocols ripng group DC2 export ripng.out
```

Verification:

```
root@srx4# run show interfaces ge-0/0/4.41
  Logical interface ge-0/0/4.41 (Index 81) (SNMP ifIndex 731)
    Description: L3 connection to DC2
    Flags: SNMP-Traps 0x0 VLAN-Tag [ 0x8100.41 ]  Encapsulation: ENET2
    Input packets : 247
    Output packets: 259
    Security: Zone: Null
    Protocol inet6, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: fd17:f0f4:f691:2::c/126, Local:
fd17:f0f4:f691:2::d
      Addresses, Flags: Is-Preferred
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
        Destination: fe80::/64, Local: fe80::3e61:400:298d:be84

root@srx4#run show route advertising-protocol ripng
fe80::3e61:400:298d:be84

inet6.0: 58 destinations, 61 routes (58 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691::/48*[Aggregate/130] 02:10:35
                    Reject
```

Keep in mind for verification you have to use link local address in the above command.

4) Enable OSPFv2 on the srx2 and ex2 interfaces to DC1. The Area ID is now known. Ensure that 16 IPv4 prefixes are received from DC1.

We do not know the OSPF Area ID so we simply configure Area 0:

```
root@srx4#set protocols ospf area 0.0.0.0 interface ge-0/0/4.25
interface-type p2p

root@srx2# run show ospf neighbor

[edit]
```

After committing the configuration, the router will attempt to form an OSPF neighbor relationship across the link. Since we simply guessed the Area ID, most likely it will not match with the one that the neighbor is using. Because of this, an OSPF neighbor relationship is highly unlikely for form.

What we can get out of this is insight into what the neighboring router would be expecting. We can see this after we enable traceoptions and investigate the errors that show up in the log:

```
root@srx2# set protocols ospf traceoptions flag error detail

root@srx2# set protocols ospf traceoptions flag hello detail
```

After a little while, we can examine the messages that the router has placed inside the log file:

```
root@srx2# run show log OSPF.log
Dec 21 10:51:39 trace_on: Tracing to "/var/log/OSPF.log" started
Dec 21 10:51:47.798575   checksum 0xffa, authtype 0
Dec 21 10:51:47.798628   mask 255.255.255.252, hello ivl 10, opts 0x12, prio
128
Dec 21 10:51:47.798694   dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0
Dec 21 10:51:48.576505 OSPF periodic xmit from 172.30.0.25 to 224.0.0.5 (IFL
83 area 0.0.0.0)
Dec 21 10:51:56.285638 OSPF periodic xmit from 172.30.0.25 to 224.0.0.5 (IFL
83 area 0.0.0.0)
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Dec 21 10:51:56.468942 OSPF packet ignored: area mismatch (0.0.0.10) from
172.30.0.26 on intf ge-0/0/4.25 area 0.0.0.0
Dec 21 10:51:56.469055 OSPF rcvd Hello 172.30.0.26 -> 224.0.0.5 (ge-0/0/4.25
IFL 83 area 0.0.0.0)
Dec 21 10:51:56.469125   Version 2, length 44, ID 172.30.47.1, area 0.0.0.10
Dec 21 10:51:56.469175   checksum 0xffa, authtype 0
Dec 21 10:51:56.469243   mask 255.255.255.252, hello_ivl 10, opts 0x12, prio
128
```

We have found out that the neighboring router is sending out OSPF Hellos that contain Area ID 0.0.0.10. Now we simply change the Area ID like this:

```
root@srx2# rename protocols ospf area 0.0.0.0 to area 0.0.0.10
```

```
root@srx2# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.26      ge-0/0/4.25            Full      172.30.47.1      128
35
```

5) Make sure that all internal routers use optimal routing to reach the DC1 prefixes.

In this task, we have to export routes from DC1 to OSPF. By default, when we export routing information into OSPF, the information is exported as an external type 2 route. This means that the metric is never incremented. To ensure optimal routing, we change the external route to a type 1:

```
root@srx2# set policy-options policy-statement from.ospfv2.to.ospfv3
term 1 from protocol ospf

root@srx2# set policy-options policy-statement from.ospfv2.to.ospfv3
term 1 from route-filter 172.31.32.0/20 orlonger

root@srx2# set policy-options policy-statement from.ospfv2.to.ospfv3
term 1 then external type 1

root@srx2# set policy-options policy-statement from.ospfv2.to.ospfv3
term 1 then accept

root@srx2# set policy-options policy-statement from.ospfv2.to.ospfv3
term 3 then reject

root@srx2# set protocols ospf3 realm ipv4-unicast export
from.ospfv2.to.ospfv3
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification:

```
root@srx2#run show route 172.31.32.0/20
172.31.32.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.33.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.34.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.35.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.36.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.37.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.38.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.39.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.40.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.41.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.42.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.43.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.44.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.45.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.46.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
172.31.47.0/24     *[OSPF3/150] 03:36:47, metric 100, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
```

6) Advertise all routes from srx2 and ex2 to DC1.

For this task, we need to export all internal and direct routes to DC1 using an export policy. We also need to make sure that the routes we advertise are not advertised back into our domain by DC1.

First, we create a policy that exports the internal OSPF and direct routes:

On srx2 and ex2:

```
root@srx2# set policy-options policy-statement from.ospfv3.to.ospfv2
term 1 from protocol ospf3
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set policy-options policy-statement from.ospfv3.to.ospfv2
term 1 from route-filter 172.30.0.0/24 orlonger

root@srx2# set policy-options policy-statement from.ospfv3.to.ospfv2
term 1 then accept

root@srx2# set policy-options policy-statement from.ospfv3.to.ospfv2
term 2 from protocol direct

root@srx2# set policy-options policy-statement from.ospfv3.to.ospfv2
term 2 from route-filter 172.30.0.0/24 orlonger

root@srx2# set policy-options policy-statement from.ospfv3.to.ospfv2
term 2 then accept

root@srx2# set policy-options policy-statement from.ospfv3.to.ospfv2
term 3 then reject
```

We then configure this policy as an export policy:

```
root@srx2# set protocols ospf export from.ospfv3.to.ospfv2
```

In order to make sure that we do not accept any of the routes when they are advertise back into our domain, we create another policy. This policy should match all the routes we previously exported:

```
root@srx2# set policy-options policy-statement filter.ospfv2 term 1
from route-filter 172.30.0.0/24 orlonger

root@srx2# set policy-options policy-statement filter.ospfv2 term 1
then reject
```

By applying this policy as an import policy, we can filter out the any of the routes we advertised to DC1:

```
root@srx2# set protocols ospf import filter.ospfv2
```

Verification can be done by checking the routing-table on the VR.

Verification:

```
root@vrdevice# run show route table DC1.inet.0 protocol ospf
root@vrdevice# run show route table DC1.inet.0
DC1.inet.0: 40 destinations, 40 routes (40 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

172.30.0.0/30      *[OSPF/150] 01:13:47, metric 0, tag 0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                    > to 172.30.0.25 via ge-0/0/1.25
172.30.0.4/30       *[OSPF/150] 01:13:47, metric 200, tag 0
                      to 172.30.0.25 via ge-0/0/1.25
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.8/30       *[OSPF/150] 01:13:47, metric 0, tag 0
                    > to 172.30.0.25 via ge-0/0/1.25
                      to 172.30.0.29 via ge-0/0/1.61
172.30.0.12/30      *[OSPF/150] 01:13:47, metric 0, tag 0
                    > to 172.30.0.25 via ge-0/0/1.25
172.30.0.16/30      *[OSPF/150] 03:13:41, metric 0, tag 0
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.20/30      *[OSPF/150] 01:13:47, metric 200, tag 0
                    > to 172.30.0.25 via ge-0/0/1.25
                      to 172.30.0.29 via ge-0/0/1.61
172.30.0.32/27      *[OSPF/150] 01:13:47, metric 400, tag 0
                    > to 172.30.0.25 via ge-0/0/1.25
172.30.0.32/30      *[OSPF/150] 03:13:41, metric 300, tag 0
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.36/30      *[OSPF/150] 03:13:41, metric 0, tag 0
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.40/30      *[OSPF/150] 03:13:41, metric 200, tag 0
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.48/30      *[OSPF/150] 03:13:41, metric 0, tag 0
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.52/30      *[OSPF/150] 03:13:41, metric 0, tag 0
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.64/30      *[OSPF/150] 01:13:47, metric 1100, tag 0
                    > to 172.30.0.25 via ge-0/0/1.25
172.30.0.96/30      *[OSPF/150] 03:13:41, metric 0, tag 0
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.100/30     *[OSPF/150] 03:13:41, metric 200, tag 0
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.104/30     *[OSPF/150] 01:13:47, metric 300, tag 0
                      to 172.30.0.25 via ge-0/0/1.25
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.128/30     *[OSPF/150] 03:13:41, metric 1200, tag 0
                    > to 172.30.0.29 via ge-0/0/1.61
172.30.0.132/30     *[OSPF/150] 01:13:47, metric 0, tag 0
                    > to 172.30.0.25 via ge-0/0/1.25
224.0.0.5/32        *[OSPF/10] 03:14:55, metric 1
                     MultiRecv
```

7) DC1 IPv4 prefixes should not be present on srx3 and srx4 and DC1 IPv4 prefixes should not be present in the Area 1 LSDB.

To complete this task, task we filter DC1 and DC3 on srx3 and srx4. We filter the routes by using an import policy. In the second part of the configuration displayed below, we prevent translating LSAs on ex2.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set policy-options policy-statement filter.dc1.v4 term 1
from route-filter 172.31.32.0/20 orlonger

root@srx3# set policy-options policy-statement filter.dc1.v4 term 1
then reject
```

On ex2:

```
root@ex2# set protocols ospf3 realm ipv4-unicast no-nssa-abr
```

Verification:

```
root@srx4# run show route protocol ospf3 table inet.0
inet.0: 32 destinations, 32 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF3/150] 01:39:34, metric 200, tag 0
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.0.0/30      *[OSPF3/10] 01:39:34, metric 300
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.0.4/30      *[OSPF3/10] 01:39:26, metric 300
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.0.8/30      *[OSPF3/10] 01:39:34, metric 200
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.0.12/30     *[OSPF3/10] 01:39:34, metric 300
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.0.16/30     *[OSPF3/10] 01:39:34, metric 200
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.0.20/30     *[OSPF3/10] 01:39:21, metric 300
                      to 172.30.0.101 via ge-0/0/4.34
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.0.32/27     *[OSPF3/10] 01:39:32, metric 400
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.0.64/30     *[OSPF3/10] 01:39:26, metric 200
                    > to 172.30.0.101 via ge-0/0/4.34
172.30.0.104/30    *[OSPF3/10] 01:39:26, metric 200
                    > to 172.30.0.101 via ge-0/0/4.34
172.30.0.128/30    *[OSPF3/10] 01:39:21, metric 1100
                    > to 172.30.0.101 via ge-0/0/4.34
172.30.0.132/30    *[OSPF3/10] 01:39:21, metric 1100
                    > to 172.30.0.101 via ge-0/0/4.34
172.30.15.1/32     *[OSPF3/10] 01:39:34, metric 300
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.15.2/32     *[OSPF3/10] 01:39:34, metric 200
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.15.3/32     *[OSPF3/10] 01:39:26, metric 100
                    > to 172.30.0.101 via ge-0/0/4.34
172.30.15.5/32     *[OSPF3/10] 01:39:26, metric 200
                    > to 172.30.0.97 via ge-0/0/4.64
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
172.30.15.6/32     *[OSPF3/10] 01:39:34, metric 100
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.15.7/32     *[OSPF3/10] 01:39:32, metric 300
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.15.8/32     *[OSPF3/10] 01:39:34, metric 200
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.15.9/32     *[OSPF3/10] 01:39:21, metric 200
                    > to 172.30.0.101 via ge-0/0/4.34
172.30.15.10/32    *[OSPF3/10] 01:39:26, metric 200
                    > to 172.30.0.101 via ge-0/0/4.34
172.30.80.0/30     *[OSPF3/10] 01:39:34, metric 400
                    > to 172.30.0.97 via ge-0/0/4.64
172.30.80.4/30     *[OSPF3/10] 01:39:34, metric 300
                    > to 172.30.0.97 via ge-0/0/4.64

root@ex3# run show route protocol ospf
inet.0: 32 destinations, 48 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF3/150] 01:06:48, metric 110, tag 0
                    > to 172.30.0.33 via ge-0/0/14.57
172.30.0.36/30     *[OSPF3/10] 01:06:36, metric 200
                    > to 172.30.0.42 via ge-0/0/14.78
172.30.0.52/30     *[OSPF3/150] 01:06:36, metric 0, tag 0
                    > to 172.30.0.42 via ge-0/0/14.78
172.30.15.8/32     *[OSPF3/10] 01:06:36, metric 100
                    > to 172.30.0.42 via ge-0/0/14.78
```

8) Enable RIP on ex3 and ex4:

EX3:

```
root@ex3# set protocols rip group DC3 neighbor ge-0/0/14.71
```

EX4:

```
root@ex4# set protocols rip group DC3 neighbor ge-0/0/14.81
```

Verification:

```
root@ex3# run show rip neighbor
                    Local  Source           Destination      Send
Receive   In
Neighbor            State  Address          Address          Mode   Mode
Met
--------            -----  -------          -----------      ----   ------
-  ---
ge-0/0/14.71          Up 172.30.0.49      224.0.0.9        mcast  both
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
1
root@ex3# run show route protocol rip

inet.0: 32 destinations, 48 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.1.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.2.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.3.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.4.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.5.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.6.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.7.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.8.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.9.0/24      *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.10.0/24     *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.11.0/24     *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.12.0/24     *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.13.0/24     *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.14.0/24     *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
172.31.15.0/24     *[RIP/100] 04:15:08, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/14.71
224.0.0.9/32       *[RIP/100] 04:21:15, metric 1
                      MultiRecv
```

9) Ensure that a single type 5 LSA, summarizing DC3 prefixes, is advertised to Area 0.

First, we export DC3 prefixes using an export policy that we configure for OSPF.

On ex3 and ex4:

```
root@ex3# set policy-options policy-statement to.ospf term 1 from
protocol rip
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3# set policy-options policy-statement to.ospf term 1 then
accept

root@ex3# set policy-options policy-statement to.ospf term 2 from
protocol direct

root@ex3# set policy-options policy-statement to.ospf term 2 from
route-filter 172.30.0.48/30 exact

root@ex3# set policy-options policy-statement to.ospf term 2 then
accept

root@ex3# set protocols ospf3 realm ipv4-unicast export to.ospf
```

After that, we summarize the routes using the **area-range** command on the ABRs. Don't forget that this is an NSSA Area. Because of this, we need to specify the **nssa area-range** command like this:

On ex1 and ex2:

```
root@ex1# set protocols ospf3 realm ipv4-unicast area 0.0.0.1 nssa
area-range 172.31.0.0/20
```

Verification:

```
root@srx1# run show route table inet.0 172.31.0.0/20 exact
172.31.0.0/20      *[OSPF3/150] 05:36:46, metric 3, tag 0
                    > to 172.30.0.2 via ge-0/0/4.12
                      to 172.30.0.6 via ge-0/0/4.13
```

10) Advertise a default route from both Area 1 ASBRs to DC for redundancy.

This task is combined with task 11.

11) Ensure that ex3 and ex4 use optimal routing to reach OSPF destinations outside Area 1.

We export the default route from OSPF to RIP using an export policy. We run the risk of learning the default route from the RIP domain again. To ensure optimal routing, we mark the routes we advertise with tag 100. We then filter the routes we marked through the use of an import policy:

On ex3 and ex4:

```
root@ex3# set policy-options policy-statement to.rip term 1 from
protocol ospf

root@ex3# set policy-options policy-statement to.rip term 1 from
route-filter 0.0.0.0/0 exact
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3# set policy-options policy-statement to.rip term 1 then tag
100

root@ex3# set policy-options policy-statement to.rip term 1 then
accept

root@ex3# set policy-options policy-statement from.rip term no.tag100
from protocol rip

root@ex3# set policy-options policy-statement from.rip term no.tag100
from tag 100

root@ex3# set policy-options policy-statement from.rip term no.tag100
then reject

root@ex3# set policy-options policy-statement from.rip term 1 from
protocol rip

root@ex3# set policy-options policy-statement from.rip term 1 then
accept


root@ex3# set protocols rip group DC3 export to.rip

root@ex3# set protocols rip group DC3 import from.rip
```

Verification:

```
root@vrdevice# run show route table DC3.inet.0

DC3.inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0          *[RIP/100] 03:55:27, metric 2, tag 100
                      to 172.30.0.49 via ge-0/0/1.71
                    > to 172.30.0.53 via ge-0/0/1.81
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 4: Protocol Independent Routing and Policies

1) Configure static routes to ISP1 and ISP2 on srx1 and srx2:

```
root@srx1# set routing-options rib inet6.0 static route
fc09:c0:ffee:abed::/64 next-hop fd17:f0f4:f691:80::1
```

```
root@srx2# set routing-options static route 192.168.168.0/24 next-hop
172.30.80.1
```

2) Ensure the advertisement of a conditionalized static.
3) Make sure the cost for the advertised default route will include the cost of the internal links.

To complete task 2 and 3, we need to configure a generate route. A generate route can be used as a route of last resort. The generate route is active for as long as contributing routes exist. By default, a contributing route is an active and more specific match of the generate route itself. We can control what routes are allowed to be a contributing route through the configuration of a routing policy.

The generate route is very similar to the aggregate route, the generate route even shows up in the routing table listed as an aggregate route. The main different though, is that a generate route will use the next hop of its primary contributing route.

For this assignment, we configure the policy in such a way that only the previously configured static routes can become contributing routes. We do this by specifically matching the static and setting the last term in this policy reject.

On srx1 and srx2:

```
root@srx1# set routing-options generate route 0.0.0.0/0 policy
contribute.filter
```

```
root@srx1# set routing-options rib inet6.0 generate route ::/0 policy
contribute.filter.v6
```

```
root@srx1# set policy-options policy-statement contribute.filter term
1 from route-filter 192.168.168.0/24 exact
```

```
root@srx1# set policy-options policy-statement contribute.filter term
1 then accept
```

```
root@srx1# set policy-options policy-statement contribute.filter term
2 then reject
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set policy-options policy-statement contribute.filter.v6
term 1 from route-filter fc09:c0:ffee:abed::/64 exact

root@srx1# set policy-options policy-statement contribute.filter.v6
term 1 then accept

root@srx1# set policy-options policy-statement contribute.filter.v6
term 2 then reject
```

To be able to export the generate to OSPF, we need to configure a second routing policy. In this policy, we need to match the previously configured generate route. We can do this by using the '**from protocol aggregate**' match condition. We also need to change the OSPF external to type 1. This way, the metric of the route will also include the costs of the internal links.

On srx1:

```
root@srx1# set policy-options policy-statement default.to.ospf term 1
from protocol aggregate

root@srx1# set policy-options policy-statement default.to.ospf term 1
from route-filter 0.0.0.0/0 exact

root@srx1# set policy-options policy-statement default.to.ospf term 1
then external type 1

root@srx1# set policy-options policy-statement default.to.ospf term 1
then accept

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 from protocol aggregate

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 from rib inet6.0

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 from route-filter ::/0 exact

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 then external type 1

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 then accept

root@srx1# set protocols ospf3 realm ipv4-unicast export
default.to.ospf

root@srx1# set protocols ospf3 export default.to.ospf.v6
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On srx2:

```
root@srx1# set policy-options policy-statement from.ospfv2.to.ospfv3
term 2 from protocol aggregate

root@srx1# set policy-options policy-statement from.ospfv2.to.ospfv3
term 2 from route-filter 0.0.0.0/0 exact

root@srx1# set policy-options policy-statement from.ospfv2.to.ospfv3
term 2 then external type 1

root@srx1# set policy-options policy-statement from.ospfv2.to.ospfv3
term 2 then accept

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 from protocol aggregate

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 from rib inet6.0

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 from route-filter ::/0 exact

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 then external type 1

root@srx1# set policy-options policy-statement default.to.ospf.v6 term
1 then accept

root@srx1# set protocols ospf3 export default.to.ospf.v6
```

In the routing table, we can see the generate route appear as an 'Aggregate' route. Take a look at the next-hop of the route as well, it is set to the statics next-hop IP address:

Verification:

```
root@srx1# run show route 0/0 exact

inet.0: 58 destinations, 60 routes (58 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 06:25:53
                    > to 172.30.80.1 via ge-0/0/4.14
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

4) To ensure the default gateway is used to reach the loopback addresses of SRX4b in Area 0, we should use the **restrict** option with the **area-range** command on srx1 and srx2:

```
root@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.3 area-
range 172.30.15.10/32 restrict

root@srx1# set protocols ospf3 area 0.0.0.3 area-range
fd17:f0f4:f691:15::a/128 restrict
```

Verification:

```
root@ex1# run show route protocol ospf 172.30.15.10

inet.0: 72 destinations, 72 routes (72 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF3/150] 06:08:23, metric 100, tag 0
                    > to 172.30.0.5 via ge-0/0/14.13
                      to 172.30.0.13 via ge-0/0/14.23
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Chapter 4: BGP Routing

A JNCIE-ENT candidate should be familiar with the following topics:

- BGP configuration and implementation
- BGP route selection
- 4 Byte AS number
- Route policy
- AS-path and community regular expressions
- Monitoring and troubleshooting

BGP overview.

BGP is an Exterior Gateway Protocol used to exchange routing-information between Autonomous Systems (AS). An AS is a set of devices under a common administration and each AS represents a different domain on the Internet.

BGP is a path-vector protocol and uses the AS path to evaluate or gauge the distance to a destination prefix. The AS path is also the method BGP uses to prevent routing-loops. If a BGP speaking router encounters its own path in the AS path, the route is not eligible for use.

BGP uses TCP (port 179) to provide for reliable communication. BGP neighbors are explicitly configured and as soon as a neighbor relationship is set up, the BGP peers start exchanging Network Layer Reachability Information (NLRI).

There are two different types of BGP neighbor relationships. The peering session between routers that belong to a different AS is called an External BGP (EBGP) peering session. The peering session between routers belonging to the same AS peer is called an Internal BGP (IBGP) peering session. The rules for prefix advertisements differ for EBGP and IBGP peers:

- IBGP peers advertise routes received from EBGP peers to other IBGP peers.
- EBGP peers advertise routes learned from IBGP or EBGP peers to other EBGP peers.
- IBGP peers do not advertise routes learned from IBGP peers to other IBGP peers.

That last rule is needed to prevent loops inside an AS. This is because the AS path is only updated across EBGP peering sessions. Because IBGP peers do not re-advertise IBGP learned routes, the IBGP neighbors should be fully meshed. Scaling methods to circumvent this requirement include route-reflection and confederation (neither are a topic on the exam).

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

The AS path is not the only attribute used by BGP. The attributes are grouped into four main classes:

- **Well-known mandatory**: must be supported by all BGP neighbors and must be included in all BGP updates.
- **Well-known discretionary**: must be supported by all BGP neighbors, do not have to be included in all updates.
- **Optional transitive**: support is not mandatory. When the attribute is encountered and it is not recognized, it must be passed to other neighbors.
- **Optional non-transitive**: support is not mandatory. When the attribute is encountered and it is not recognized, it can be ignored.

The following attributes are some of the most used attributes in BGP Enterprise environments:

- **Next-Hop**: this a well-known mandatory attribute. It represents the IP address of the advertising router. For a router to install a BGP route into the routing-table, this IP address must be reachable. EBGP peers update this attribute automatically to the IP address of their outgoing interface. IBGP peers do not update this attribute automatically, but they can be configured to do so.
- **AS Path**: well-known mandatory attribute that represents the path to the source of the prefix. The AS-path is a list of autonomous-systems. BGP also uses the AS-path to prevent loops. The AS-path is updated automatically across EBGP sessions. The router that sends the routing-update prepends its own AS to the list before updating its neighbor.
- **Origin**: well-known mandatory attribute. The origin describes the way in which the route was originally injected into BGP. The values can be IGP (I), EGP (E) and Incomplete (?). By default, JUNOS will set this value to IGP.
- **Local Preference**: Well-known discretionary attribute that is used to define the most preferable way out of the local AS. The route with the highest Local Preference wins. The local-preference attribute is only send across IBGP sessions. the default value is 100.
- **Multi-Exit Discriminator**: optional non-transitive attribute, used to influence the preferred path into the AS when there are multiple links between the two ASs. After the MED is send across the EBGP session into the neighboring AS, that AS will not advertise it further across its EBGP sessions. When the MED is absent, the Junos OS treats it as though the value was set to 0.
- **Community**: optional transitive attribute that is used to tag a prefix or a group of prefixes. After tagging them, they can be easily identified later on and they can help simplify policies applied to routes. Even between different ASs.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

BGP configuration and implementation

IBGP sessions are usually sourced from the loopback interface, which is specified using the **local-address** configuration command. In case of IBGP, we can either configure the group/neighbor to be **internal**, or we can configure the group/neighbor with the **peer-as** statement. In addition to this, we also need to specify the AS of the system itself:

set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.0.0.1
set protocols bgp group ibgp neighbor 10.0.0.2
set protocols bgp group ibgp neighbor 10.0.0.3
set protocols bgp group ibgp neighbor 10.0.0.4

set routing-options autonomous-system 65000

After setting the AS in the IBGP configuration, adding a standard EBGP peer would require the following configuration:

set protocols bgp group ebgp neighbor 192.168.1.2 peer-as 65001

The EBGP peer will set the next-hop attribute to the IP address that it used to source the EBGP session. To be able to use the routes that are advertised by this peer, the rest of the network needs to be able to resolve this next-hop IP address. This can be achieved in one of several ways:

- Use a next-hop self policy towards the IBGP neighbors
- Export the subnet used for the EBGP peering into the IGP through the use of a policy
- Configure the interface used for EBGP as a passive interface under the IGP
- By configuring a static route to the EBGP next-hop on each IBGP router
- Configuring an IGP adjacency with the EBGP peer

The solution that is most-often used (and recommended) is to apply a next-hop self policy on the IBGP sessions. IBGP sessions do not change the next-hop IP address by default. Through the use of a policy, you can change this behavior. When the IBGP session are sourced from the loopback interface, this IP address will (probably) be redistributed into the IGP already. This way, the other IBGP neighbors will not have any problems resolving the next-hop. To implement this solution, you can use the following policy and export it to the IBGP:

set policy-options policy-statement *nhs* term 1 then next-hop self
set protocols bgp group ibgp export *nhs*

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Load sharing across multiple BGP sessions can be achieved using the **multipath** configuration command. In case both EBGP speakers are using the loopback IP address to source the EBGP session, don't forget the **multihop** feature to overcome BGP's default TTL of 1. IP connectivity towards the loopback of the other BGP speaker can be provided through the configuration of a static route or by creating an IGP adjacency between the routers and injecting the BGP speakers' loopback IP address. The **local-address** command should be used to source the local EBGP session using the loopback IP.

Both methods of load sharing, multiple sessions or a single multihop session, require the configuration of a policy that is applied to the forwarding table:

set policy-options policy-statement *load-balancing* then load-balance per-packet

set routing-options forwarding-table export *load-balancing*

In some situations, especially when we use a link to share IPv6 routes over an IPv4 BGP session, we need to accept a remote next hop that does not belong to the same subnet. In such cases, we can use the **accept-remote-nexthop** configuration knob.

Note that we cannot configure **multihop** and **accept-remote-nexthop** at the same time.

Some common configuration items related to BGP you should know of include the following:

set protocols bgp log-updown

Configure the system to generate a logging message for every BGP state transition. This is also configurable per-group.

set protocols bgp group xx authentication-key workbook

Configures all peering session within the group with an MD5 authentication key.

set protocols bgp group xx hold-time 100

Specifies the advertised hold-time value. The default value is 90. The hold-time is three times the interval at which keepalives will be send. The lowest hold-time of the two peers will be used.

set protocols bgp group xx family inet unicast

Configures the BGP peers for the inet unicast address-family. This address family is active by default, but has to be specified as soon as another address family is configured.

set protocols bgp group xx family inet6 unicast

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Configures the BGP peers for the inet6 unicast address family.

BGP route selection.

BGP determines the best path based on the following comprehensive route-selection procedure:

1. lowest **route preference**
2. highest **local-preference**
3. shortest **AS-PATH**
4. lowest **origin value** (IGP over EGP over incomplete)
5. lowest **multiple exit discriminator** (no MED is treated as 0)
    a. checked only for routes from same AS, changed with 'always-compare-med'
6. **EBGP** over **IBGP**. When there are only EBGP routes, skip the next step.
7. lowest **IGP** cost towards BGP next-hop
8. Differs for internal and external routes:
    a. For internal routes: prefer the route with the lowest router-id
    b. For external routes: select the oldest (or currently) active route
9. (when route-reflection is used) shortest cluster list
10. lowest peer IP address

4 Byte AS number

AS numbers are assigned by the Regional Internet Registers. Initially, they were 16-bits (2-bytes). In anticipation of the 2-byte AS numbers' depletion, 32-bits (or 4-bytes) AS numbers were introduced. These 4-byte AS number are made backwards compatible with the old 2-byte AS numbers. The way in which BGP provides for this backwards compatibility is described in RFC 4893.

In that RFC, two new optional transitive attributes are defined, the AS4_PATH and the AS4_AGGREGATOR. Both these attributes were developed to make sure that 4-byte AS information (AS path and aggregates) could be propagated across networks containing 'old' BGP speakers that do not offer support for a 4-byte AS. The RFC also stated that a 2-octet AS number (AS 23456) is to be reserved.

The way in which this all works is when an old and a new BGP speaker peer with each other, the new BGP speaker notices that the old BGP speaker did not send a 4-byte AS number capability along with its open message. The new BGP speaker will present itself as AS 23456 to the old BGP speaker and it will encode the AS-path information into the new communities.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

When the old BGP speaker advertises the path information to another new BGP speaker, that BGP speaker can reconstruct the AS path information from the added attributes. This way, the information was effectively tunneled across the old BGP speaker.

Notice that there are three formats to display 4-byte AS numbers:

- **asplain**: decimal representation of the AS, ranging from 0 4294967295.
- **asdot+**: the 32-bit number is broken up into a lower-order half and a high-order half. If you start incrementing the AS numbers at a low value, you can count from '0.1' up until '0.65535'. If you were to increment that AS number by 1, you would change the high-order bit to 1 and the low order bits to 0. You'd end up with '1.0' From then on, you can increment the low-order value until you reach 65535 again ('1.65535').
- **asdot**: mixture of asplain and asdot+. Values in the range of asplain are written in asplain. The values that exceed this range are written in asdot+.

Note: Junos OS allows for the configuration of BGP peers in asplain as well as asdot. Junos OS can also be configured to have the output of the BGP show commands display the AS in asdot.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Route policy

Before applying any routing-policies, it is important to know that the BGP protocol uses three different tables: RIB-IN, RIB-LOCAL and RIB-OUT.

The RIB-IN table is created for every established peer. The routes received from peers are stored in this table. The import policies are applied after the routes are stored.

The RIB-LOCAL is the route-table of the device. This information is used to populate the forwarding table.

The RIB-OUT is created for every peer and stores the routes that are advertised to the peer. Routes in this table have already passed the export policy.

The BGP protocol is policy driven protocol. We can use a routing-policy to determine how we want traffic to flow. The following BGP attributes can be manipulated using the **then** action in a routing policy:

- **next hop - then next-hop**
  - we can alter the next-hop across IBGP sessions to make sure other routers can resolve the advertised routes
  - the next-hop can be set to a discard interface, in case we want to black hole traffic
  - in case an IPv4 BGP session is used to advertise information for both IPv4 and IPv6, we can alter the next-hop for IPv6 routing information
- **local preference - then local-preference**
  - the highest local preference is preferred. Used to determine the exit point for outbound traffic inside the AS.
- **origin - then origin**
  - lowest origin value will win (IGP (I) over EGP (E) over incomplete(?))
- **MED - then metric**
  - used to manipulate inbound traffic into the AS in case multiple EBPG sessions exist with another AS
- **AS path - then as-path-prepend**
  - In an import policy, it can be used to manipulate outbound traffic
  - In an export policy, it can be used to manipulate inbound traffic
- **Community - then community**
  - add: adds a community to an empty or existing list of communities
  - set: overwrite the list of existing communities (removing previously applied communities, if any)
  - delete: deletes the reference community
  - well-known communities can be configure by specifying the name:
    - **no-export**: used to prevent route-advertisements to other AS
    - **no-advertise:** used to prevent route-advertisement to other BGP peers

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

The following attributes can be manipulated through BGP peering configuration without the use of a policy:

- **as-override**
  - used when an AS number is used on different sites that are connected to the same network
  - replace the peer AS with your own AS before advertising the route to a peer
- **remove-private**
  - strips all private AS numbers from the AS path. The router stops stripping private AS numbers from left to right until it encounters a public AS
- **metric-out**
  - explained earlier
- **local-preference**
  - explained earlier

AS-path and community regular expressions

You can use AS-path regular expressions to locate or identify a set of routes in a policy or when issuing a 'show route aspath-regex' command. AS-path regular expressions are configured under the [**policy-options as-path**] stanza and can be referenced as a **from** condition inside a routing-policy.

Commonly used AS-path Regex:

.	Any one AS (not a single character).
*****	0 or more repetitions of the AS.
{m,n}	Match at least m repetitions and at most n repetitions of the term.
{m}	Matches exactly m repetitions of the term.
{m,}	Matches m or more repetitions of the term.
+	Matches one or more instances of the term.
?	Zero or 1 repetitions of the term.
|	Matches the term on either side of this sign, performing an OR.
-	Between a starting and an ending range.
()	matches a null-path.

Note:

^	Added implicitly at the beginning of the regex, matches the start of the string.
$	Added implicitly at the end of the regex, matches the end of the string.

Some examples:

1234*	matches 1234, 1234 1234, null AS.
.*456.*	matches any length path that contains the list 4,5,6.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

64{2,6} 65? matches 2 to 6 instances of AS 64 followed by zero or 1 instance of AS 65.

.* 62 63 64 .* matches a path that contains the list '62 63 64' anywhere.

"()" matches null AS (can be used to advertise only the routes that originated from your own AS).

You can configure and apply communities to tag, or group together, certain prefixes. Additionally, you can use the configured communities to gather output for the 'show route community-name' command. When you are configuring communities, you can choose to use regex.

Commonly used community Regex:

. = Any single digit

***** = Any single community

.* = Everything in complex form

Examples:

"* : *" matches all communities (in case you want to strip routes from any communities)

600:* match any community which first part is 600

111:50.0 matches 111:5000, 111:5010, 111:5020, …

"^56: (2.*)$" matches 56:2111, 56:26765,….

Please refer to Juniper web site for more information about regular expressions.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Monitoring and troubleshooting

Common BGP issues related to the BGP sessions include issues with:

- Connectivity (no IP connectivity between the hosts)
- BGP misconfiguration (wrong AS, wrong source address, TTL of 1)
- Authentication mismatch (wrong password, no password)
- Firewall filters (BGP peer address no allowed in the filter used to protect the control plane)

In some cases, your BGP session will be established, but you will find that there is still no communication between the autonomous systems. In these cases, the issue could be related to the following:

- Policy misconfiguration (too much or the wrong prefixes are filtered)
- BGP next hop reachability (the next-hop IP cannot be resolved)

Troubleshooting BGP usually starts with the **show bgp summary** command. Here we can quickly spot which peer is not in the established state. We should then check for connectivity issues. If there is IP connectivity between the IP addresses used to setup the peering session, we can see if **show log message** outputs anything related to the issue.

If the log leaves us with nothing, we can either use **traceoptions** for BGP or monitor **traffic interfaces *<interface name>* detail no-resolve**

Be aware of the fact that hidden routes can be in that state for the following reasons:

- next-hop is not reachable
- policy rejects the route

The following commands are helpful troubleshooting and verifying BGP:

show bgp neighbor

show route advertising-protocol bgp

show route receive-protocol bgp

show route protocol bgp

show route hidden

show route resolution unresolved

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 1: Base Network Deployment

1) Load the base configuration on the devices:

```
root@srx1# load override terminal
```

One Ctrl+D is needed to load configuration then enter and commit.

Do not forget to reboot SRX devices in order to activate packet mode.

2) The same procedure is used to load the base configuration on the virtual device and AS switch.

You can use the *lab* username and *lab123* password to connect to these devices.

3) We configure the additional interfaces towards the EBGP neighbors.

On srx1:

```
root@srx1# set interfaces ge-0/0/4 unit 101 description "L3 connection
to AS1912"

root@srx1# set interfaces ge-0/0/4 unit 101 vlan-id 101

root@srx1# set interfaces ge-0/0/4 unit 101 family inet address
172.16.101.2/30
```

On srx2:

```
root@srx2# set interfaces ge-0/0/4 unit 202 description "L3 connection
to AS512 No.2"

root@srx2# set interfaces ge-0/0/4 unit 202 vlan-id 202

root@srx2# set interfaces ge-0/0/4 unit 202 family inet address
172.16.202.2/30

root@srx2# set interfaces ge-0/0/4 unit 202 family inet6 address
2001:f0f4:f691:202::2/126
```

On ex2:

```
root@ex2# set interfaces ge-0/0/14 unit 103 description "L3 connection
to P2"

root@ex2# set interfaces ge-0/0/14 unit 103 vlan-id 103
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set interfaces ge-0/0/14 unit 103 family inet6
```

On ex4:

```
root@ex4# set interfaces ge-0/0/14 unit 104 description "L3 connection
to P2"

root@ex4# set interfaces ge-0/0/14 unit 104 vlan-id 104

root@ex4# set interfaces ge-0/0/14 unit 104 family inet address
172.16.104.1/30
```

On srx4:

```
root@srx4# set interfaces ge-0/0/4 unit 105 description "L3 connection
to AS16535"

root@srx4# set interfaces ge-0/0/4 unit 105 vlan-id 105

root@srx4# set interfaces ge-0/0/4 unit 105 family inet6 address
2001:f0f4:f691:105::2/126

root@srx4# set interfaces ge-0/0/4 unit 106 description "L3 connection
to AS30987"

root@srx4# set interfaces ge-0/0/4 unit 106 vlan-id 106

root@srx4# set interfaces ge-0/0/4 unit 106 family inet6 address
2001:f0f4:f691:106::2/126
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 2: Internal BGP Configuration

1) Configuring iBGP:

A full mesh of iBGP sessions needs to be configured on the devices under your administration. We need to separate the groups into IPv4 and IPv6, the ex3 and ex4 switches only participate in IPv4 iBGP.

First, we specify what autonomous-system the local device belongs to. We configure this on every device:

```
root@srx1# set routing-options autonomous-system 9876
```

Then, we configure the groups. Inside the group configuration, we can choose to configure '**type internal**' or '**peer-as 9876**'. With iBGP peering, it is also common to source the iBGP session from the local loopback interface, so we should not forget the **local-address** statement.

First, the IPv4 group:

```
root@srx1# set protocols bgp group internal4 type internal
root@srx1# set protocols bgp group internal4 local-address 172.30.15.1
root@srx1# set protocols bgp group internal4 neighbor 172.30.15.2
root@srx1# set protocols bgp group internal4 neighbor 172.30.15.3
root@srx1# set protocols bgp group internal4 neighbor 172.30.15.4
root@srx1# set protocols bgp group internal4 neighbor 172.30.15.5
root@srx1# set protocols bgp group internal4 neighbor 172.30.15.6
root@srx1# set protocols bgp group internal4 neighbor 172.30.15.7
root@srx1# set protocols bgp group internal4 neighbor 172.30.15.8
root@srx1# set protocols bgp group internal4 neighbor 172.30.15.9
root@srx1# set protocols bgp group internal4 neighbor 172.30.15.10
```

Then, the IPv6 group:

```
root@srx1# set protocols bgp group internal6 type internal
root@srx1# set protocols bgp group internal6 local-address
fd17:f0f4:f691:15::1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:15::2

root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:15::3

root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:15::4

root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:15::5

root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:15::6

root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:15::9

root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:15::a
```

The BGP configuration is applied on all the devices, except for the ex3 and ex4 switch. Here, only IPv4 is required:

```
root@ex3# set protocols bgp group internal4 type internal

root@ex3# set protocols bgp group internal4 local-address 172.30.15.7

root@ex3# set protocols bgp group internal4 neighbor 172.30.15.1

root@ex3# set protocols bgp group internal4 neighbor 172.30.15.2

root@ex3# set protocols bgp group internal4 neighbor 172.30.15.3

root@ex3# set protocols bgp group internal4 neighbor 172.30.15.4

root@ex3# set protocols bgp group internal4 neighbor 172.30.15.5

root@ex3# set protocols bgp group internal4 neighbor 172.30.15.6

root@ex3# set protocols bgp group internal4 neighbor 172.30.15.8

root@ex3# set protocols bgp group internal4 neighbor 172.30.15.9

root@ex3# set protocols bgp group internal4 neighbor 172.30.15.10
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification:

On srx1:

```
root@srx1# run show bgp summary
Groups: 2 Peers: 16 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.0
                     491        446          0          0          0
0
inet6.0
                     155        137          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Accepted/Damped...
172.30.15.2            9876         55        448       0       0
23:05 3/3/3/0              0/0/0/0
172.30.15.3            9876         47        411       0       0
20:40 0/0/0/0              0/0/0/0
172.30.15.4            9876         42        407       0       0
18:44 0/0/0/0              0/0/0/0
172.30.15.5            9876         41        405       0       0
17:58 0/0/0/0              0/0/0/0
172.30.15.6            9876         20        386       0       0
17:06 0/0/0/0              0/0/0/0
172.30.15.7            9876         38        400       0       0
15:37 0/1/1/0              0/0/0/0
172.30.15.8            9876         37        399       0       0
14:53 8/9/9/0              0/0/0/0
172.30.15.9            9876         40        407       0       0
18:36 0/0/0/0              0/0/0/0
172.30.15.10           9876         41        407       0       0
18:39 0/0/0/0              0/0/0/0
fd17:f0f4:f691:15::2        9876         52        137       0       0
22:53 Establ
  inet6.0: 16/16/16/0
fd17:f0f4:f691:15::3        9876         46        127       0       0
20:38 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:15::4        9876         43        123       0       0
18:44 Establ
  inet6.0: 2/2/2/0
fd17:f0f4:f691:15::5        9876         42        121       0       0
17:58 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:15::6        9876         21        102       0       0
17:06 Establ
  inet6.0: 16/16/16/0
fd17:f0f4:f691:15::9        9876         41        123       0       0
18:48 Establ
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  inet6.0: 0/0/0/0
fd17:f0f4:f691:15::a        9876          40         123        0        0
18:40 Establ
  inet6.0: 0/0/0/0
```

On ex3:

```
root@ex3# run show bgp summary
Groups: 1 Peers: 9 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.0               447        445          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Accepted/Damped...
172.30.15.1            9876        404         42       0       0
17:46 434/435/435/0        0/0/0/0
172.30.15.2            9876         40         41       0       0
17:42 3/3/3/0              0/0/0/0
172.30.15.3            9876         40         41       0       0
17:43 0/0/0/0              0/0/0/0
172.30.15.4            9876         38         40       0       0
17:34 0/0/0/0              0/0/0/0
172.30.15.5            9876         40         40       0       0
17:33 0/0/0/0              0/0/0/0
172.30.15.6            9876         21         22       0       0
17:40 0/0/0/0              0/0/0/0
172.30.15.8            9876         41         39       0       0
17:02 8/9/9/0              0/0/0/0
172.30.15.9            9876         38         41       0       0
17:42 0/0/0/0              0/0/0/0
172.30.15.10           9876         39         41       0       0
17:43 0/0/0/0              0/0/0/0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

2) Make sure all iBGP sessons are authenticated with MD5 authentication:

Properties of a BGP peering session can be configured for the peer, for the group, or globally. The most specific configuration will take effect. In this case, we configure the authentication per group. This way, we only have to configure the '**authentication-key**' twice:

```
root@srx1# set protocols bgp group internal4 authentication-key
workbook

root@srx1# set protocols bgp group internal6 authentication-key
workbook
```

Do not forget to check the neighborship by issuing the **show bgp summary** on all of the devices to see if the sessions become establish again. The command **show bgp neighbor <*neighbor*>** will display a list of 'Options'. When BGP sessions are configured with MD5 authentication, the list of options will include 'AuthKey'.

3) iBGP state changes should be logged on all the devices:

For this assignment, we configure the option globally. By doing so, we affect all configured BGP sessions in every group.

```
root@srx1# set protocols bgp log-updown
```

Verification:

```
root@srx1# run show bgp neighbor 172.30.15.6
Peer: 172.30.15.6+53826 AS 9876 Local: 172.30.15.1+179 AS 9876
  Type: Internal    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ nhs ] Import: [ rtbh ]
  Options: <Preference LocalAddress HoldTime AuthKey LogUpDown
Refresh>
```

4) Make sure that all iBGP sessions on ex2 have a 60 seconds keepalive: should be set to 60 Sec on ex2:

The keepalive is not a value you can configure directly. What we can configure directly is the '**hold-time**' value. Devices calculate the keepalive interval by dividing the hold-time value by three. The hold-time

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

value is advertised in the BGP Open message when peers negotiate a connection. The smallest of the hold-time values is used by both BGP peers.

To achieve a 60 second keepalive interval, a hold-time of 180 is required. The Junos OS default hold time of 90 is lower then 180, so if we were to only configure the ex2 switch with a hold-time of 180, the active hold-time would still be 90. Therefore, we need to configure ex2 and the peering devices with a hold-time of 180:

```
root@ex2# set protocols bgp group internal4 hold-time 180

root@ex2# set protocols bgp group internal6 hold-time 180
```

Verification:

```
root@ex2# run show bgp neighbor
Peer: 172.30.15.1+179 AS 9876  Local: 172.30.15.6+53826 AS 9876
  Type: Internal    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Import: [ rtbh ]
  Options: <Preference LocalAddress HoldTime AuthKey LogUpDown
Refresh>
  Authentication key is configured
  Local Address: 172.30.15.6 Holdtime: 180 Preference: 170
  Number of flaps: 0
  Peer ID: 172.30.15.1     Local ID: 172.30.15.6       Active
Holdtime: 180
  Keepalive Interval: 60          Peer index: 3
```

5) This assignment requires us to accomplish two objectives. The first objective is the redistribution of an aggregate route into iBGP, while tagging it with a specific community. The second objective is making sure that srx1 and srx2 use the iBGP route without changing the local preference:

To complete our first objective, we create and apply an export policy that will redistribute the aggregate route into BGP. When configuring the policy, we try to be as specific as possible. In this case, we match the exact prefix (fd17:f0f4:f691:f::/64) and protocol source (aggregate).

We are also required to attach a community to the route. We configure the community, in this case named 'DC2.aggregate', in the [**policy-options community**] stanza. We then reference it in the our routing-policy using the 'set DC2.aggregate'. Using the **set** keyword implies that all other communities are replaced with our DC2.aggregate community. If **add** would have been used, the community would have been added to a (possibly) already existing list of communities.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On srx4:

```
root@srx4# set policy-options policy-statement aggregate.to.bgp term 1
from protocol aggregate

root@srx4# set policy-options policy-statement aggregate.to.bgp term 1
from route-filter fd17:f0f4:f691:f::/64 exact

root@srx4# set policy-options policy-statement aggregate.to.bgp term 1
then community set DC2.aggregate

root@srx4# set policy-options policy-statement aggregate.to.bgp term 1
then accept

root@srx4# set policy-options community DC2.aggregate members 9876:100

root@srx4# set protocols bgp group internal6 export aggregate.to.bgp
```

The second objective is to make sure that srx1 and srx2 actually use the iBGP route. The route was redistributed into OSPF, so we need to alter the preference value to something low enough that ensure the routers to use it. By default, OSPF external routes have a preference of 150 and BGP routes (iBGP and eBGP alike) have a preference of 170.

We can lower the preference of the BGP routes by configuring a policy that matches the previously tagged BGP routes (tagged with the community DC2.aggregate) and then altering the preference using 'then preference 30':

On srx1 and srx2:

```
root@srx1# set policy-options policy-statement lower.preference term 1
from protocol bgp

root@srx1# set policy-options policy-statement lower.preference term 1
from community DC2.aggregate

root@srx1# set policy-options policy-statement lower.preference term 1
then preference 30


root@srx1# set policy-options community DC2.aggregate members 9876:100


root@srx1# set protocols bgp group internal6 import lower.preference
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Note: we configured the policy without a '**then accept**'. Because of this, any additionally configured term or policy can also be used to further alter the same routes. For example, if we were to configure an additional term that matches all routes and changes the local-preference, the same 'DC2.aggregate' tagged BGP routes would also see their local-preference altered.

Verification:

```
root@srx1# run show route fd17:f0f4:f691:f::/64
inet6.0: 189 destinations, 215 routes (188 active, 0 holddown, 1
hidden)
+ = Active Route, - = Last Active, * = Both
fd17:f0f4:f691:f::/64
                   *[BGP/30] 01:01:52, localpref 100, from
fd17:f0f4:f691:15::4
                     AS path: I
                   > to fe80::7a19:f700:c18:3d84 via ge-0/0/4.12
                     to fe80::3e94:d500:d46:7311 via ge-0/0/4.13
                    [OSPF3/150] 01:01:57, metric 0, tag 0
                   > to fe80::7a19:f700:c18:3d84 via ge-0/0/4.12
                     to fe80::3e94:d500:d46:7311 via ge-0/0/4.13
```

The above printout shows us we were successful, the BGP route fd17:f0f4:f691:f::/64 was selected because of the lower preference value we set using our import-policy.

6) For this assignment, we need to redistribute the IPv4 RIP routes to iBGP on ex3 and ex4. We also need to load balance the routes on srx2:

We start out redistributing the IPv4 RIP routes into iBGP. We create a policy and try to be as specific as possible:

On ex3 and ex4:

```
root@ex3# set policy-options policy-statement aggregate.to.bgp term 1
from protocol aggregate

root@ex3# set policy-options policy-statement aggregate.to.bgp term 1
from route-filter 172.31.0.0/20 exact

root@ex3# set policy-options policy-statement aggregate.to.bgp term 1
then accept

root@ex3# set protocols bgp group internal4 export aggregate.to.bgp
```

After redistributing the routes on both ex3 and ex4, we need to make sure that srx2 will load-balance the routes. This requires two steps. First, we need to enable the iBGP sessions srx2 has with ex3 and ex4 for multipathing. We do this by configuring the peering session with the keyword '**multipath**'. Normally,

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

the BGP path selection singles out 1 best route. This part of the configuration will allow the srx to consider the BGP routes to be load-balanced.

This will not make the srx load-balance yet. To make the srx actually starts using both BGP routes, we need to configure the forwarding table to do so. This is done by applying a policy to the forwarding table.

On srx2:

```
root@srx2# set protocols bgp group internal4 neighbor 172.30.15.7
multipath

root@srx2# set protocols bgp group internal4 neighbor 172.30.15.8
multipath


root@srx2# set policy-options policy-statement load.balance term 1
then load-balance per-packet

root@srx2# set routing-options forwarding-table export load.balance
```

Verification:

```
root@srx2# run show route forwarding-table destination 172.31.0.0/20
Routing table: default.inet
Internet:
Destination       Type RtRef Next hop        Type Index NhRef Netif
172.31.0.0/20     user     0                 ulst 262156    1
                                 indr 262154    2
                  172.30.0.14    ucst   719    13 ge-0/0/4.23
                                 indr 262155    10
                  172.30.0.10    ucst   730    12 ge-0/0/4.24
```

Note: the '**ulst**' is a list of unicast next-hops.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 3: External BGP Configuration

1) We need to configure an EBGP session between srx1 and ISP1. This single session needs to support both IPv6 and IPv4 routes:

We enable the EBGP session with AS1912 with both address families 'inet unicast' as well as 'inet6 unicast'. Even though the 'inet unicast' address family might be enabled by default, we still need to configure it. This is because as soon as any other address family is configured, the 'inet unicast' family is no longer automatically enabled.

When we use a single IPv4 BGP session to exchange both IPv4 and IPv6 routes, we will run into an issue where the router will not be able to resolve any of the IPv6 routes. This will cause for the routes to be hidden. By default, the IPv6 routes will have their next-hop attribute set to an IPv4-mapped IPv6 address. To make sure that srx1 can resolve this route, we configure an IPv4-mapped IPv6 address on the interface:

```
root@srx1# set protocols bgp group ISP1 family inet unicast

root@srx1# set protocols bgp group ISP1 family inet6 unicast

root@srx1# set protocols bgp group ISP1 peer-as 1912

root@srx1# set protocols bgp group ISP1 neighbor 172.16.101.1

root@srx1# set interfaces ge-0/0/4 unit 101 family inet6 address
::ffff:172.16.101.2/126
```

Verification:

```
root@srx1# run show bgp summary
Groups: 3 Peers: 17 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.0
                     491        446          0          0          0
0
inet6.0
                     155        137          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Accepted/Damped...
172.16.101.1           1912        564         63       0       0
24:47 Establ
  inet.0: 435/478/438/0
  inet6.0: 103/121/120/0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# run show route receive-protocol bgp 172.16.101.1 table
inet6.0
inet6.0: 189 destinations, 215 routes (188 active, 0 holddown, 1
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 2001:240::/32           ::ffff:172.16.101.1                     1912
1620 33112 60177 60637 30709 57763 I
* 2001:250:204::/48       ::ffff:172.16.101.1                     1912
1620 33112 5697 30284 25817 58254 24799 ?
* 2001:254::/32           ::ffff:172.16.101.1                     1912
1620 61671 52411 64333 44013 43197 43321 6648 I
```

2) In this task, we configure an EBGP session between srx2 and AS512. We need load balance the traffic across both links to ISP2:

Our objective can be accomplished by establishing an EBGP multihop session between the loopback IP addresses of the srx2 and the router in AS512. To enable IP connectivity across both links, we configure a static route to the loopback IP address of the router in AS512 across both links.

```
root@srx2# set routing-options rib inet6.0 static route
2001:f0f4:f691:fe::1/128 next-hop 2001:f0f4:f691:102::1

root@srx2# set routing-options rib inet6.0 static route
2001:f0f4:f691:fe::1/128 next-hop 2001:f0f4:f691:202::1

root@srx2# set routing-options static route 172.18.12.1/32 next-hop
172.16.102.1

root@srx2# set routing-options static route 172.18.12.1/32 next-hop
172.16.202.1
```

After enabling the IP connectivity, we configure the EBGP session. Here, we need to focus on two things. By default, a BGP session is sourced using the IP address that is closest to the destination of the BGP peer address. Since we want to have a peering session between the loopback IP addresses of the routers, we first need to configure srx2 to use the loopback IP to source the BGP session.

The second thing we need to keep in mind is the fact that packets used to set up an EBGP session have a default TTL of 1. For the BGP packets to make it to the loopback interface of the other device, we need to enable the BGP session with the '**multihop**' keyword. You can configure the exact TTL with this configuration command. If you omit the TTL, the Junos OS device will use a value of 64 by default.

```
root@srx2# set protocols bgp group ISP2 multihop

root@srx2# set protocols bgp group ISP2 local-address 172.30.15.2

root@srx2# set protocols bgp group ISP2 peer-as 512
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set protocols bgp group ISP2 neighbor 172.18.12.1

root@srx2# set protocols bgp group ISP2-v6 multihop

root@srx2# set protocols bgp group ISP2-v6 local-address
fd17:f0f4:f691:15::2

root@srx2# set protocols bgp group ISP2-v6 peer-as 512

root@srx2# set protocols bgp group ISP2-v6 neighbor
2001:f0f4:f691:fe::1
```

Verification:

```
root@srx2# run show bgp summary
Groups: 4 Peers: 18 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.0
                     923        448          0          0          0
0
inet6.0
                     242        137          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Accepted/Damped...
172.18.12.1             512        481         57       0       0
24:14 3/478/438/0          0/0/0/0
2001:f0f4:f691:fe::1         512        141         58       0       0
24:10 Establ
  inet6.0: 16/121/120/0
```

3) Configure an IPv6 EBGP session, using the link-local address, between ex2 and P2 without knowing the AS:

To make sure that ex2 uses it's link-local address, we configure the session with the '**local-interface**' configuration command. Since we do not know what AS to use, we simply configure a random AS and troubleshoot the session later on:

```
root@ex2# set protocols bgp group P2 peer-as 66666

root@ex2# set protocols bgp group P2 neighbor fe80::21f:1200:672e:6f81
local-interface ge-0/0/14.103
```

Troubleshooting:

```
root@ex2# run show bgp summary
Groups: 4 Peers: 17 Down peers: 1
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
fe80::21f:1200:672e:6f81       66666          0          2        0
0          18 Active
root@ex2# run monitor traffic interface ge-0/0/14.103 detail no-
resolve
Address resolution is OFF.
Listening on ge-0/0/14.103, capture size 1514 bytes

18:42:21.153789  In IP6 (class 0xc0, flowlabel 0x0f869, hlim 1, next-
header: TCP (6), length: 44) fe80::21f:1200:672e:6f81.56297 >
fe80::2e6b:f500:678f:87d1.179: S 1724170742:1724170742(0) win 16384
<mss 1440,nop,wscale 0,nop,nop,timestamp 7706404 0,sackOK,eol>
          Version 4, my AS 64512, Holdtime 90s, ID 1.1.1.1
          Optional parameters, length: 30
            Option Capabilities Advertisement (2), length: 6
              Multiprotocol Extensions (1), length: 4
                AFI IPv6 (2), SAFI Unicast (1)
            Option Capabilities Advertisement (2), length: 2
              Route Refresh (Cisco) (128), length: 0
            Option Capabilities Advertisement (2), length: 6
              32-Bit AS Number (65), length: 4
                 4 Byte AS 64512
```

The session was configured and the BGP peers started their attempt to set up a BGP session. We issued the '**monitor traffic interface ge-0/0/14.103 detail no-resolve**' command. This command allows us to see traffic directed to or sourced from the routing-engine. In this case, we spotted the other system sending a message stating 'my AS 64512'.

Now, we simply replace the 'peer-as 66666' configuration statement with the 'peer-as 64512' configuration statement:

```
root@ex2# set protocols bgp group P2 peer-as 64512
```

```
root@ex2# run show bgp summary
Groups: 4 Peers: 17 Down peers: 0
fe80::21f:1200:672e:6f81       64512         32        117       0
0       13:20 Establ
  inet6.0: 16/16/16/0
```

As an alternative to the '**monitor traffic**' command, we could have configured traceoptions under the [**protocols bgp group P2**] stanza.

4) Establish an EBGP session with AS1352.365 on ex4 and assume that ex4 does not support a 4-byte AS.:

To accomplish this task, we configure the peer as to be '23456'. The theory behind this is explained at the beginning of this chapter.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex4# set protocols bgp group P1 peer-as 23456

root@ex4# set protocols bgp group P1 neighbor 172.16.104.2
```

```
root@ex4# run show bgp summary
Groups: 2 Peers: 10 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.0               449        447          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Accepted/Damped...
172.16.104.2          23456         40        409       0       0
17:18 8/8/8/0                0/0/0/0
```

5) Configure an IPv6 eBGP session between srx4 and both ISP3 and ISP4. This task is combined with task 6.
6) Make srx4 write a syslog message if more then 10 prefixes are received from ISP3 or ISP4:

By default, the Junos OS will accept any number of prefixes across a BGP session. To achieve the objective of this task, we should use the '**prefix-limit maximum**' configuration command that is nested under the '**family inet6 unicast** ' stanza inside the BGP neighbor configuration stanza. This will make the device log a message as soon as the limit is exceeded.

```
root@srx4# set protocols bgp group ISP3 peer-as 16535

root@srx4# set protocols bgp group ISP3 neighbor 2001:f0f4:f691:105::1
family inet6 unicast prefix-limit maximum 10

root@srx4# set protocols bgp group ISP4 peer-as 30987

root@srx4# set protocols bgp group ISP4 neighbor 2001:f0f4:f691:106::1
family inet6 unicast prefix-limit maximum 10
```

```
root@srx4# run show bgp summary
2001:f0f4:f691:105::1       16535        350        432       0
0     2:41:02 Establ
  inet6.0: 1/1/1/0
2001:f0f4:f691:106::1       30987        350        432       0
0     2:40:52 Establ
  inet6.0: 0/1/1/0
```

Note: the '**prefix-limit teardown**' option can be used to not only log a message when the maximum prefix limit is crossed, but to terminate the BGP session as well for a short time. The '**idle-timeout**' option can be used to further specify how long the BGP session should remain down.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

7) Ensure that BGP state changes are logged to syslog:

```
root@srx1# set protocols bgp log-updown
```

The **'log-updown'** statement can be applied on a protocol, a group or on peer level. By configuring the statement on a protocol level, we ensure that any BGP session configured on the device will have its state changes logged.

The '**show bgp neighbor x.x.x.x | match Options**' command can be used to verify whether or not state changes are logged. If the state changes are logged, you'll see '**LogUpDown**' listed among the other options.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 4: BGP Routing policies

1) We need to export an aggregate that represents the internal IPv4 and IPv6 range to the EBGP neighbors:

On srx1:

On this device, a single BGP session is used to exchange both IPv4 as well as IPv6 routes. Because of this, we configure a single policy with two terms. One term for the IPv4 aggregate and a second term for the IPv6 aggregate:

```
root@srx1# set routing-options rib inet6.0 aggregate route
fd17:f0f4:f691::/48

root@srx1# set routing-options aggregate route 172.30.0.0/16

root@srx1# set policy-options policy-statement local.aggregate term 1
from protocol aggregate

root@srx1# set policy-options policy-statement local.aggregate term 1
from route-filter 172.30.0.0/16 exact

root@srx1# set policy-options policy-statement local.aggregate term 1
then accept

root@srx1# set policy-options policy-statement local.aggregate term 2
from protocol aggregate

root@srx1# set policy-options policy-statement local.aggregate term 2
from rib inet6.0

root@srx1# set policy-options policy-statement local.aggregate term 2
from route-filter fd17:f0f4:f691::/48 exact

root@srx1# set policy-options policy-statement local.aggregate term 2
then accept

root@srx1# set protocols bgp group ISP1 export local.aggregate
```

Verification:

```
root@srx1# run show route advertising-protocol bgp 172.16.101.1
172.30.0.0/16
inet.0: 546 destinations, 554 routes (507 active, 0 holddown, 40
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 172.30.0.0/16           Self                                    I
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# run show route advertising-protocol bgp 172.16.101.1 table
inet6.0
inet6.0: 189 destinations, 215 routes (188 active, 0 holddown, 1
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* fd17:f0f4:f691::/48     Self                                    I
```

On srx2:

```
root@srx2# set routing-options rib inet6.0 aggregate route
fd17:f0f4:f691::/48

root@srx2# set routing-options aggregate route 172.30.0.0/16

root@srx2# set policy-options policy-statement local.aggregate term 1
from protocol aggregate

root@srx2# set policy-options policy-statement local.aggregate term 1
from route-filter 172.30.0.0/16 exact

root@srx2# set policy-options policy-statement local.aggregate term 1
then accept

root@srx2# set policy-options policy-statement local.aggregate.v6 term
1 from protocol aggregate

root@srx2# set policy-options policy-statement local.aggregate.v6 term
1 from route-filter fd17:f0f4:f691::/48 exact

root@srx2# set policy-options policy-statement local.aggregate.v6 term
1 then accept

root@srx2# set protocols bgp group ISP2 export local.aggregate

root@srx2# set protocols bgp group ISP2-v6 export local.aggregate.v6
```

On ex2:

```
root@ex2# set routing-options rib inet6.0 aggregate route
fd17:f0f4:f691::/48

root@ex2# set policy-options policy-statement local.aggregate term 1
from protocol aggregate

root@ex2# set policy-options policy-statement local.aggregate term 1
from route-filter fd17:f0f4:f691::/48 exact
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set policy-options policy-statement local.aggregate term 1
then accept

root@ex2# set protocols bgp group P2 export local.aggregate
```

Area 1 is an NSSA area. The ex1 and ex2 ABRs (ex1 and ex2) inject default a gateway into this area. If we would choose to configure an aggregate on ex4, we would have a lack of contributing routes. Therefore, we configure an aggregate on ex1 and ex2 and export that route to ex4 using iBGP. To make sure that ex4 can export the route to P1, it needs to be able to resolve the next hop. For this reason, we alter the next-hop in the export policy on ex1 and ex2.

On ex1 and ex2:

```
root@ex1# set routing-options generate route 172.30.0.0/16

root@ex1# set policy-options policy-statement aggregate.internal.v4
term 1 from protocol aggregate

root@ex1# set policy-options policy-statement aggregate.internal.v4
term 1 from route-filter 172.30.0.0/16 exact

root@ex1# set policy-options policy-statement aggregate.internal.v4
term 1 then next-hop 172.30.0.33

root@ex1# set policy-options policy-statement aggregate.internal.v4
term 1 then accept

root@ex1# set protocols bgp group internal4 neighbor 172.30.15.8
export aggregate.internal.v4
```

Note: on ex2, the next-hop is set to 172.30.0.37.

On ex4:

After configuring ex1 and ex2, we only need to permit the prefix in the export policy on ex4:

```
root@ex4# set policy-options policy-statement local.aggregate term 1
from route-filter 172.30.0.0/16 exact

root@ex4# set policy-options policy-statement local.aggregate term 1
then accept

root@ex4# set protocols bgp group P1 export local.aggregate
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On srx4:

```
root@srx4# set routing-options rib inet6.0 aggregate route
fd17:f0f4:f691::/48

root@srx4# set policy-options policy-statement local.aggregate.v6 term
1 from protocol aggregate

root@srx4# set policy-options policy-statement local.aggregate.v6 term
1 from route-filter fd17:f0f4:f691::/48 exact

root@srx4# set policy-options policy-statement local.aggregate.v6 term
1 then accept

root@srx4# set protocols bgp group ISP3 export local.aggregate.v6
```

2) Use communities to mark EBGP learned routes. Ensure that neighbors can be uniquely identified by these communities. Delete any communities when advertising routes to Partners.

First, we configure a unique community for every individual peer. After this, we make sure that every peer will have that community added to all of the routes learned from that peer. We can do this by adding the community using an import policy. By configuring the '**community add**' statement in a policy, we add the referenced community to the list of communities that are (possibly) already attached to the BGP routes.

On srx1:

```
root@srx1# set policy-options community ISP1 members 9876:1001

root@srx1# set policy-options community ISP2 members 9876:1002

root@srx1# set policy-options community P1 members 9876:2001

root@srx1# set policy-options community P2 members 9876:2002


root@srx1# set policy-options policy-statement as1912.in term 1 then
local-preference 150

root@srx1# set policy-options policy-statement as1912.in term 1 then
community add ISP1

root@srx1# set protocols bgp group ISP1 import as1912.in
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification:

```
root@srx1# run show route detail | match 9876:1001
                Communities: 9876:1001
                Communities: 9876:1001
                Communities: 9876:1001
                Communities: 9876:1001
```

Every device that has EBGP sessions requires a similar configuration.

Note: by configuring the communities on every router, verification and troubleshooting can become a lot easier.

On ex2 and ex4:

On these devices, we have EBGP sessions with a partner. We not only need to tag the learned routes with a community, we also need to delete all communities from the routes we export to these EBGP peers. This can be achieved in three steps.

The first step is to configure a community that matches any possible community. This is done specifying the members of the community as '*:*'. The second step is to create a policy that deletes the previously configured community. The last step is to apply this policy in the proper direction. In this case, we need to apply the policy in the export direction.

```
root@ex2# set policy-options community match-all members *:*

root@ex2# set policy-options policy-statement delete.comunities term 1
then community delete match-all

root@ex2# set protocols bgp group P2 export delete.comunities
```

Verification:

We tagged routes learned from other eBGP peers with a unique community. To verify that we are advertising routes carrying that community, we can reference the '**community-name**' in the '**show route advertising-protocol bgp**' command. Since the export policy successfully deleted all communities, the following command should not return any route:

```
root@ex2#run show route advertising-protocol bgp
fe80::21f:1200:672e:6f81 community-name "ISP1"
inet6.0: 177 destinations, 183 routes (177 active, 0 holddown, 0 hidden)
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

3) Prefixes learned from ISP1 and ISP2 need to have the local AS prepended when advertised to P2 and vice versa:

BGP not only uses the AS path attribute as a loop detection mechanism, it also uses the AS path length as a way to determine the distance towards a prefix. The fact that an AS path contains multiple instance of the same AS does not matter, the only thing that matters is the length of the path. By prepending the local AS several times, we increase the AS path length and make the advertised route less attractive.

Because we have already tagged the routes from the different eBGP peers with a unique community, identifying what routes need to be prepended is easy. On the ex2 devices, we match the ISP learned routes by matching the relevant communities, and prepend those routes on export to P2. On the srx1 and srx2 device, we do the opposite. We match P2 learned routes and prepend those on export.

On ex2 and ex4:

```
root@ex2#set policy-options policy-statement internet.prepend term 1
from community ISP1

root@ex2#set policy-options policy-statement internet.prepend term 1
from community ISP2

root@ex2#set policy-options policy-statement internet.prepend term 1
then as-path-prepend "9876 9876 9876 9876"

root@ex2#set protocols bgp group P2 export internet.prepend

root@ex2#insert protocols bgp group P2 export internet.prepend before
delete.communities
```

Verfication:

```
root@ex2# run show route advertising-protocol bgp
fe80::21f:1200:672e:6f81
inet6.0: 177 destinations, 183 routes (177 active, 0 holddown, 0
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 2001:240::/32           Self                                    9876
9876 9876 9876 [9876] 1912 1620 33112 60177 60637 30709 57763 I
* 2001:250:204::/48       Self                                    9876
9876 9876 9876 [9876] 1912 1620 33112 5697 30284 25817 58254 24799 ?
* 2001:254::/32           Self                                    9876
9876 9876 9876 [9876] 1912 1620 61671 52411 64333 44013 43197 43321
6648 I
* 2001:2b8:17::/48        Self                                    9876
9876 9876 9876 [9876] 1912 1620 61671 41164 28389 7809 45703 I
* 2001:2b8:d4::/48        Self                                    9876
9876 9876 9876 [9876] 1912 1620 33112 27075 14574 55821 I
* 2001:2d8:1::/48         Self                                    9876
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
9876 9876 9876 [9876] 1912 1620 61671 49129 48750 6221 24807 I
* 2001:2e0::/32           Self                                    9876
9876 9876 9876 [9876] 1912 1620 33112 60177 23789 9452 28275 31473 ?
* 2001:3e0::/32           Self                                    9876
9876 9876 9876 [9876] 1912 1620 33112 60177 5037 32401 1078 9850 39463
42575 I
* 2001:410:102::/48       Self                                    9876
9876 9876 9876 [9876] 1912 1620 61671 41164 18823 8392 I
* 2001:418:1::/48         Self                                    9876
9876 9876 9876 [9876] 1912 1620 61671 30404 11114 53598 10591 4850
54617 I
```

On srx1 and srx2:

```
root@srx1# set policy-options community ISP1 members 9876:1001

root@srx1# set policy-options community ISP2 members 9876:1002

root@srx1# set policy-options community P1 members 9876:2001

root@srx1# set policy-options community P2 members 9876:2002


root@srx1# set policy-options policy-statement p2.prepend term 1 from
community P2

root@srx1# set policy-options policy-statement p2.prepend term 1 then
as-path-prepend "9876 9876 9876 9876"


root@srx1# set protocols bgp group ISP1 export p2.prepend
```

Verification:

```
root@srx1# run show route advertising-protocol bgp 172.16.101.1 table
inet6.0
inet6.0: 189 destinations, 215 routes (188 active, 0 holddown, 1
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* fd17:f0f4:f691::/48     Self                                    I
  fd17:f0f4:f691:d::/80
*                         Self                                    9876
9876 9876 9876 [9876] I
  fd17:f0f4:f691:d:1::/80
*                         Self                                    9876
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
9876 9876 9876 [9876] I
  fd17:f0f4:f691:d:2::/80
*                          Self                                    9876
9876 9876 9876 [9876] I
  fd17:f0f4:f691:d:3::/80
*                          Self                                    9876
9876 9876 9876 [9876] I
```

4) Making sure P1 routes are not advertised to the Internet:

Because we tagged the P1 learned routes with a unique community earlier, preventing the announcement of the P1 learned routes to the Internet is a pretty straightforward task. We simply reject all routes tagged with the P1 community on export to the ISPs:

On srx1 and srx2:

```
root@srx1# set policy-options policy-statement p1.reject term 1 from
community P1

root@srx1# set policy-options policy-statement p1.reject term 1 then
reject

root@srx1# set protocols bgp group ISP1 export p1.reject
```

Verification:

```
root@srx1#run show route advertising-protocol bgp 172.16.101.1
community-name P1
inet.0: 546 destinations, 554 routes (507 active, 0 holddown, 40
hidden)
inet6.0: 189 destinations, 215 routes (188 active, 0 holddown, 1
hidden)
```

5) Ensure P2 does not readvertise your prefixes:

We can use one of the well-known communities to ensure P2 does not readvertise our prefixes. We tag our prefixes with the 'no-export' community to indicate to the P2 autonomous system that the routes should not be readvertised. To be able to use the 'no-export' community in policy, we need to configure it first:

```
root@ex2#set policy-options community no-export members no-export

root@ex2#set policy-options policy-statement local.aggregate term 1
then community add no-export

root@ex2#insert protocols bgp group P2 export local.aggregate after
delete.communities
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification:

```
root@ex2#run show route advertising-protocol bgp
fe80::21f:1200:672e:6f81 community-name no-export
inet6.0: 177 destinations, 183 routes (177 active, 0 holddown, 0
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* fd17:f0f4:f691::/48     Self                                    I
```

6) On ex2 and ex4, only accept prefixes originated from peer autonomous systems:

To filter any prefixes that do not originate directly from within the peer autonomous systems, we create an as-path filter that matches all of the prefixes that originate from that peer autonomous system. We also factor in any possible as-prepending by configuring '+$'. The next as-path filter **"^64512+$"** matches an autonomous system path that consists of one or more repetitions of the 64512 autonomous system:

```
root@ex2#set policy-options as-path AS64512 "^64512+$"

root@ex2#set policy-options policy-statement as64152.in term 1 from
as-path AS64512
```

Verification:

```
root@ex2#run show route receive-protocol bgp fe80::21f:1200:672e:6f81
terse
inet.0: 523 destinations, 541 routes (523 active, 0 holddown, 0
hidden)
inet6.0: 177 destinations, 183 routes (177 active, 0 holddown, 0
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
  fd17:f0f4:f691:d::/80
*                         fe80::21f:1200:672e:6f81
64512 I
  fd17:f0f4:f691:d:1::/80
*                         fe80::21f:1200:672e:6f81
64512 I
  fd17:f0f4:f691:d:2::/80
*                         fe80::21f:1200:672e:6f81
64512 I
```

On ex4:

```
root@ex4#set policy-options as-path AS23456 "^23456+$"
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex4#set policy-options policy-statement as1352.365.in term 1 from
as-path AS23456
```

7) No private AS numbers should not be advertised to ISP1 and ISP2:

To have the local system strip any private AS number from the AS path attribute before advertising prefixes to peers, we use the **'remove private'** configuration command.

On srx1 and srx2:

```
root@srx1# set protocols bgp group ISP1 remove-private
```

Verification:

```
root@srx1# run show route advertising-protocol bgp 172.16.101.1
private
[edit]
```

8) Make sure that your network does not become a transit AS for ISP1 and ISP2:

To ensure that ISP1 and ISP2 will not start using your network for transit traffic, we ensure that none of the ISP1 learned prefixes are advertised to ISP2 and vice versa. We can achieve this by filtering the ISP learned routes in our eBGP export policy towards the ISPs. On export to ISP1, we filter out ISP2-learned routes. On export to ISP2, we filter our ISP1-learned routes:

On srx1 and srx2:

```
root@srx1# set policy-options policy-statement isp2.reject term 1 from
community ISP2

root@srx1# set policy-options policy-statement isp2.reject term 1 then
reject

root@srx1# set protocols bgp group ISP1 export isp2.reject
```

Verification:

```
root@srx1#run show route advertising-protocol bgp 172.16.101.1
community-name ISP1
inet.0: 546 destinations, 554 routes (507 active, 0 holddown, 40
hidden)
inet6.0: 189 destinations, 215 routes (188 active, 0 holddown, 1
hidden)
[edit]
```

9) Do not accept any default routes from IPv4 or IPv6 neighbors. Accept no IPv4 prefixes longer than /26 from ISP1 and ISP2:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On srx1 and srx2:

```
root@srx1# set policy-options policy-statement filter.routes term 1
from route-filter 0.0.0.0/0 through 0.0.0.0/32 reject

root@srx1# set policy-options policy-statement filter.routes term 1
from route-filter 0.0.0.0/0 prefix-length-range /27-/32 reject

root@srx1# set policy-options policy-statement filter.routes term 2
from family inet6

root@srx1# set policy-options policy-statement filter.routes term 2
from route-filter ::/0 through ::/128

root@srx1# set policy-options policy-statement filter.routes term 2
then reject

root@srx1# set protocols bgp group ISP1 import filter.routes

root@srx1# insert protocols bgp group ISP1 import filter.routes before
as1912.in
```

In the filter above, we used the '`route-filter 0.0.0.0/0 through 0.0.0.0/32 reject`' to filter out any IPv4 default route. This part of the statement matches the '0.0.0.0' route with any mask (0, 1, 2, 3....31 and 32). Any route other than a route containing exactly '0.0.0.0' is not matched. The same logic applies to the '`::/0 through ::/128`' statement. This was used to filter out any IPv6 route.

For IPv4, we had an additional objective. We also needed to filter out prefixes longer then /26. This was achieved using the '`route-filter 0.0.0.0/0 prefix-length-range /27-/32 reject`' configuration command inside the policy.

When multiple policies are configured for a BGP group (or peer), they are applied from top to bottom. Since another policy was already put in place, we use '`insert`' to make sure our newly created policy will be applied before the other policy.

Verification:

```
root@srx1# run show route receive-protocol bgp 172.16.101.1 0.0.0.0
exact
[edit]

root@srx1# run show route receive-protocol bgp 172.16.101.1 ::/0
exact table inet6.0
[edit]

root@srx1# run show route receive-protocol bgp 172.16.101.1 match-
prefix */27-32
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
inet.0: 546 destinations, 554 routes (507 active, 0 holddown, 40
hidden)
H1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
H3.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
inet6.0: 189 destinations, 215 routes (188 active, 0 holddown, 1
hidden)
root@srx1# run show route protocol bgp hidden

inet.0: 546 destinations, 554 routes (507 active, 0 holddown, 40
hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/4           [BGP ] 04:47:33, localpref 100
                      AS path: 1912 1620 61671 I
                    > to 172.16.101.1 via ge-0/0/4.101
0.0.0.0/5           [BGP ] 04:47:33, localpref 100
                      AS path: 1912 1620 61671 27075 I
                    > to 172.16.101.1 via ge-0/0/4.101
17.97.99.128/27     [BGP ] 04:47:33, localpref 100
```

On ex2:

```
root@ex2# set policy-options policy-statement filter.routes.v6 term 2
from family inet6

root@ex2# set policy-options policy-statement filter.routes.v6 term 2
from route-filter ::/0 through ::/128

root@ex2# set policy-options policy-statement filter.routes.v6 term 2
then reject

root@ex2# set protocols bgp group P2 import filter.routes.v6

root@ex2#  insert protocols bgp group P2 import filter.routes.v6
before as64152.in
```

10) srx1 should be the preferred exit point to the Internet and AS 1912 and srx2 should be the preferred exit point to AS 512:

In order to complete this task, we can use the local-preference attribute. The route with the highest local preference is preferred and the default local preference value is 100.

To meet our objectives, we raise the local preference for all of the ISP1-learned prefixes (both Internet and AS 1912 alike) on srx1 to 150.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On srx2, we need to make sure that prefixes originating from AS 512 are preferred over the same AS 512 originated prefixes that were learned via srx1. We do this by setting the local preference to 200, but only for AS 512 originated prefix. This is done using an as-path filter that singles out the AS 512 originated routes:

On srx1:

```
root@srx1# set policy-options policy-statement as1912.in term 1 then
local-preference 150
```

On srx2:

```
root@srx2# set policy-options as-path AS512 "^512+$"

root@srx2# set policy-options policy-statement as512.in term 1 from
as-path AS512

root@srx2# set policy-options policy-statement as512.in term 1 then
local-preference 200

root@srx2# insert policy-options policy-statement as512.in term 1
before term 2
```

Verification:

```
root@srx1# run show route terse | match 150
* 1.64.0.0/10          B 170        150             >172.16.101.1
1912 1620 61671 I
* 1.84.160.0/20        B 170        150             >172.16.101.1
1912 1620 33112 I
* 1.96.0.0/11          B 170        150             >172.16.101.1
1912 1620 33112 63164 40776 51777 I
* 1.161.192.0/21       B 170        150             >172.16.101.1
1912 1620 33112 30404 32138 45045 I
* 1.176.0.0/12         B 170        150             >172.16.101.1
1912 1620 33112 49129 16320 52954 I

root@srx2# run show route extensive | match "Localpref: 200"
    Localpref: 200
                Localpref: 200
    Localpref: 200
                Localpref: 200
    Localpref: 200
                Localpref: 200
    Localpref: 200
                Localpref: 200
    Localpref: 200
                Localpref: 200
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

11) srx1 should be the preferred entry point for inbound traffic coming from the internet:

To accomplish this task, we prepend the local AS number multiple times to any eBGP exported aggregate routes on srx2. This will make the route via srx2 seem less attractive to the rest of the Internet:

On srx2:

```
root@srx2# set policy-options policy-statement local.aggregate term 1
then as-path-prepend "9876 9876 9876 9876"

root@srx2# set policy-options policy-statement local.aggregate.v6 term
1 then as-path-prepend "9876 9876 9876 9876"
```

Verification:

```
root@srx2# run show route advertising-protocol bgp 172.18.12.1
172.30/16
inet.0: 555 destinations, 1011 routes (516 active, 0 holddown, 40
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 172.30.0.0/16           Self                                    9876
9876 9876 9876 [9876] I

root@srx2# run show route advertising-protocol bgp
2001:f0f4:f691:fe::1 fd17:f0f4:f691::/48 exact
inet6.0: 193 destinations, 308 routes (192 active, 0 holddown, 1
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* fd17:f0f4:f691::/48     Self                                    9876
9876 9876 9876 [9876] I
```

12) Make sure that Partner routes are preferred over the same routes learned from the ISPs.

We can use the local preference attribute to accomplish this task. We do have to factor in any previously used values applied elsewhere. Therefore, on ex2 and ex4, we set the local preference to a value higher than what we used in the ISP policies:

```
root@ex2# set policy-options policy-statement as64152.in term 1 then
local-preference 230
```

Verification:

```
root@ex2# run show route extensive | match "Localpref: 230"
    Localpref: 230
               Localpref: 230
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Localpref: 230
```

13) Traffic to routes marked with the custom community '*:666' should be discarded on all routers in your AS.

To accomplish our objective, any BGP learned route with the community '*:666' attached to it should be installed with a 'discard' next-hop value. To achieve this, we first configure the rtbh community. Then, we write a policy that sets the next-hop to discard for any route tagged with this rtbh community.

After this, we could apply this policy in the import direction for every group on all of the routers. To speed things up, we configure and apply a group configuration. This group configuration contains a wildcard instead of a group. Therefore, we apply this group, the configuration is applied to any configured BGP group on the device.

On all the devices:

```
root@srx1# set policy-options community rtbh members *:666

root@srx1# set policy-options policy-statement rtbh term 1 from community rtbh

root@srx1# set policy-options policy-statement rtbh term 1 then next-hop discard

root@srx1# set groups rtbh protocols bgp group <*> import rtbh

root@srx1# set protocols bgp apply-groups rtbh
```

Verification:

```
root@srx1# run show route community *:666
inet.0: 546 destinations, 554 routes (507 active, 0 holddown, 40 hidden)
+ = Active Route, - = Last Active, * = Both
128.0.0.0/4         *[BGP/170] 05:24:53, localpref 150, from 172.16.101.1
                      AS path: 1912 1620 61671 60177 39950 3300 18089 50215 46389 5586 I
                    Discard
```

Note: to verify that the configuration and application of the rtbh group was successful, we can use the '**show configuration protocols bgp | display inheritance**' command.

14) You cannot have any black hole or suboptimal routing:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On eBGP, the next-hop of the advertised routes is changed. This is not the case when routes are readvertised to iBGP neighbors. To make sure that the BGP speaking routers in the autonomous system can resolve the iBGP learned routes, we can change the next-hop of these routes using a policy. We configure a NHS policy that changes the next-hop to the loopback IP address of the advertising router. Since the loopback IP is injected into the IGP, the routers should have no problems resolving any of the routes:

On all devices:

```
root@srx1# set policy-options policy-statement nhs term 1 then next-hop self

root@srx1# set protocols bgp group internal4 export nhs

root@srx1# set protocols bgp group internal6 export nhs
```

Note: issuing the '**show route resolution unresolved**' command on every device is an excellent way to quickly spot any routes that cannot be resolved.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 5: Filter Based Forwarding

There is one walkthrough to cover all of the individual assignments given in this task.

The assignments are covered in the same steps that are required to configure filter based forwarding:

1) Defining a routing instance of the type 'forwarding'
2) Sharing routes with the routing instances through the use of a rib group
3) Defining a filter and applying it on the proper interface

We start out defining two routing-instances, one for each ISP. The routing-instances should be of the instance-type forwarding:

```
root@srx4# set routing-instances ISP3_best instance-type forwarding

root@srx4# set routing-instances ISP4_best instance-type forwarding
```

The routing-instances do not have any routes yet. We need to copy the routes from the main inet6.0 table to the routing tables that are associated with the newly created routing-instances. To achieve this, we need to use a rib-group.

Inside a rib-group configuration, we need to list the table that contains the routes first. The tables configured after that are the tables that we can copy the information to. In this case, the routes will be present in the main inet.6 table and we will be copying the routes to the '**ISP3_best.inet6.0**' and '**ISP4_best.inet6.0**' table. Therefore, we configure the rib-group as follows:

```
root@srx4# set routing-options rib-groups FBFrib import-rib inet6.0

root@srx4# set routing-options rib-groups FBFrib import-rib
ISP3_best.inet6.0

root@srx4# set routing-options rib-groups FBFrib import-rib
ISP4_best.inet6.0
```

The rib-group needs to be applied before it has any effect. To apply a rib-group, it needs to be configured on the place where the routes are coming from. There are a lot of places that you can apply rib-groups to. For instance, when you need to copy OSPF routes, you configure the rib-group under [**protocols ospf rib-group**], when you need to copy interface routes, you configure the rib-group under [**routing-options interface-routes rib-group**].

In our case, we need to copy routing information that we learn from a BGP peer. Therefore, we apply the rib-group under the corresponding BGP groups:

```
root@srx4# set protocols bgp group ISP3 family inet6 unicast rib-group
FBFrib
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx4# set protocols bgp group ISP4 family inet6 unicast rib-group
FBFrib
```

We then configure a routing-policy that sets a local-preference high enough to achieve what we want. We have already used a local preference value of 150 for ISP1 and a value of 200 for ISP2. To ensure that ISP3 and ISP4 are the exit points for the prefixes, we can use a local preference of 250 here:

```
root@srx4# set policy-options policy-statement fbf.isp3.in term 1 to
rib ISP3_best.inet6.0

root@srx4# set policy-options policy-statement fbf.isp3.in term 1 then
local-preference 250

root@srx4# set policy-options policy-statement fbf.isp3.in term 2 from
route-filter fdfc:b846:cb44:7350::/64 exact

root@srx4# set policy-options policy-statement fbf.isp3.in term 2 then
accept


root@srx4# set policy-options policy-statement fbf.isp4.in term 1 to
rib ISP4_best.inet6.0

root@srx4# set policy-options policy-statement fbf.isp4.in term 1 then
local-preference 250

root@srx4# set policy-options policy-statement fbf.isp4.in term 2 from
route-filter fdfc:b846:cb44:7350::/64 exact

root@srx4# set policy-options policy-statement fbf.isp4.in term 2 then
accept
```

We apply the newly created policies as follows:

```
root@srx4# set protocols bgp group ISP3 import fbf.isp3.in

root@srx4# set protocols bgp group ISP4 import fbf.isp4.in
```

The final steps of this task include the creation of a firewall filter and the application of the firewall filter on the relevant interfaces.

We need to use the firewall filter to match traffic sourced from the ex1 IPv6 loopback address (fd17:f0f4:f691:15::5) and send it out via AS16535 (ISP3). Traffic sourced from the srx4a IPv6 loopback (fd17:f0f4:f691:15::9) address needs to be send via AS30987 (ISP4). This is achieved by configuring two terms that match the relevant IPv6 address and place the traffic into the desired routing-instance. We also configure a named counter to make things easier during our verification. Lastly, we conclude the firewall filter with a permit all:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx4# set firewall family inet6 filter FBF term ISP3_best from
source-address fd17:f0f4:f691:15::5/128

root@srx4# set firewall family inet6 filter FBF term ISP3_best then
count ISP3_best

root@srx4# set firewall family inet6 filter FBF term ISP3_best then
routing-instance ISP3_best

root@srx4# set firewall family inet6 filter FBF term ISP4_best from
source-address fd17:f0f4:f691:15::9/128

root@srx4# set firewall family inet6 filter FBF term ISP4_best then
count ISP4_best

root@srx4# set firewall family inet6 filter FBF term ISP4_best then
routing-instance ISP4_best

root@srx4# set firewall family inet6 filter FBF term permit_all then
accept
```

The firewall filter is applied in the input direction on the following interfaces:

```
root@srx4# set interfaces ge-0/0/4 unit 34 family inet6 filter input
FBF

root@srx4# set interfaces ge-0/0/4 unit 64 family inet6 filter input
FBF
```

Verification:

To verify our configuration, we first issue several pings sourced from the relevant loopback IP addresses on ex1 and srx4a. We then check to see if our named counter increments. If it does, we can conclude that the firewall filter is correctly matching the traffic:

```
root@ex1# run ping fdfc:b846:cb44:7350::1 source fd17:f0f4:f691:15::5
PING6(56=40+8+8 bytes) fd17:f0f4:f691:15::5 --> fdfc:b846:cb44:7350::1
16 bytes from 2001:f0f4:f691:105::1, icmp_seq=0 hlim=62 time=5.674 ms
16 bytes from 2001:f0f4:f691:105::1, icmp_seq=1 hlim=62 time=4.606 ms
16 bytes from 2001:f0f4:f691:105::1, icmp_seq=2 hlim=62 time=6.282 ms
root@srx4# run ping routing-instance SRX4a fdfc:b846:cb44:7350::1
source fd17:f0f4:f691:15::9
PING6(56=40+8+8 bytes) fd17:f0f4:f691:15::9 --> fdfc:b846:cb44:7350::1
16 bytes from 2001:f0f4:f691:106::1, icmp_seq=0 hlim=62 time=4.830 ms
16 bytes from 2001:f0f4:f691:106::1, icmp_seq=1 hlim=62 time=3.483 ms
16 bytes from 2001:f0f4:f691:106::1, icmp_seq=2 hlim=62 time=4.296 ms

root@srx4# run show firewall counter ISP3_best filter FBF
Filter: FBF
Counters:
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Name                                Bytes              Packets
ISP3_best                          199827                 2190

root@srx4# run show firewall counter ISP4_best filter FBF
Filter: FBF
Counters:
Name                                Bytes              Packets
ISP4_best                          196871                 2146
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Chapter 5 : Multicast Routing and Switching

A JNCIE-ENT candidate should be familiar with the following topics:

- Multicast addressing
- RPF check
- Multicast routing tables
- IGMP
- PIM
- SSM
- Anycast-RP configuration

Multicast Overview

Multicast traffic is either a one-to-many traffic flow, or a many-to-many traffic flow. A multicast source sends out traffic using its own source IP. The destination IP that the traffic is send to is an IP multicast group address. This is an address that falls within the class D range (224.0.0.0/4).

Receivers or hosts that are interested in this multicast traffic have to inform the network that multicast traffic for a certain group needs to be send their way. IPv4 hosts can express their interest using IGMP, IPv6 hosts can use MLD. A multicast routing protocol, PIM or DVMRP, is used to build and maintain multicast forwarding trees.

There are 2 Service models used in multicast. In the Any Source Multicast (**ASM**) model, receivers can join a multicast group without specifying the source. In the Source Specific Multicast (**SSM**) mode, receivers always specify the multicast source they want to join.

There are 2 modes of distribution: **sparse** and **dense**. In dense mode, multicast is forwarded to all parts of a network and routers have to prune themselves if they don't want to receive the traffic any longer. In sparse mode, multicast traffic is only forwarded when receivers have expressed an interest for a particular multicast group.

There are two types of distribution trees, the Shortest Path Tree (**SPT**) and the shared tree, or RP-Tree (**RPT**). The SPT is rooted at the multicast source. Dense mode always uses the SPT. In sparse mode, it's used only after the receiver learns about the source. The RPT is a shared tree, only used in ASM. It's the tree between the RP and the receivers of a multicast group and it is used only in sparse mode.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Multicast addressing

The class D range of IP addresses, 224/4, can be used for multicast.

The base address of the range, 224.0.0.0, is reserved and cannot be used.

The 224.0.0.0/24 range is reserved for local wire use.

The default SSM address block is 232.0.0.0-232.255.255.255 (232/8).

RPF check

The RPF (Reverse Path Forwarding) mechanism is used to prevent looped and duplicated multicast packets. This mechanism inspects the source address and the incoming interface on which a multicast packet is received. These are then compared with the outgoing next-hop entry in the unicast routing table towards the source of the multicast packet.

For multicast packets to pass the RPF check, the outgoing interface in the unicast routing table must match the interface on which the multicast traffic is received. If there is a pass, the packet is forwarded. If the RPF check fails, the packet is discarded.

To check all the entries used for the RPF calculations, you can use the **show multicast rpf** command. The **show multicast rpf summary** will display the table that is used to calculate the RPF check.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Multicast routing tables

By default, the **inet.0** and the **inet6.0** are used for the RPF check. The **inet.1** and the **inet6.1** are used to store the multicast forwarding cache.

It is possible to have the router perform an RPF check based on information from another table. This way, you can build a multicast routing-topology that differs from your unicast routing topology. You can decide to use the **inet.2** table for this. By default, this table is empty. In order to have the RPF check use the **inet.2** table, we would have to take two steps:

- **Step 1**: configure a rib-group that copies information to this table and apply it to where the information should come from.

Example in which OSPF and interface routes are copied to the inet.2 table:

set routing-options rib-groups to-mcast import-rib inet.0
set routing-options rib-groups to-mcast import-rib inet.2

set protocols ospf rib-group inet to-mcast
set routing-options interface-routes rib-group inet to-mcast

- **Step 2**: configure a rib-group that only references the inet.2 table, and configure PIM to start using this table.

Example in which PIM is configured to start using the inet.2 table:

set routing-options rib-groups mcast import-rib inet.2
set protocols pim rib-group inet mcast

IGMP

IGMP (Internet Group Management Protocol) is used between hosts interested in receiving multicast packets and multicast routers.

There are three IGMP versions.

IGMP v1: the older (and inefficient) version that can only be used used for ASM.

IGMP v2: the Junos default. IGMP v2 supports explicit leave-group message to speed up convergence. It is used for ASM and can be used for SSM through SSM-mapping policies that convert IGMP v2 to IGMP v3.

IGMP v3: the default version for SSM (can support ASM as well).

Note: in IPv6 we use **MLD** instead of IGMP. MLDv1 is similar to IGMPv2 and MLDv2 is similar to IGMP v3

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

PIM

PIM (Protocol Independent Multicast) is the predominantly used multicast routing protocol. Two versions of PIM exist. These are PIM version 1 and PIM version 2. The default version that is used, is based on the location where PIM is configured. Version 1 is the default for RP mode and v2 is the default for interface mode. The version can be configured explicitly to override these defaults.

PIM can operate in several modes:

- **dense mode**: in this mode, the assumption is that all devices in the network are interested in receiving multicast traffic. Routers stop forwarding multicast traffic only after a prune message is received.
- **sparse mode**: sparse distribution of receivers is assumed, routers must explicitly join a multicast group before any traffic is received.
- **sparse-dense mode**: a combination of both. Interfaces can operate in either sparse or dense mode on a group-by-group basis.

PIM sparse mode in combination with ASM requires an **RP** (Rendezvous Point). The RP is the meeting place for originators of multicast traffic and interested receivers.

There are three ways to inform the devices in the network about the location of the RP:

1. **Static RP**: the RP is statically configured on every router.

2. **Auto-RP**: using auto-RP, the non-RP routers can learn the location of the RP automatically. Auto-RP requires PIM to be configured for sparse-dense mode. The dense-mode groups required are 224.0.1.39 and 224.0.1.40. These need to be configured on every router.

3. **BSR** (Bootstrap Router): BSR messages are forwarded hop-by-hop, informing the routers about the location of the RP. BSR requires PIM v2 and it can be configured per multicast group. The router with the highest priority (or IP address) is elected to be the bootstrap router. By default, routers have a priority of 0, making them ineligible to become the BSR.

A BSR-learned RP is preferred over an auto-RP learned RP.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Auto-RP configuration example:

RP routers:

set protocols pim dense-groups 224.0.1.39/32
set protocols pim dense-groups 224.0.1.40/32
set protocols pim rp local address *<ip>*
set protocols pim rp auto-rp mapping
set protocols pim interface all mode sparse-dense

Non-RP routers:

set protocols pim dense-groups 224.0.1.39/32
set protocols pim dense-groups 224.0.1.40/32
set protocols pim rp auto-rp discovery
set protocols pim interface all mode sparse-dense

BSR configuration example:

RP routers:

set protocols pim rp local address *<ip>*
set protocols pim rp bootstrap family inet priority 175

Non-RP routers:

set protocols pim interface all mode sparse

SSM

SSM uses PIM sparse mode and IGMPv3 to allow a client to receive multicast traffic directly from the source.

The default address range that can be used by SSM is 232/8. To be able to also use this range for ASM, we can configure the **asm-override-ssm** configuration command under the [**routing-options multicast**] stanza. To add an additional class D address to the default SSM range, we can configure the **ssm-group** under the [**routing-options multicast**] stanza.

In SSM, the multicast source is known to the receiver. Therefore, an RP does not need to be implemented.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

For hosts that do no support IGMPv3, we can use SSM mapping. SSM mapping can translate IGMPv1 and IGMPv2 reports to an IGMPv3 report. To enable SSM, we need to create a policy that matches a multicast group. That policy must then be referenced using the **ssm-map** configuration statement under the [**routing options multicast**] stanza.

Anycast-RP configuration

In order to have multiple active RPs per multicast group, anycast-RP needs to be deployed. When anycast-RP is used, the anycast-RP routers will share an IP address that they both announce into the network. The anycast-RP devices both perform the RP role. The devices in the network will use the RP that is closest to them. When the RP closest to the device fails, the other anycast-RP router can be used.

The anycast-RP routers need to keep each other updated on the interested receivers and the active multicast sources. There are two anycast RP mechanism that can be implemented to do this. The native **PIM anycast** can be implemented for both IPv4 and IPv6. **MSDP** can be implemented only for **IPv4**.

The steps to implement anycast-RP are the following:

1. Configure a unique address on the loopback interface and specify that IP address to be the **primary** address. This address is used for routing purposes (OSPF, BGP, etc.).
2. Configure a non-unique address that is used for anycast-RP on the RPs loopback address.
3. Configure the RP discovery mechanism of choice, if any (typically a static RP is configured).
4. Configure the anycast-RP mechanism of choice (PIM or MSDP).

PIM anycast configuration example wherein router 1 uses 10.0.0.1, router 2 uses 10.0.0.2 and the anycast-RP address is 10.0.0.255:

PIM anycast router 1:

set protocols pim rp local family inet address 10.0.0.255
set protocols pim rp local family inet anycast-pim rp-set address 10.0.0.2
set protocols pim rp local family inet anycast-pim local-address 10.0.0.1
set protocols pim interface all mode sparse

PIM anycast router 2:

set protocols pim rp local family inet address 10.0.0.255
set protocols pim rp local family inet anycast-pim rp-set address 10.0.0.1
set protocols pim rp local family inet anycast-pim local-address 10.0.0.2
set protocols pim interface all mode sparse

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

MSDP anycast router configuration example wherein router 1 uses 10.0.0.1, router 2 uses 10.0.0.2 and the anycast-RP address is 10.0.0.255:

MSDP anycast router 1:

set protocols pim rp local address 10.0.0.255
set protocols msdp group any-cast peer 10.0.0.2
set protocols msdp group any-cast local-address 10.0.0.1

MSDP anycast router 2:

set protocols pim rp local address 10.0.0.255
set protocols msdp group any-cast peer 10.0.0.1
set protocols msdp group any-cast local-address 10.0.0.2

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 1: Base configuration

1) Load the chapter 5 base configuration on all devices:

```
lab@srx1#load override terminal
```

Use Ctrl+D at the end, commit and reload the SRXs.

2) Load base configuration of VR-device and AS switch with username lab and password lab123:

```
lab@vrdevice#load override terminal
```

Use Ctrl+D at the end, commit and reload VR.

3) Configure additional interfaces:

On srx1:

```
lab@srx1#set interfaces ge-0/0/4 unit 40 description "Connection to
Src1"
lab@srx1#set interfaces ge-0/0/4 unit 40 vlan-id 40
lab@srx1#set interfaces ge-0/0/4 unit 40 family inet address
192.168.40.254/24
lab@srx1#set interfaces ge-0/0/4 unit 40 family inet6
```

On srx2:

```
lab@srx2#set interfaces ge-0/0/4 unit 46 description "Connection to
Src2"
lab@srx2#set interfaces ge-0/0/4 unit 46 vlan-id 46
lab@srx2#set interfaces ge-0/0/4 unit 46 family inet6 address
fd08:a131:17d8:46::1/64
```

4) Configure VLAN 1400 on ex2 and ex4:

On ex2:

```
lab@ex2#set vlans Delta vlan-id 1400
lab@ex2#set vlans Delta l3-interface vlan.1400
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
lab@ex2#set interfaces vlan unit 1400 family inet address
192.168.32.254/24
```

On ex4:

```
lab@ex4#set vlans Delta vlan-id 1400

lab@ex4#set vlans Delta interface ge-0/0/9.0

lab@ex4#set interfaces ge-0/0/9 unit 0 family ethernet-switching
```

5) VLAN 1400 should be exported to iBGP. The other subnets should be exported to OSPFv3:

```
lab@srx1#set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface
ge-0/0/4.40 passive

lab@srx2#set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.46
passive

lab@ex2#set policy-options policy-statement aggregate.multicast term 1
from protocol aggregate

lab@ex2#set policy-options policy-statement aggregate.multicast term 1
from route-filter 192.168.32.0/24 exact

lab@ex2#set policy-options policy-statement aggregate.multicast term 1
then accept

lab@ex2#set protocols bgp group internal4 export aggregate.multicast

lab@ex2#set protocols bgp group internal4 neighbor 172.30.15.8 export
aggregate.multicast
```

As this neighbor has an export policy applied on the group level already, we have to configure the export policy '**aggregate.multicast**' on the neighbor level for this peer. This is because export policies configured on the neighbor level take precedence over export policies configured on the group level.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 2: Multicast Configuration

1) Enable igmp-snooping on ex4 for VLAN1400. Make sure that interface to multicast first hop router is statically configured. Ensure that the switch stops forwarding traffic immediately after the last multicast group member leaves the group:

IGMP snooping gives a device insight into the exchange of IGMP traffic between hosts and routers. This additional insight can be used to make more informed decisions. We enable igmp-snooping with the immediate-leave configuration option. This enables host tracking and causes for the switch to be able to determine when the last host interested in a particular multicast group leaves that multicast group. When that happens, the switch immediately stops forwarding multicast traffic to the interface.

Additionally, to statically configure the interface towards the multicast router, we configure ge-0/0/19.0 with the 'multicast-router-interface' configuration command:

```
lab@ex4#set protocols igmp-snooping vlan Delta immediate-leave

lab@ex4#set protocols igmp-snooping vlan Delta interface ge-0/0/19.0
multicast-router-interface
```

```
root@ex4# run show igmp-snooping membership detail
VLAN: Delta Tag: 1400 (Index: 3)
    Router interfaces:
        ge-0/0/9.0 dynamic Uptime: 01:19:16 timeout: 248
        ge-0/0/19.0 static Uptime: 01:21:39
```

2) Make sure that IGMP/MLD is configured explicitly on the interfaces to the receivers. Configure PIM version 2 using sparse mode on all devices. Make sure that PIM neighbors are established for IPv4 and IPv6. PIM should not be running on the OoB interfaces.

In our lab topology, we do not have any real receiver hosts. Therefore, we configure various devices with a static IGMP group membership. This way, we will be able to test multicast forwarding later on:

On ex2:

root@ex2# **set protocols igmp interface vlan.1400 static group 239.255.255.2**

On ex3:

root@ex3# **set protocols igmp interface ge-0/0/14.43 static group 239.255.255.3 source 192.168.40.100**

root@ex3# **set protocols igmp interface ge-0/0/14.44 static group 239.255.255.2**

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On srx4:

root@srx4# **set protocols mld interface ge-0/0/4.45 static group ff1e::2010**

We also need to enable PIM. We need to enable sparse mode on every device for all interfaces except for the out of band interface. The easiest way to do this is by first enabling PIM for all interfaces, and then disabling it for the one OoB interface.

We also need to keep in mind that PIM messages need to be accepted by the IPv6 firewall filter.

On all SRX devices:

```
root@srx1# set protocols pim interface all mode sparse

root@srx1# set protocols pim interface all version 2

root@srx1# set protocols pim interface ge-0/0/0.0 disable

root@srx1# set firewall family inet6 filter protect-re-ipv6 term pim
from next-header pim

root@srx1# set firewall family inet6 filter protect-re-ipv6 term pim
then accept

root@srx1# insert firewall family inet6 filter protect-re-ipv6 term
pim before term icmp
```

The OoB interface on the ex differs from the one on the srx. On the ex, we configure the following:**On ex1, ex2:**

```
root@ex1# set protocols pim interface all mode sparse

root@ex1# set protocols pim interface all version 2

root@ex1# set protocols pim interface me0.0 disable
```

PIM verification on all devices:

```
root@srx1# run show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit
Interface           IP V Mode        Option      Uptime Neighbor addr
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
ge-0/0/4.12          4  2                HPLGT       06:18:29 172.30.0.2
ge-0/0/4.13          4  2                HPLGT       06:15:48 172.30.0.6
gr-0/0/0.0           4  2                HPLGT       06:11:05 172.30.0.130
vlan.200             4  2                HPLGT       06:18:27
192.168.30.253
vlan.400             4  2                HPLGT       06:18:00
192.168.31.253
ge-0/0/4.12          6  2                HPLGT       06:18:29
fe80::7a19:f700:c18:3d84
ge-0/0/4.13          6  2                HPLGT       06:15:46
fe80::3e94:d500:d46:7311
gr-0/0/0.0           6  2                HPLGT       06:11:05
fe80::3e8a:b000:5c:7aa0
vlan.300             6  2                HPLGT       06:14:32 fe80::3
```

3) Make sure that srx1, srx2, ex1 and ex2 use inet.2 for IPv4 RPF checks:

By default, the inet.2 table does not contains any routes. To be able to use the inet.2 table for RPF checks, we need to place routes into this table ourselves. To do this, we create the **inet0.to.inet2** rib-group. We will use this rib-group to copy routes from the inet.0 table to the inet.2 table.

When we configure the rib-group, we first import the rib we want to copy the routes from. This will be the inet.0 table. We then configure the rib we want to copy the routes to. This will be the inet.2 table:root@srx1# **set routing-options rib-groups inet0.to.inet2 import-rib inet.0**

root@srx1# **set routing-options rib-groups inet0.to.inet2 import-rib inet.2**

This configuration alone will not do anything. For the rib-group we configured to have any effect, we need to apply it to where we want to copy the routes from. Since we need to copy the interface routes, the iBGP routes and the IPv4 OSPF3 routes, we configure the following:

root@srx1# **set routing-options interface-routes rib-group inet inet0.to.inet2**

root@srx1# **set protocols bgp group internal4 family inet unicast rib-group inet0.to.inet2**

root@srx1# **set protocols ospf3 realm ipv4-unicast rib-group inet0.to.inet2**

We then proceed to configure a second rib-group, the **multicast.rpf** rib-group. We do not need to use this rib-group to copy routes to or from any table. We only need to have this rib-group reference the

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

inet.2 table and apply it to the PIM protocol. By doing so, we instruct PIM to alter the lookup from the default inet.0 table to the inet.2 table:

```
root@srx1# set routing-options rib-groups multicast.rpf import-rib
inet.2

root@srx1# set protocols pim rib-group inet multicast.rpf
```

We perform this configuration on every device that is running PIM.

Verification:

```
root@srx1# run show multicast rpf summary
Multicast RPF table: inet.2 , 74 entries

Multicast RPF table: inet6.0 , 196 entries
```

4) Configure ex3 to convert any IGMPv1/IGMPv2 joins from hosts that do not support IGMPv3 for multicast group G3. Assume that IGMPv1/IGMPv2 hosts are connected to interface ge-0/0/6.0 and ge-0/0/7.0. Make sure that any SSM group address can be used for ASM distribution as well:

First we complete our objective on ex3. SSM mapping can be used to translate IGMPv1/IGMPv2 to IGMPv3. We first configure a policy that matches the relevant multicast group, G3:

```
root@ex3# set policy-options policy-statement ssm.mapping term 1 from
route-filter 239.255.255.3/32 exact

root@ex3# set policy-options policy-statement ssm.mapping term 1 then
accept
```

The newly configured policy is applied under the [**routing-options multicast**] stanza. Here we also need to specify the IP address of src1:

```
root@ex3# set routing-options multicast ssm-map translate.to.ssm
policy ssm.mapping

root@ex3# set routing-options multicast ssm-map translate.to.ssm
source 192.168.40.100
```

Finally, we apply the SSM-map under the [**protocols igmp**] stanza:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3# set protocols igmp interface ge-0/0/6.0 ssm-map
translate.to.ssm

root@ex3# set protocols igmp interface ge-0/0/7.0 ssm-map
translate.to.ssm
```

The second objective requires configuration on every device. Here, we configure the device with an additional SSM-group and we enable the default SSM range for ASM:

On all devices:

```
root@srx1# set routing-options multicast ssm-groups 239.0.0.0/8

root@srx1# set routing-options multicast asm-override-ssm
```

5) Configure the BSR mechanism for IPv4 RP discovery. Nominate srx1 and srx2 as RP candidates. Configure srx1 as BSR router and ex2 as backup. Make sure that other devices have learned the RP via the bootstrap mechanism.

For this task, we first configure both srx1 as well as srx2 to be an RP:

```
root@srx1# set protocols pim rp local family inet address 172.30.15.1

root@srx2# set protocols pim rp local family inet address 172.30.15.2
```

Next is making sure that the rest of the network will learn of these RPs through the use of a BSR. The router with the highest priority (or IP address) is elected to become the bootstrap router. The default bootstrap priority of a Juniper device is 0, making the device ineligible to become the BSR. To make sure that srx1 is elected BSR, we make sure that it is configured with the highest priority:

```
root@srx1# set protocols pim rp bootstrap family inet priority 200
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

To enable the ex2 to be a backup BSR, it is configured with a non-zero priority that is lower then what we previously configured srx1 with:root@ex2# **set protocols pim rp bootstrap family inet priority 100**

Bootstrap Verification:

Multicast routing protocols are slow. It may take a while for the network to converge and for all the device to learn about the RP.

```
root@srx1# run show pim bootstrap
Instance: PIM.master
BSR                  Pri Local address          Pri State
Timeout
172.30.15.1          200 172.30.15.1            200 Elected
10
root@srx2# run show pim bootstrap
Instance: PIM.master
BSR                  Pri Local address          Pri State
Timeout
172.30.15.1          200 172.30.15.2              0 Ineligible
95
root@srx2# run show pim rps
Instance: PIM.master
Address family INET
RP address             Type        Holdtime Timeout Groups Group
prefixes
172.30.15.1            bootstrap        150     103      0
224.0.0.0/4
172.30.15.2            bootstrap        150    None      0
172.30.15.2            static             0    None      0
224.0.0.0/4
root@ex2# run show pim rps
Instance: PIM.master
Address family INET
RP address             Type        Holdtime Timeout Groups Group
prefixes
172.30.15.1            bootstrap        150     121      0
224.0.0.0/4
```

IPv4 multicast verification:

Since our lab topology does not have any real receiver hosts, we configured various devices with a static IGMP group membership. This simulates the existence of interested receivers and it enables us to verify a number of things. For one, we can verify whether or not the joins that are send out make it to the RP. Another thing we will be able to verify is whether or not multicast traffic for the configured groups will be forwarded to the devices that have such a statically configured IGMP group membership.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

In addition to our own statically configured IGMP group membership, the SAP protocol has been configured on VR device. This will not only simulate interested receivers, but it will also make the VR device respond to ICMP packets send to the multicast group that the VR is listening to.

To simulate a multicast stream, we can start a ping from the VR device with a multicast group as a destination address. We can leave this ping on and examine how the multicast stream flows through the rest of the network.

We can also simply generate a few packets to different multicast groups and examine how the network responded to that. In the following example, we start two multicast streams from Src1 on the VR. After this, we verify the multicast configuration and operation:

```
root@vrdevice# run ping routing-instance Src1 ttl 20 bypass-routing
239.255.255.2 interface ge-0/0/1.40
PING 239.255.255.2 (239.255.255.2): 56 data bytes
64 bytes from 192.168.43.10: icmp_seq=0 ttl=61 time=123.066 ms
64 bytes from 192.168.43.10: icmp_seq=1 ttl=61 time=4.006 ms
64 bytes from 192.168.43.10: icmp_seq=2 ttl=61 time=4.907 ms
root@vrdevice# run ping routing-instance Src1 ttl 20 bypass-routing
239.255.255.3 interface ge-0/0/1.40
PING 239.255.255.3 (239.255.255.3): 56 data bytes
64 bytes from 192.168.43.10: icmp_seq=0 ttl=61 time=60.395 ms
64 bytes from 192.168.43.10: icmp_seq=1 ttl=61 time=3.943 ms
root@srx1# run show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.255.255.2
    Source: *
    RP: 172.30.15.1
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 06:14:20
    Downstream neighbors:
        Interface: ge-0/0/4.12
            172.30.0.2 State: Join Flags: SRW Timeout: 171
            Uptime: 01:24:39 Time since last Join: 00:00:39
    Downstream neighbors:
        Interface: ge-0/0/4.13
            172.30.0.6 State: Join Flags: SRW Timeout: 179
            Uptime: 06:14:20 Time since last Join: 00:00:30

Group: 239.255.255.3
    Source: *
    RP: 172.30.15.1
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Upstream state: Local RP
    Uptime: 06:14:20
    Downstream neighbors:
        Interface: ge-0/0/4.13
            172.30.0.6 State: Join Flags: SRW Timeout: 179
            Uptime: 06:14:20 Time since last Join: 00:00:30

Group: 239.255.255.3
    Source: 192.168.40.100
    Flags: sparse
    Upstream interface: ge-0/0/4.40
    Upstream neighbor: Direct
    Upstream state: None, Local Source, Local RP
    Keepalive timeout:
    Uptime: 06:14:30
    Downstream neighbors:
        Interface: ge-0/0/4.13
            172.30.0.6 State: Join Flags: S Timeout: 179
            Uptime: 06:14:30 Time since last Join: 00:00:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

root@ex2# run show pim join extensive inet
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.255.255.2
    Source: *
    RP: 172.30.15.1
    Flags: sparse,rptree,wildcard
    Upstream interface: ge-0/0/14.24
    Upstream neighbor: 172.30.0.9
    Upstream state: Join to RP
    Uptime: 01:25:08
    Downstream neighbors:
        Interface: vlan.1400
            192.168.32.254 State: Join Flags: SRW  Timeout: Infinity
            Uptime: 01:25:08 Time since last Join: 01:25:08

root@ex2# run show multicast route extensive
Instance: master Family: INET

Group: 239.255.255.2
    Source: 192.168.40.100/32
    Upstream interface: ge-0/0/14.24
    Downstream interface list:
        vlan.1400
    Session description: Organisational Local Scope
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 131080
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
```

6) Configure srx1, srx2, ex1 and ex2 to load balance PIM join messages over equal cost path:

PIM join messages are forwarded towards the source based on the RPF check. If there are equal-cost paths, only one interface is chosen. Since only one interface is chosen, the multicast traffic will only use this one path.

To be able to spread multicast traffic across equal-cost paths, we simply enable the device to load balance the PIM join messages. This can be done using only one configuration command:

```
root@srx1# set protocols pim join-load-balance
```

After this has been configured, new PIM join messages will be distributed evenly across equal cost-paths. The command '**show pim interfaces**' or '**show pim neighbors detail**' can be used to verify the configuration. Both commands enable you to see the number of joins that were sent across a link.

7) Configure Anycast-RP for IPv6 multicast on srx1 and srx2. Use Anycast RP address 'fd17:f0f4:f691:255::1'. Configure this address statically on all other devices:

We start this task configuring the Anycast-RP addresses first. We need to add the address that we will use as an Anycast-RP address to the loopback interfaces of srx1 and srx2. We also configure the IPv6 address that is already active on the srx1 and srx2 loopback interface to be the primary address:

On srx1:

```
root@srx1# set interfaces lo0 unit 0 family inet6 address
fd17:f0f4:f691:15::1/128 primary

root@srx1# set interfaces lo0 unit 0 family inet6 address
fd17:f0f4:f691:255::1/128
```

On srx2:

```
root@srx2# set interfaces lo0 unit 0 family inet6 address
fd17:f0f4:f691:15::2/128 primary

root@srx2# set interfaces lo0 unit 0 family inet6 address
fd17:f0f4:f691:255::1/128
```

We then move to the PIM configuration. Here, we configure both srx1 and srx2 to be an RP. We also enable the PIM Anycast mechanism on both srx1 and srx2. We do this by configuring both the **rp-set address** as well as the **local-address**. The **rp-set address** is the non-RP loopback address of the device

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

the router will form an Anycast-RP pair with. The **local-address** is the non-RP loopback address that is configured on the local device.

On srx1:

```
root@srx1# set protocols pim rp local family inet6 address
fd17:f0f4:f691:255::1

root@srx1# set protocols pim rp local family inet6 anycast-pim rp-set
address fd17:f0f4:f691:15::2

root@srx1# set protocols pim rp local family inet6 anycast-pim local-
address fd17:f0f4:f691:15::1
```

On srx2:

```
root@srx2# set protocols pim rp local family inet6 address
fd17:f0f4:f691:255::1

root@srx2# set protocols pim rp local family inet6 anycast-pim rp-set
address fd17:f0f4:f691:15::1

root@srx2# set protocols pim rp local family inet6 anycast-pim local-
address fd17:f0f4:f691:15::2
```

On all other devices:

As per our assignment requirements, the rest of the devices in the network are configured with a statically assigned RP address.

```
root@ex1# set protocols pim rp static address fd17:f0f4:f691:255::1
```

IPv6 multicast verification:

```
root@srx1# run show pim rps inet6
Instance: PIM.master
Address family INET6
RP address                Type        Holdtime Timeout Groups Group
prefixes
fd17:f0f4:f691:255::1     static             0    None      0 ff00::/8

root@srx4# run show pim rps inet6
Instance: PIM.master
Address family INET6
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
RP address                Type        Holdtime Timeout Groups Group
prefixes
fd17:f0f4:f691:255::1     static             0    None      1 ff00::/8

root@vrdevice# run ping routing-instance Src2 ff1e::2010 bypass-
routing ttl
PING6(56=40+8+8 bytes) fd08:a131:17d8:46::100 --> ff1e::2010
^C
--- ff1e::2010 ping6 statistics ---
28 packets transmitted, 0 packets received, 100% packet loss

root@srx2# run show pim join inet6 extensive
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff1e::2010
    Source: *
    RP: fd17:f0f4:f691:255::1
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 05:11:49
    Downstream neighbors:
        Interface: ge-0/0/4.24
            fe80::2e6b:f500:188f:87d1 State: Join Flags: SRW Timeout:
161
            Uptime: 05:11:49 Time since last Join: 00:00:49

Group: ff1e::2010
    Source: fd08:a131:17d8:46::100
    Flags: sparse,spt
    Upstream interface: ge-0/0/4.46
    Upstream neighbor: Direct
    Upstream state: None, Local Source, Local RP
    Keepalive timeout: 357
    Uptime: 00:01:26
    Downstream neighbors:
        Interface: ge-0/0/4.24
            fe80::2e6b:f500:188f:87d1 State: Join Flags: S Timeout:
161
            Uptime: 00:01:26 Time since last Join: 00:00:49

root@srx4# run show pim join inet6 extensive
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff1e::2010
    Source: *
    RP: fd17:f0f4:f691:255::1
    Flags: sparse,rptree,wildcard
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Upstream interface: ge-0/0/4.64
    Upstream neighbor: fe80::2e6b:f500:408f:87d1
    Upstream state: Join to RP
    Uptime: 05:12:52
    Downstream neighbors:
        Interface: ge-0/0/4.45
            fe80::3e61:400:2d8d:be84 State: Join Flags: SRW  Timeout:
Infinity
            Uptime: 05:12:52 Time since last Join: 05:12:52

Group: ff1e::2010
    Source: fd08:a131:17d8:46::100
    Flags: sparse,spt
    Upstream interface: ge-0/0/4.64
    Upstream neighbor: fe80::2e6b:f500:408f:87d1
    Upstream state: None, Join to Source
    Keepalive timeout: 304
    Uptime: 00:02:19
    Downstream neighbors:
        Interface: ge-0/0/4.45
            fe80::3e61:400:2d8d:be84 State: Join Flags: S   Timeout:
Infinity
            Uptime: 00:02:19 Time since last Join: 00:02:19

root@srx4# run show multicast route
Instance: master Family: INET

Instance: master Family: INET6

Group: ff1e::2010
    Source: fd08:a131:17d8:46::100/128
    Upstream interface: ge-0/0/4.64
    Downstream interface list:
        ge-0/0/4.45
```

8) Limit the number of groups learned on interface vlan.1400 to no more than 50:

We can configure a limit for the number of multicast groups that can join an interface using the 'group-limit' command:

```
root@ex2# set protocols igmp interface vlan.1400 group-limit 50
```

9) Disable the sending and receiving of bootstrap messages to and from both ISP1 and ISP2:

We create a policy that matches the ISP facing interface and configure the reject action. We then apply the policy under the [**protocols pim rp bootstrap family inet**] hierarchy in both the import as well as the export direction.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On srx1:

```
root@srx1# set policy-options policy-statement block-bsr term 1 from
interface ge-0/0/4.14

root@srx1# set policy-options policy-statement block-bsr term 1 then
reject

root@srx1# set protocols pim rp bootstrap family inet import block-bsr

root@srx1# set protocols pim rp bootstrap family inet export block-bsr
```

On srx2:

```
root@srx2# set policy-options policy-statement block-bsr term 1 from
interface ge-0/0/4.22

root@srx2# set policy-options policy-statement block-bsr term 1 then
reject

root@srx2# set protocols pim rp bootstrap family inet import block-bsr

root@srx2# set protocols pim rp bootstrap family inet export block-bsr
```

10) Ensure that RPs in ISP1 or ISP2 cannot become the RP for your AS. Reject the multicast data traffic from group-range 239/8 received from ISP1 and ISP2. You are not allowed to use policies to accomplish this requirement.

Since we cannot use policies, we turn to multicast administrative scoping. By configuring a multicast scope, we can control the propagation of multicast messages as well as data-forwarding.

For this assignment, we define two different scopes. One scope for the groups used by auto-RP and another for the local multicast groups. We configure the groups for the relevant prefixes and we specify the boundary of the administrative scope by configuring the ISP interface:

On srx1:

```
root@srx1# set routing-options multicast scope auto-rp-discovery
prefix 224.0.1.40/32

root@srx1# set routing-options multicast scope auto-rp-discovery
interface ge-0/0/4.14
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set routing-options multicast scope auto-rp-announce prefix
224.0.1.39/32

root@srx1# set routing-options multicast scope auto-rp-announce
interface ge-0/0/4.14

root@srx1# set routing-options multicast scope local-groups prefix
239.0.0.0/8

root@srx1# set routing-options multicast scope local-groups interface
ge-0/0/4.14
```

On srx2:

```
root@srx2# set routing-options multicast scope auto-rp-discovery
prefix 224.0.1.40/32

root@srx2# set routing-options multicast scope auto-rp-discovery
interface ge-0/0/4.22

root@srx2# set routing-options multicast scope auto-rp-announce prefix
224.0.1.39/32

root@srx2# set routing-options multicast scope auto-rp-announce
interface ge-0/0/4.22

root@srx2# set routing-options multicast scope local-groups prefix
239.0.0.0/8

root@srx2# set routing-options multicast scope local-groups interface
ge-0/0/4.22
```

Verification:

```
root@srx1# run show multicast scope

Resolve
Scope name           Group prefix       Interface
rejects
auto-rp-announce     224.0.1.39/32      ge-0/0/4.14
inactive
auto-rp-discovery    224.0.1.40/32      ge-0/0/4.14
inactive
local-groups         239.0.0.0/8        ge-0/0/4.14
inactive
local                239.255.0.0/16     ge-0/0/4.14
inactive
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Chapter 6: Class of Service

A JNCIE-ENT candidate should be familiar with the following topics:

- Classification
- Loss priority
- Scheduling
- Drop profiles and RED
- Rewrite
- Policing
- CoS-Based Forwarding
- EX switch CoS templates

COS overview.

In today's networking world, numerous applications and protocols are using the same packet switched network. These applications and protocols are very different in nature, but they have to share a common infrastructure.

Packet loss, delay and jitter can severely impact these applications and protocols. A set of QoS mechanisms can be used to give preferential treatment to sensitive applications and/or protocols. To be able to use the different QoS mechanisms on Juniper devices, we use the Class of Service (CoS) configuration construct.

The first CoS processing task that is performed on every router in the network is input traffic classification. By distributing packets in different virtual containers called Forwarding Classes (FC), we can define the order in which packets leave the router at a later stage. Packet classification is performed by means of a BA classifier or multi-field (MF) classifier.

A BA classifier examines the headers of incoming packets and then determines the Forwarding Class of those packets. The classification can be based on the value of any of the following fields: DSCP, IP Precedence and 802.1p. You can also define one more locally significant parameter called Packet Loss Priority (PLP). There are 4 possible values: high, medium-high, medium-low and low. PLP only comes into play during times of congestion and it can be used to determine what packets are dropped first. By default, a BA classifier is applied to every logical interface that is configured on the device. This can be changed by configuring and applying another BA classifier.

The MF classifier offers a more granular approach as you can use numerous firewall filter parameters as match criteria. The MF classifier can be used to assign packets to different Forwarding Classes and it can be used to set PLP values. The MF classifier is commonly used on edge devices. Since they are processed after BA classifiers, the MF classifier can be used to overwrite any previous classifications.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Policing is a mechanism that can be used to limit the amount of traffic (mostly incoming) for a configured bandwidth value. Policing uses a token bucket algorithm which allows for traffic bursts as long as there are enough tokens in the bucket. Tokens are accumulated during physical interface silence and spent when the interface transmits packets at wire speed. Policers can be set to a hard-limit, which means that out-of-profile traffic is dropped. You can also configure soft-limit Policers. When a soft-limit Policer is used, out-of-profile traffic is assigned to a different (or worse) forwarding class and it is given a higher PLP.

Before transit traffic is sent through the fabric of the device, it can be subjected to CoS-based forwarding. This mechanism changes the way in which the outgoing interface is identified. Instead of a traditional destination IP lookup, the router uses class-to-next-hop mappings to forward traffic destined to a configured prefix. By using this approach, you can 'push' traffic across a more 'reliable' and 'faster' path.

Schedulers define how outgoing interface queues will be serviced. Every physical interface has 4 or 8 hardware queues that are mapped to forwarding classes. Every scheduler has a priority, interface bandwidth and buffer size configured. Each queue is serviced on a round-robin basis (from the highest priority to the lowest) according to its configured weight. That means that every Forwarding Class has a guaranteed rate (CIR) while unallocated bandwidth can be shared equally between other queues. You can also reference drop profiles in the JUNOS scheduler configuration, which can be used to enable the congestion avoidance mechanism called WRED.

As packets leave the interface of the device, the rewrite rule is the last mechanism applied during CoS processing. The rewrite rule can be used to set/ change different CoS field values according to the packet's forwarding class and PLP.

Classification

The BA and the multifield classifier are applied at different stages during CoS processing, and they are configured in different configuration stanzas as well.

BA classifier: The BA classifier is configured, and applied, under the [**class-of-service classifiers**] stanza. In this stanza, you first choose what type of classifier you want to configure and then you name the classifier. Inside the classifier configuration, you simply associate forwarding classes with code-points and a loss-priority.

The following creates a classifier named IPv4, and associates a forwarding class called BE with a loss-priority of low and a code-point of '000000':

set class-of-service classifiers dscp IPV4 forwarding-class BE loss-priority low code-points 000000

You can, optionally, configure code-point-aliases under the [**class-of-service code-point-aliases**] stanza. This way, you can reference names instead of code-point values.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Multifield classifier: The multifield classifier is configured using a firewall filter. Firewall filters can be configured for different address families, and the match conditions offer a lot more granularity then the BA classifier. Some match conditions include source IP, destination IP, source port, destination port and protocol. After specifying the match criteria (using **from**), the '**then forwarding-class**' can be used to actually classify the packet.

On EX, you can also easily classify VoIP VLAN traffic by configuring the [**ethernet-switching-options voip interface xx forwarding-class xxx**] command.

Loss Priority

Packet Loss Priority (**PLP**) can be used to decide what traffic is to be dropped first when congestion occurs. The PLP can be set using multifield classifiers, BA classifiers and policers. Usually, a higher loss priority is configured to be dropped first during congestion. How exactly this is done can be specified using a drop profile.

Scheduling

Schedulers, configured under the [**class-of-service schedulers**] stanza, define how traffic inside a single queue is treated. Inside this configuration stanza, we can defined the **transmission rate**, **queue priority**, **buffer size** and **drop profile** that is associated with the scheduler.

Scheduler maps, configured under the [**class-of-service scheduler-maps**] stanza, can be used to apply the schedulers to a forwarding class. Scheduler-maps can be applied a logical interface directly, or they can be applied using a traffic-control-profile.

Example configuration of a scheduler called '**silver-sched'**:

set class-of-service schedulers **silver-sched** transmit-rate percent 20
set class-of-service schedulers **silver-sched** buffer-size percent 16
set class-of-service schedulers **silver-sched** priority medium-high
set class-of-service schedulers **silver-sched** drop-profile-map loss-priority high protocol any drop-profile high

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

A scheduler-map, called '**example**', that applies several schedulers to forwarding classes:

set class-of-service scheduler-maps **example** forwarding-class bronze scheduler **bronze-sched**
set class-of-service scheduler-maps **example** forwarding-class silver scheduler **silver-sched**
set class-of-service scheduler-maps **example** forwarding-class gold scheduler **gold-sched**

Drop profiles and RED

By default, Junos applies a 100% drop rate when there is a 100% fill. To start using Weighting RED (**WRED**), we have to define a drop profile. The drop profile is configured under the [**class-of-service drop-profiles**] stanza. Inside this configuration stanza, you can configure the drop probability for a fill-level.

By default, the drop profile is **segmented**. This means the drop profile will apply the drop probability as soon as a fill level is reached. It will do so up until the fill level rises to another configured level. When we specify the drop profile to be **interpolated**, we get a smooth graph. In this case, Junos automatically interpolates the different values that were configured.

For a drop profile to become active, it needs to be referenced in a scheduler.

The following is an example of an interpolated drop profile called '**HIGH**':

set class-of-service drop-profiles **HIGH** interpolate fill-level 50
set class-of-service drop-profiles **HIGH** interpolate fill-level 75
set class-of-service drop-profiles **HIGH** interpolate fill-level 85
set class-of-service drop-profiles **HIGH** interpolate drop-probability 0
set class-of-service drop-profiles **HIGH** interpolate drop-probability 50
set class-of-service drop-profiles **HIGH** interpolate drop-probability 100

Rewrite rules

Rewrite rules are used to mark packets on egress. Downstream nodes will be able to perform classification using only the BA classifier. The purpose of the rewrite rules is to convey a packet CoS profile to the next-hop router based on forwarding class and PLP. Rewrite rules are configured in the [**class-of-service rewrite-rules**] stanza and they are applied under the [**class-of-service interfaces**] stanza.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Policing

The Junos OS supports the following policers:

- Single-rate two-color

This policer consists of a CIR (Committed Information Rate) and a CBS (Committed Burst Size). Traffic exceeding the CIR and CBS can be marked or discarded.

- Single-rate tricolor

This policer uses the CIR, the CBS. It also introduces an additional burst size, the EBS (Excess Burst Size). The three different colors related to this policer are green, yellow and red. Green traffic is all traffic that falls within the CIR + CBS limit. Yellow exceeds the CBS, but falls within the limits set by the EBS. Red is traffic that exceeds the EBS. Different packet loss priorities can be applied to each of these colors.

- Two-rate tricolor

A two-rate tricolor policers takes the CIR and CBS from the single-rate two color policer, and adds both a Peak Information Rate (PIR) and a Peak Burst Size (PBS). Green is what conforms to the CIR + CBS. Yellow exceeds this, but it does not exceed the PIR + CBS. Red is what exceeds PIR + PBS.

A single-rate tricolor policer offers an additional burst-size over a single-rate two-color policer. A two-rate tricolor offers an additional rate over the additional burst size that a single-rate tricolor policer offers.

In the following example, we configure a single-rate two-color policer:

```
[edit firewall]
policer policer-name {
    if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    then  {
        policer-action (discard, set forwarding class, set PLP)
    }
}
```

We can apply the policer directly on an interface (under the family hierarchy) or within a firewall filter using the **then** statement. By default, when we reference a policer in a firewall term, the policer is 'term specific'. This means that the policer is an instance for that specific term. When can also define a 'filter-specific' policer. In this case, a single instance of a policer will police several terms.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

CoS-Based Forwarding

CoS-Based Forwarding configuration requires the following configuration steps:

A) Configuring a [**class-of-service forwarding-policy next-hop map]**
B) Configuring a **routing-policy** to match interesting routes (**from protocol**, **route-filter** etc.) and point it to the previously configured forwarding policy (**then cos-next-hop-map**)
C) Exporting the configured routing policy to the **forwarding-table**

EX switch CoS templates

EX switches come with several predefined CoS templates. They are called EZQoS VoIP templates and they are stored in **/etc/config/ezqos-voip.conf**

To be able to use these templates, we need to follow these steps:

1- **load merge /etc/config/ezqos-voip-conf**
2- **apply-groups ezqos-voip**
3- Apply the scheduler-map, classifier and rewrite rules to the relevant interfaces

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 1: Base Network Deployment

1) Load base configuration:

On all devices:

```
lab@srx1# load override terminal
```

Enter Ctrl+D at the end, commit and reload the device.

2) Use username lab and password lab123 on the VR-device and the AS switch:

```
lab@srx1# load override terminal
```

Enter Ctrl+D at the end, commit and reload the device.

3) Configure additional interfaces and configure OSPFv3 to advertise these interfaces:

On srx1 and srx2:

```
lab@srx1# set interfaces ge-0/0/4 unit 3001 description "Connection to
Video Server"

lab@srx1# set interfaces ge-0/0/4 unit 3001 vlan-id 3001

lab@srx1# set interfaces ge-0/0/4 unit 3001 family inet address
192.168.231.1/24

lab@srx1# set interfaces ge-0/0/4 unit 3001 family inet6 address
fd08:a131:17d8:231::1/64

lab@srx1# set protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.3001 passive

lab@srx1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.3001
passive
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 2: SRX Class of Service

1) Configure forwarding classes and map them to the output queues based on the task:

We use the following configuration to name our forwarding classes and map them to the proper output queue. We do this on all srx devices:

On srx devices:

```
lab@srx1# set class-of-service forwarding-classes queue 0 best-effort

lab@srx1# set class-of-service forwarding-classes queue 1 server-
bronze

lab@srx1# set class-of-service forwarding-classes queue 2 server-
silver

lab@srx1# set class-of-service forwarding-classes queue 3 server-gold

lab@srx1# set class-of-service forwarding-classes queue 5 voice

lab@srx1# set class-of-service forwarding-classes queue 7 control
```

Verification

```
root@srx1# run show class-of-service forwarding-class
Forwarding class                  ID                        Queue
Policing priority    SPU priority
  best-effort                              0                        0
normal           low
  server-bronze                            1                        1
normal           low
  server-silver                            2                        2
normal           low
  server-gold                              3                        3
normal           low
  voice                                    5                        5
normal           low
  control                                  7                        7
normal           low
```

2) Configure schedulers and a scheduler map on all SRX devices:

In the previous task, we named our forwarding classes and mapped them to an output queue. We did not define how traffic inside each queue should be treated. To define how traffic inside a queue should be treated, we use schedulers. Schedulers are configured inside the [**class-of-service schedulers**

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

scheduler-name] configuration stanza. In this configuration stanza, we can define our schedulers with the following parameters:

Transmit-rate:

With the '**transmit-rate rate *bps***' configuration command, we specify the rate in bits per second. If we use the '**transmit-rate percent *percent* **' configuration command, we allocate the scheduler a percentage of the available interface bandwidth.

It is important to note that these limits apply during times of congestion. When there is no congestion on an interface, schedulers configured with a transmit-rate will be able to use the maximum bandwidth. If we were to add the optional **rate-limit** configuration command to a scheduler, the configured transmit rate would become the limit. Even when there is no congestion.

As an alternative to the '**transmit-rate rate**' and the '**transmit-rate percent**' statements, we can also use the '**transmit-rate remainder**' option. This simply means that the scheduler is assigned all bandwidth that is not yet configured to be allocated to other schedulers.

Buffer-size:

Schedulers can be allocated a buffer-size. We can allocate the buffer size as a percentage of the total buffer using the '**buffer-size percent *percent***' statement. If we want to choose a temporal value instead of a percentage, we can use the '**buffer-size temporal *microseconds***' to specify the buffer in microseconds.

Same as with the transmit rate, here we can also choose to allocate the remaining buffer using the '**buffer-size remainder**' configuration command.

Priority:

Used to assign relative priorities to queues. We have five priorities available on JUNOS devices:

- priority *low*
- priority *medium-low*
- priority *medium-high*
- priority *high*
- priority *strict-high*

Note that the strict high queue is always services before traffic in other queues. This might lead to traffic starvation for the other non strict-high queues. If you want to prevent starvation in other queues, you should rate-limit this queue.

Drop profile:

The drop profile itself is configured under the [**class-of-service drop-profiles**] stanza. Inside the scheduler configuration, you can assign the drop profile to a specific packet loss priority or to a specific

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

protocol. This is done by using the ' **drop-profile-map loss-priority *loss-priority* protocol *protocol* drop-profile *profile-name***' configuration command.

We proceed to configure the schedulers with the parameters given in table 9. We configure the following on all srx devices:

```
root@srx1# set class-of-service schedulers control-sched transmit-rate
percent 5

root@srx1# set class-of-service schedulers control-sched buffer-size
percent 10

root@srx1# set class-of-service schedulers control-sched priority
medium-high

root@srx1# set class-of-service schedulers voice-sched transmit-rate
percent 10

root@srx1# set class-of-service schedulers voice-sched buffer-size
temporal 20k

root@srx1# set class-of-service schedulers voice-sched priority
strict-high

root@srx1# set class-of-service schedulers server-gold-sched transmit-
rate percent 25

root@srx1# set class-of-service schedulers server-gold-sched buffer-
size percent 10

root@srx1# set class-of-service schedulers server-gold-sched buffer-
size exact

root@srx1# set class-of-service schedulers server-gold-sched priority
high

root@srx1# set class-of-service schedulers server-silver-sched
transmit-rate percent 35

root@srx1# set class-of-service schedulers server-silver-sched buffer-
size percent 50

root@srx1# set class-of-service schedulers server-silver-sched buffer-
size exact

root@srx1# set class-of-service schedulers server-silver-sched
priority medium-high
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set class-of-service schedulers server-silver-sched drop-profile-map loss-priority high protocol any drop-profile relaxed

root@srx1# set class-of-service schedulers server-bronze-sched transmit-rate remainder

root@srx1# set class-of-service schedulers server-bronze-sched buffer-size remainder

root@srx1# set class-of-service schedulers server-bronze-sched priority medium-low

root@srx1# set class-of-service schedulers server-bronze-sched drop-profile-map loss-priority low protocol any drop-profile relaxed

root@srx1# set class-of-service schedulers server-bronze-sched drop-profile-map loss-priority high protocol any drop-profile aggresive

root@srx1# set class-of-service schedulers best-effort-sched transmit-rate remainder

root@srx1# set class-of-service schedulers best-effort-sched buffer-size remainder

root@srx1# set class-of-service schedulers best-effort-sched priority low

root@srx1# set class-of-service schedulers best-effort-sched drop-profile-map loss-priority any protocol any drop-profile aggressive
```

To be able to use the schedulers, we need to configure a scheduler-map and apply that scheduler-map to the interfaces. Inside a scheduler map, we can assign the schedulers we previously configured to a specific forwarding class:

```
root@srx1# set class-of-service scheduler-maps cos-sched-map forwarding-class control scheduler control-sched

root@srx1# set class-of-service scheduler-maps cos-sched-map forwarding-class voice scheduler voice-sched

root@srx1# set class-of-service scheduler-maps cos-sched-map forwarding-class server-gold scheduler server-gold-sched

root@srx1# set class-of-service scheduler-maps cos-sched-map forwarding-class server-silver scheduler server-silver-sched
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set class-of-service scheduler-maps cos-sched-map
forwarding-class server-bronze scheduler server-bronze-sched

root@srx1# set class-of-service scheduler-maps cos-sched-map
forwarding-class best-effort scheduler best-effort-sched
```

We will apply the scheduler map later on.

Note that you will not be able to commit the configuration until after the configuration of the drop profile that is referenced by several of the schedulers.

3) Configure two interpolated drop profiles listed in table 10:

By default, a drop-profile is segmented. This means that a graph is created that will apply the drop probability as soon as the configured fill level is reached. For this assignment, we need to configure both drop profiles to be interpolated. This means that JUNOS will automatically interpolates the different values we configure to create a smooth graph.

We configure the drop profiles as follows:

```
root@srx1# set class-of-service drop-profiles relaxed interpolate
fill-level 75

root@srx1# set class-of-service drop-profiles relaxed interpolate
fill-level 90

root@srx1# set class-of-service drop-profiles relaxed interpolate
fill-level 100

root@srx1# set class-of-service drop-profiles relaxed interpolate
drop-probability 0

root@srx1# set class-of-service drop-profiles relaxed interpolate
drop-probability 75

root@srx1# set class-of-service drop-profiles relaxed interpolate
drop-probability 100

root@srx1# set class-of-service drop-profiles aggresive interpolate
fill-level 50

root@srx1# set class-of-service drop-profiles aggresive interpolate
fill-level 75

root@srx1# set class-of-service drop-profiles aggresive interpolate
fill-level 90
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set class-of-service drop-profiles aggresive interpolate
fill-level 100

root@srx1# set class-of-service drop-profiles aggresive interpolate
drop-probability 0

root@srx1# set class-of-service drop-profiles aggresive interpolate
drop-probability 50

root@srx1# set class-of-service drop-profiles aggresive interpolate
drop-probability 75

root@srx1# set class-of-service drop-profiles aggresive interpolate
drop-probability 100
```

After configuring the drop profile, we can commit and verify our configuration. We can verify the configuration by closely comparing the following output to the objectives stated in our assignment:

```
root@srx1# run show class-of-service scheduler-map

< output omitted >

Scheduler map: cos-sched-map, Index: 8778

  Scheduler: best-effort-sched, Forwarding class: best-effort, Index:
2582
    Transmit rate: remainder, Rate Limit: none, Buffer size:
remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         27372    aggresive
      Medium low      any         27372    aggresive
      Medium high     any         27372    aggresive
      High            any         27372    aggresive

  Scheduler: server-bronze-sched, Forwarding class: server-bronze,
Index: 13758
    Transmit rate: remainder, Rate Limit: none, Buffer size:
remainder,
    Buffer Limit: none, Priority: medium-low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any          8233    relaxed
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
      High            any          27372    aggresive

  Scheduler: server-silver-sched, Forwarding class: server-silver,
Index: 34223
    Transmit rate: 35 percent, Rate Limit: none, Buffer size: 50
percent,
    Buffer Limit: none, Priority: medium-high
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium low      any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any        8233    relaxed

  Scheduler: server-gold-sched, Forwarding class: server-gold, Index:
40991
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 10
percent,
    Buffer Limit: none, Priority: high
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium low      any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

  Scheduler: voice-sched, Forwarding class: voice, Index: 47632
    Transmit rate: 10 percent, Rate Limit: none, Buffer size: 20000
us,
    Buffer Limit: none, Priority: strict-high
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium low      any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

  Scheduler: control-sched, Forwarding class: control, Index: 34455
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 10
percent,
    Buffer Limit: none, Priority: medium-high
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium low      any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

4) Apply the schedulers to all physical interfaces:

To achieve our objective, we need to apply the scheduler-map to all of the interfaces. The easiest and quickest way to do this, is by configuring a wildcard, like this:

```
root@srx1# set class-of-service interfaces ge-* scheduler-map cos-sched-map
```

Verification:

```
root@srx1# run show class-of-service interface ge-0/0/4
Physical interface: ge-0/0/4, Index: 139
Queues supported: 8, Queues in use: 6
  Scheduler map: cos-sched-map, Index: 8778
  Congestion-notification: Disabled
```

5) Configure IPv4 and IPv6 Behavior Aggregate classifiers. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that the default DSCP classifiers are used.

For this task, we need to configure a BA classifier for IPv4 as well as for IPv6. The BA classifiers is a simple and efficient way to classify traffic that is received on an interface. The BA classifier will classify all traffic based on the existing markings.

We configure our classifiers with the '**import default**' configuration statement. By doing this, we import the system default classifier mapping configuration as a template into our own classifier. Because of this, anything we have not specifically configured in our own classifier will be classified according to the values present in the default classifier.

We then proceed to configure the values as they are mentioned in the task. Any of the entries we configure ourselves will be added to the entries that were already imported. If any of the code-points that are explicitly configured in the classifiers match a value that is imported, the explicitly configured value will replace the value that was imported.

We first configure the DSCP classifier for IPv4:

```
root@srx1# set class-of-service classifiers dscp ba-classifier import default
```

```
root@srx1# set class-of-service classifiers dscp ba-classifier forwarding-class control loss-priority low code-points nc1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set class-of-service classifiers dscp ba-classifier
forwarding-class control loss-priority low code-points nc2

root@srx1# set class-of-service classifiers dscp ba-classifier
forwarding-class voice loss-priority low code-points ef

root@srx1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-gold loss-priority low code-points af41

root@srx1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-silver loss-priority low code-points af21

root@srx1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-silver loss-priority high code-points af23

root@srx1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-bronze loss-priority low code-points af11

root@srx1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-bronze loss-priority high code-points af13

root@srx1# set class-of-service classifiers dscp ba-classifier
forwarding-class best-effort loss-priority low code-points be
```

After IPv4, we configure the classifier for IPv6:

```
root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
import default

root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class control loss-priority low code-points nc1

root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class control loss-priority low code-points nc2

root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class voice loss-priority low code-points ef

root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-gold loss-priority low code-points af41

root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-silver loss-priority low code-points af21

root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-silver loss-priority high code-points af23
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-bronze loss-priority low code-points af11

root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-bronze loss-priority high code-points af13

root@srx1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class best-effort loss-priority low code-points be
```

6) Apply the classifiers to all logical interfaces:

The easiest and quickest way to apply the classifiers is by using the same strategy as the one we used when we applied the scheduler-map. Keep in mind that we need to apply the classifier to all logical interfaces. Because of this, we use a wildcard when we specify the interface as well as the unit:

```
root@srx1# set class-of-service interfaces ge-* unit * classifiers
dscp ba-classifier

root@srx1# set class-of-service interfaces ge-* unit * classifiers
dscp-ipv6 ba-classifier-v6
```

Verification:

```
root@srx1# run show class-of-service interface ge-0/0/4
Physical interface: ge-0/0/4, Index: 139
Queues supported: 8, Queues in use: 6
  Scheduler map: cos-sched-map, Index: 8778
  Congestion-notification: Disabled

  Logical interface: ge-0/0/4.101, Index: 85
    Object                  Name                    Type
Index
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786

  Logical interface: ge-0/0/4.12, Index: 81
    Object                  Name                    Type
Index
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786

  Logical interface: ge-0/0/4.13, Index: 82
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Object                  Name                    Type
Index
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786

  Logical interface: ge-0/0/4.14, Index: 83
    Object                  Name                    Type
Index
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786

  Logical interface: ge-0/0/4.3001, Index: 86
    Object                  Name                    Type
Index
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786

  Logical interface: ge-0/0/4.32767, Index: 87

  Logical interface: ge-0/0/4.40, Index: 84
    Object                  Name                    Type
Index
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786
```

7) Configure the rewrite rules to mark the outgoing traffic with the correct DSCP value on the outgoing interfaces of all SRXs:

The forwarding class of a packet is determined by a BA classifier or a Multifield classifier. During the different CoS stages, the forwarding class can be altered (by a policer for example). When traffic leaves the device, the rewrite rule checks the forwarding class that the traffic belongs to. Based on this value, the rewrite-rule does a lookup and writes the configured CoS value into the packet header.

All we need to do here is configure the rewrite-rule based on the table that is provided:

```
root@srx1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class control loss-priority low code-point nc1

root@srx1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class control loss-priority high code-point nc2
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class voice loss-priority low code-point ef

root@srx1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-gold loss-priority low code-point af41

root@srx1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-silver loss-priority low code-point af21

root@srx1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-silver loss-priority high code-point af23

root@srx1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-bronze loss-priority low code-point af11

root@srx1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-bronze loss-priority high code-point af13

root@srx1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class best-effort loss-priority high code-point be

root@srx1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class control loss-priority low code-point nc1

root@srx1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class control loss-priority high code-point nc2

root@srx1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class voice loss-priority low code-point ef

root@srx1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-gold loss-priority low code-point af41

root@srx1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-silver loss-priority low code-point af21

root@srx1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-silver loss-priority high code-point af23

root@srx1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-bronze loss-priority low code-point af11

root@srx1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-bronze loss-priority high code-point af13

root@srx1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class best-effort loss-priority high code-point be
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification:

```
root@srx1# run show class-of-service rewrite-rule

Rewrite rule: rewrite-v4, Code point type: dscp, Index: 4683
  Forwarding class                Loss priority       Code point
  best-effort                     high                000000
  server-bronze                   low                 001010
  server-bronze                   high                001110
  server-silver                   low                 010010
  server-silver                   high                010110
  server-gold                     low                 100010
  voice                           low                 101110
  control                         low                 110000
  control                         high                111000

Rewrite rule: rewrite-v6, Code point type: dscp-ipv6, Index: 37450
  Forwarding class                Loss priority       Code point
  best-effort                     high                000000
  server-bronze                   low                 001010
  server-bronze                   high                001110
  server-silver                   low                 010010
  server-silver                   high                010110
  server-gold                     low                 100010
  voice                           low                 101110
  control                         low                 110000
  control                         high                111000
```

8) Apply the rewrite rules on all logical interfaces:

We use the same strategy here as when we configured the classifiers:

```
root@srx1# set class-of-service interfaces ge-* unit * rewrite-rules
dscp rewrite-v4

root@srx1# set class-of-service interfaces ge-* unit * rewrite-rules
dscp-ipv6 rewrite-v6
```

Verification:

```
root@srx1# run show class-of-service interface ge-0/0/4
Physical interface: ge-0/0/4, Index: 139
Queues supported: 8, Queues in use: 6
  Scheduler map: cos-sched-map, Index: 8778
  Congestion-notification: Disabled
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Logical interface: ge-0/0/4.101, Index: 85
    Object                  Name                    Type
Index
    Rewrite                 rewrite-v4              dscp
4683
    Rewrite                 rewrite-v6              dscp-ipv6
37450
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786

  Logical interface: ge-0/0/4.12, Index: 81
    Object                  Name                    Type
Index
    Rewrite                 rewrite-v4              dscp
4683
    Rewrite                 rewrite-v6              dscp-ipv6
37450
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786

  Logical interface: ge-0/0/4.13, Index: 82
    Object                  Name                    Type
Index
    Rewrite                 rewrite-v4              dscp
4683
    Rewrite                 rewrite-v6              dscp-ipv6
37450
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786

  Logical interface: ge-0/0/4.14, Index: 83
    Object                  Name                    Type
Index
    Rewrite                 rewrite-v4              dscp
4683
    Rewrite                 rewrite-v6              dscp-ipv6
37450
    Classifier              ba-classifier           dscp
19373
    Classifier              ba-classifier-v6        dscp-ipv6
61786

  Logical interface: ge-0/0/4.3001, Index: 86
    Object                  Name                    Type
Index
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Rewrite                 rewrite-v4            dscp
4683
    Rewrite                 rewrite-v6            dscp-ipv6
37450
    Classifier              ba-classifier         dscp
19373
    Classifier              ba-classifier-v6      dscp-ipv6
61786

  Logical interface: ge-0/0/4.32767, Index: 87

  Logical interface: ge-0/0/4.40, Index: 84
    Object                  Name                  Type
Index
    Rewrite                 rewrite-v4            dscp
4683
    Rewrite                 rewrite-v6            dscp-ipv6
37450
    Classifier              ba-classifier         dscp
19373
    Classifier              ba-classifier-v6      dscp-ipv6
61786
```

9) Configure the srx1 interface facing video server to classify incoming IPv4 traffic using a Multifield classifier, based on the criteria provided :

Multifield classifiers provide the most granular classification possible. It involves the configuration of a firewall filter to match interesting traffic. After matching the relevant traffic, the **'then forwarding-class'** statement can be used to assign the traffic to the desired forwarding-class.

Mind the fact that a firewall filter has a default last term that rejects all traffic. In this case, we configure a final term that accepts all traffic. Also make sure that you apply the MFC in the proper direction:

```
root@srx1# set firewall family inet filter mf-video term control-
traffic from protocol udp

root@srx1# set firewall family inet filter mf-video term control-
traffic from protocol tcp

root@srx1# set firewall family inet filter mf-video term control-
traffic from port 558

root@srx1# set firewall family inet filter mf-video term control-
traffic then loss-priority low

root@srx1# set firewall family inet filter mf-video term control-
traffic then forwarding-class control
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set firewall family inet filter mf-video term control-traffic then accept

root@srx1# set firewall family inet filter mf-video term data-traffic from protocol tcp

root@srx1# set firewall family inet filter mf-video term data-traffic from protocol udp

root@srx1# set firewall family inet filter mf-video term data-traffic from port 16384-32767

root@srx1# set firewall family inet filter mf-video term data-traffic then loss-priority low

root@srx1# set firewall family inet filter mf-video term data-traffic then forwarding-class server-gold

root@srx1# set firewall family inet filter mf-video term last then accept

root@srx1# set interfaces ge-0/0/4 unit 3001 family inet filter input mf-video
```

10) Configure the srx2 interface facing the database server to classify incoming traffic using a MF classifier. Configure a single-rate two-color policer to rate limit the database traffic to 300Mbit/s and allow a burst of 3MB. Excess traffic should have the Loss Priority raised to High:

We start out with the configuration of the policer. We need a single-rate two-color policer that raises the loss priority to high when traffic is out of contract. Traffic exceeding the bandwidth-limit should still have the forwarding-class set to 'server-silver' :

```
root@srx2# set firewall policer database-limit if-exceeding bandwidth-limit 300m

root@srx2# set firewall policer database-limit if-exceeding burst-size-limit 3m

root@srx2# set firewall policer database-limit then loss-priority high

root@srx2# set firewall policer database-limit then forwarding-class server-silver
```

We then configure a firewall filter that matches the relevant traffic and applies the previously configured policer. The term 'database' will assign all matched traffic to the 'server-silver' class with a

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

low loss-priority. At the same time, all traffic is subjected to the previously configured policer. Whenever traffic exceeds the rate of that policer, traffic is assigned the 'server-silver' class and it will have the loss-priority set to high.

```
root@srx2# set firewall family inet filter mf-database term database
from protocol tcp

root@srx2# set firewall family inet filter mf-database term database
from port 9876

root@srx2# set firewall family inet filter mf-database term database
then policer database-limit

root@srx2# set firewall family inet filter mf-database term database
then loss-priority low

root@srx2# set firewall family inet filter mf-database term database
then forwarding-class server-silver

root@srx2# set firewall family inet filter mf-database term database
then accept
```

We conclude the firewall filter with a term accepting all other traffic and assigning it to the 'server-bronze' forwarding class with the proper priority:

```
root@srx2# set firewall family inet filter mf-database term all-other
then loss-priority low

root@srx2# set firewall family inet filter mf-database term all-other
then forwarding-class server-bronze

root@srx2# set firewall family inet filter mf-database term all-other
then accept
```

11) Configure srx2 to use Class based forwarding so that video traffic to subnet 192.168.231.0/24 will be send out interface ge-0/0/4.12 and best-effort traffic will be sent out interface ge-0/0/4.23:

The configuration of CoS-based forwarding requires three steps.

Step 1: the configuration of a next-hop map.

```
root@srx2# set class-of-service forwarding-policy next-hop-map cbf-map
forwarding-class server-gold next-hop 172.30.0.1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set class-of-service forwarding-policy next-hop-map cbf-map
forwarding-class best-effort next-hop 172.30.0.14
```

Step 2: the configuration of a route-policy that matches interesting routes and refers to the previously configured cos-next-hop-map.

```
root@srx2# set policy-options policy-statement cbf.policy term 1 from
route-filter 192.168.231.0/24 exact
```

```
root@srx2# set policy-options policy-statement cbf.policy term 1 then
cos-next-hop-map cbf-map
```

Step 3: applying the configured policy to the forwarding table.

```
root@srx2# set routing-options forwarding-table export cbf.policy
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 3: EX class of Service

The CoS construct to configure quality of service is very similar on different devices running JUNOS. The previous chapter has already provided detailed explanations on certain CoS topics. This chapter will reiterate some of the explanations, but not in the same level of detail.

1) Configure forwarding classes and map them to output queues on all EX switches:

We use the following configuration to name our forwarding classes and map them. We apply the configuration on every ex device:

```
root@ex1# set class-of-service forwarding-classes queue 0 best-effort

root@ex1# set class-of-service forwarding-classes queue 1 server-
bronze

root@ex1# set class-of-service forwarding-classes queue 2 server-
silver

root@ex1# set class-of-service forwarding-classes queue 3 server-gold

root@ex1# set class-of-service forwarding-classes queue 5 voice

root@ex1# set class-of-service forwarding-classes queue 7 control
```

```
root@ex1# run show class-of-service forwarding-class

Verification:
Forwarding class                      ID                      Queue
Policing priority    SPU priority
  best-effort                         0                       0
normal           low
  voice                               1                       5
normal           low
  server-bronze                       2                       1
normal           low
  control                             3                       7
normal           low
  server-silver                       4                       2
normal           low
  server-gold                         5                       3
normal           low
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

2) Configure schedulers and scheduler map on all EX switches:

We named and configured our forwarding classes and we mapped them to the different output queues. To define how traffic inside a queue should be treated, we now turn to the scheduler configuration.

Schedulers are configured inside the [**class-of-service schedulers scheduler-name**] configuration stanza. Here, we can configure parameters such as the transmit-rate, buffer-size, priority and drop profiles. We already discussed the schedulers in more detail when we performed the second assignment of the previous task 'SRX Class of Service'. We will not go over them in detail again, we'll only point out that there is one parameter in use here that we did not use previously encounter. That parameter is the 'shaping-rate'.

The transmit-rate will offer a minimum rate to a scheduler during times of congestion. When there is no congestion, the scheduler can use more than it was configured for.

The shaping-rate is different. When a shaping-rate is configured, the rate at which packets are send is always throttled to the maximum that is configured.

We proceed to configure the schedulers:

```
root@ex1# set class-of-service schedulers control-sched-core shaping-rate percent 5

root@ex1# set class-of-service schedulers control-sched-core buffer-size percent 10

root@ex1# set class-of-service schedulers control-sched-core priority strict-high

root@ex1# set class-of-service schedulers control-sched-user shaping-rate percent 2

root@ex1# set class-of-service schedulers control-sched-user buffer-size percent 10

root@ex1# set class-of-service schedulers control-sched-user priority strict-high

root@ex1# set class-of-service schedulers voice-sched-core shaping-rate percent 10

root@ex1# set class-of-service schedulers voice-sched-core buffer-size percent 5

root@ex1# set class-of-service schedulers voice-sched-core buffer-size exact
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set class-of-service schedulers voice-sched-core priority
strict-high

root@ex1# set class-of-service schedulers voice-sched-user shaping-
rate percent 5

root@ex1# set class-of-service schedulers voice-sched-user buffer-size
percent 5

root@ex1# set class-of-service schedulers voice-sched-user buffer-size
exact

root@ex1# set class-of-service schedulers voice-sched-user priority
strict-high

root@ex1# set class-of-service schedulers server-gold-sched shaping-
rate percent 25

root@ex1# set class-of-service schedulers server-gold-sched buffer-
size percent 10

root@ex1# set class-of-service schedulers server-gold-sched buffer-
size exact

root@ex1# set class-of-service schedulers server-gold-sched priority
strict-high

root@ex1# set class-of-service schedulers server-silver-sched
transmit-rate percent 35

root@ex1# set class-of-service schedulers server-silver-sched buffer-
size percent 50

root@ex1# set class-of-service schedulers server-silver-sched priority
low

root@ex1# set class-of-service schedulers server-bronze-sched
transmit-rate remainder

root@ex1# set class-of-service schedulers server-bronze-sched buffer-
size remainder

root@ex1# set class-of-service schedulers server-bronze-sched priority
low

root@ex1# set class-of-service schedulers best-effort-sched transmit-
rate remainder
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set class-of-service schedulers best-effort-sched buffer-
size remainder

root@ex1# set class-of-service schedulers best-effort-sched priority
low
```

We have configured core and user schedulers. To be able to apply these schedulers later on, we configure two different scheduler-maps: 'cos-sched-map-core' and 'cos-sched-map-user'.

First, the 'cos-sched-map-core':

```
root@ex1# set class-of-service scheduler-maps cos-sched-map-core
forwarding-class control scheduler control-sched-core

root@ex1# set class-of-service scheduler-maps cos-sched-map-core
forwarding-class voice scheduler voice-sched-core

root@ex1# set class-of-service scheduler-maps cos-sched-map-core
forwarding-class server-gold scheduler server-gold-sched

root@ex1# set class-of-service scheduler-maps cos-sched-map-core
forwarding-class server-silver scheduler server-silver-sched

root@ex1# set class-of-service scheduler-maps cos-sched-map-core
forwarding-class server-bronze scheduler server-bronze-sched

root@ex1# set class-of-service scheduler-maps cos-sched-map-core
forwarding-class best-effort scheduler best-effort-sched
```

Then, the 'cos-sched-map-user':

```
root@ex1# set class-of-service scheduler-maps cos-sched-map-user
forwarding-class control scheduler control-sched-user

root@ex1# set class-of-service scheduler-maps cos-sched-map-user
forwarding-class voice scheduler voice-sched-user

root@ex1# set class-of-service scheduler-maps cos-sched-map-user
forwarding-class server-gold scheduler server-gold-sched

root@ex1# set class-of-service scheduler-maps cos-sched-map-user
forwarding-class server-silver scheduler server-silver-sched

root@ex1# set class-of-service scheduler-maps cos-sched-map-user
forwarding-class server-bronze scheduler server-bronze-sched
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set class-of-service scheduler-maps cos-sched-map-user
forwarding-class best-effort scheduler best-effort-sched
```

Verification:

```
root@ex1# run show class-of-service scheduler-map
Scheduler map: <default>, Index: 2

  Scheduler: <default-be>, Forwarding class: best-effort, Index: 21
    Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: <default-nc>, Forwarding class: control, Index: 23
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

Scheduler map: <default-chassis>, Index: 4

  Scheduler: <default-ch-scheduler>, Forwarding class: best-effort,
Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: <default-ch-scheduler>, Forwarding class: voice, Index:
25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: <default-ch-scheduler>, Forwarding class: server-bronze,
Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: <default-ch-scheduler>, Forwarding class: control, Index:
25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

Scheduler map: <default-chassis-derived>, Index: 5

  Scheduler: <default-ch-scheduler>, Forwarding class: best-effort,
Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: <default-ch-scheduler>, Forwarding class: voice, Index:
25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: <default-ch-scheduler>, Forwarding class: server-bronze,
Index: 25
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: <default-ch-scheduler>, Forwarding class: control, Index:
25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

Scheduler map: cos-sched-map-core, Index: 42113

  Scheduler: best-effort-sched, Forwarding class: best-effort, Index:
2582
    Transmit rate: remainder, Rate Limit: none, Buffer size:
remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: voice-sched-core, Forwarding class: voice, Index: 4529
    Transmit rate: remainder, Rate Limit: none, Buffer size: 5
percent,
    Buffer Limit: exact, Priority: strict-high
    Excess Priority: unspecified
    Shaping rate: 10 percent,
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: server-bronze-sched, Forwarding class: server-bronze,
Index: 13758
    Transmit rate: remainder, Rate Limit: none, Buffer size:
remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: control-sched-core, Forwarding class: control, Index:
8136
    Transmit rate: remainder, Rate Limit: none, Buffer size: 10
percent,
    Buffer Limit: none, Priority: strict-high
    Excess Priority: unspecified
    Shaping rate: 5 percent,
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: server-silver-sched, Forwarding class: server-silver,
Index: 34223
    Transmit rate: 35 percent, Rate Limit: none, Buffer size: 50
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: server-gold-sched, Forwarding class: server-gold, Index:
40991
    Transmit rate: remainder, Rate Limit: none, Buffer size: 10
percent,
    Buffer Limit: exact, Priority: strict-high
    Excess Priority: unspecified
    Shaping rate: 25 percent,
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

Scheduler map: cos-sched-map-user, Index: 54909

  Scheduler: best-effort-sched, Forwarding class: best-effort, Index:
2582
    Transmit rate: remainder, Rate Limit: none, Buffer size:
remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Scheduler: voice-sched-user, Forwarding class: voice, Index: 25421
    Transmit rate: remainder, Rate Limit: none, Buffer size: 5
percent,
    Buffer Limit: exact, Priority: strict-high
    Excess Priority: unspecified
    Shaping rate: 5 percent,
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: server-bronze-sched, Forwarding class: server-bronze,
Index: 13758
    Transmit rate: remainder, Rate Limit: none, Buffer size:
remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: control-sched-user, Forwarding class: control, Index:
27956
    Transmit rate: remainder, Rate Limit: none, Buffer size: 10
percent,
    Buffer Limit: none, Priority: strict-high
    Excess Priority: unspecified
    Shaping rate: 2 percent,
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: server-silver-sched, Forwarding class: server-silver,
Index: 34223
    Transmit rate: 35 percent, Rate Limit: none, Buffer size: 50
percent,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: server-gold-sched, Forwarding class: server-gold, Index:
40991
    Transmit rate: remainder, Rate Limit: none, Buffer size: 10
percent,
    Buffer Limit: exact, Priority: strict-high
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Excess Priority: unspecified
    Shaping rate: 25 percent,
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1   <default-drop-profile>
      High            TCP             1   <default-drop-profile>
```

3) Apply the schedulers to the physical interfaces. Make sure that the OoB interfaces are not affected:

Using wildcards can save a lot of time. So where we can, we use them. On ex1 and ex2 for example:

On ex1 and ex2:

```
root@ex1# set class-of-service interfaces ge-* scheduler-map cos-sched-map-core
```

On the other devices, we simply specify the scheduler-map for every interface:

On ex3 and ex4:

```
root@ex3# set class-of-service interfaces ge-0/0/6 scheduler-map cos-sched-map-user

root@ex3# set class-of-service interfaces ge-0/0/7 scheduler-map cos-sched-map-user

root@ex3# set class-of-service interfaces ge-0/0/8 scheduler-map cos-sched-map-user

root@ex3# set class-of-service interfaces ge-0/0/14 scheduler-map cos-sched-map-core

root@ex3# set class-of-service interfaces ge-0/0/18 scheduler-map cos-sched-map-core
```

4) Configure an IPv4 and an IPv6 Behavior Aggregate classifier. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that default DSCP classifiers are used:

We configure two BA classifiers that can be used to classify all incoming traffic based on the existing markings. We add the '**import default**' configuration statement to the classifiers to import the system default classifier mapping configuration as a template.

The IPv4 classifier:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set class-of-service classifiers dscp ba-classifier import
default

root@ex1# set class-of-service classifiers dscp ba-classifier
forwarding-class control loss-priority low code-points nc1

root@ex1# set class-of-service classifiers dscp ba-classifier
forwarding-class control loss-priority low code-points nc2

root@ex1# set class-of-service classifiers dscp ba-classifier
forwarding-class voice loss-priority low code-points ef

root@ex1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-gold loss-priority low code-points af41

root@ex1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-silver loss-priority low code-points af21

root@ex1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-silver loss-priority high code-points af23

root@ex1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-bronze loss-priority low code-points af11

root@ex1# set class-of-service classifiers dscp ba-classifier
forwarding-class server-bronze loss-priority high code-points af13

root@ex1# set class-of-service classifiers dscp ba-classifier
forwarding-class best-effort loss-priority low code-points be
```

The IPv6 classifier:

```
root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
import default

root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class control loss-priority low code-points nc1

root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class control loss-priority low code-points nc2

root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class voice loss-priority low code-points ef

root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-gold loss-priority low code-points af41
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-silver loss-priority low code-points af21

root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-silver loss-priority high code-points af23

root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-bronze loss-priority low code-points af11

root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class server-bronze loss-priority high code-points af13

root@ex1# set class-of-service classifiers dscp-ipv6 ba-classifier-v6
forwarding-class best-effort loss-priority low code-points be
```

5) Apply the classifiers to the core logical interfaces. Make sure that the OoB interfaces are not affected:

The easiest and quickest way to apply classifiers is by configuring them using a wildcard:

On ex1 and ex2:

```
root@ex1# set class-of-service interfaces ge-* unit * classifiers dscp
ba-classifier

root@ex1# set class-of-service interfaces ge-* unit * classifiers
dscp-ipv6 ba-classifier-v6
```

On ex3 and ex4:

```
root@ex3# set class-of-service interfaces ge-0/0/14 unit * classifiers
dscp ba-classifier

root@ex3# set class-of-service interfaces ge-0/0/14 unit * classifiers
dscp-ipv6 ba-classifier-v6

root@ex3# set class-of-service interfaces ge-0/0/18 unit * classifiers
dscp ba-classifier

root@ex3# set class-of-service interfaces ge-0/0/18 unit * classifiers
dscp-ipv6 ba-classifier-v6
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

6) Configure rewrite-rules to mark outgoing traffic with the correct DSCP values:

As soon as the traffic exits the device, the rewrite-rule will check the forwarding class that the packet belongs to. A lookup is done and the CoS value configured in the rewrite rule is set in the packet header.

We first configure the IPv4 rewrite-rule:

```
root@ex1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class control loss-priority low code-point nc1

root@ex1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class control loss-priority high code-point nc2

root@ex1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class voice loss-priority low code-point ef

root@ex1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-gold loss-priority low code-point af41

root@ex1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-silver loss-priority low code-point af21

root@ex1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-silver loss-priority high code-point af23

root@ex1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-bronze loss-priority low code-point af11

root@ex1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class server-bronze loss-priority high code-point af13

root@ex1# set class-of-service rewrite-rules dscp rewrite-v4
forwarding-class best-effort loss-priority high code-point be
```

We then configure the IPv6 rewrite-rule:

```
root@ex1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class control loss-priority low code-point nc1

root@ex1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class control loss-priority high code-point nc2

root@ex1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class voice loss-priority low code-point ef

root@ex1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-gold loss-priority low code-point af41
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-silver loss-priority low code-point af21

root@ex1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-silver loss-priority high code-point af23

root@ex1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-bronze loss-priority low code-point af11

root@ex1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class server-bronze loss-priority high code-point af13

root@ex1# set class-of-service rewrite-rules dscp-ipv6 rewrite-v6
forwarding-class best-effort loss-priority high code-point be
```

7) Apply the rewrite-rules to the core interfaces. Make sure that the OoB interfaces are not affected:

On ex1 and ex2:

```
root@ex1# set class-of-service interfaces ge-* unit * rewrite-rules
dscp rewrite-v4

root@ex1# set class-of-service interfaces ge-* unit * rewrite-rules
dscp-ipv6 rewrite-v6
```

On ex3 and ex4:

```
root@ex3# set class-of-service interfaces ge-0/0/14 unit * rewrite-
rules dscp rewrite-v4

root@ex3# set class-of-service interfaces ge-0/0/14 unit * rewrite-
rules dscp-ipv6 rewrite-v6

root@ex3# set class-of-service interfaces ge-0/0/18 unit * rewrite-
rules dscp rewrite-v4

root@ex3# set class-of-service interfaces ge-0/0/18 unit * rewrite-
rules dscp-ipv6 rewrite-v6
```

8) Configure ex3 and ex4 to classify incoming traffic on user VLANs using an MF classifier. Rate limit the traffic in the best-effort forwarding-class to 100Mbit/s with a burst-size of 1MB:

We start with the configuration of the policer, so that we can use that later on:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3# set firewall policer best-effort-limit if-exceeding
bandwidth-limit 100m

root@ex3# set firewall policer best-effort-limit if-exceeding burst-
size-limit 1m

root@ex3# set firewall policer best-effort-limit then discard
```

We then configure the firewall filter for the ethernet-switching family:

```
root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-1 from protocol tcp

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-1 from protocol udp

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-1 from source-port 5060

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-1 then accept

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-1 then forwarding-class voice

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-1 then loss-priority low


root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-2 from protocol tcp

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-2 from protocol udp

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-2 from destination-port 5060

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-2 then accept

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-2 then forwarding-class voice

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-control-2 then loss-priority low
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-1 from protocol udp

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-1 from source-port 5004

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-1 from source-port 5005

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-1 then accept

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-1 then forwarding-class voice

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-1 then loss-priority low


root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-2 from protocol udp

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-2 from destination-port 5004

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-2 from destination-port 5005

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-2 then accept

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-2 then forwarding-class voice

root@ex3# set firewall family ethernet-switching filter mf-user term
voice-data-2 then loss-priority low


root@ex3# set firewall family ethernet-switching filter mf-user term
database-1 from protocol tcp

root@ex3# set firewall family ethernet-switching filter mf-user term
database-1 from source-port 9876

root@ex3# set firewall family ethernet-switching filter mf-user term
database-1 then accept
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3# set firewall family ethernet-switching filter mf-user term
database-1 then forwarding-class server-silver

root@ex3# set firewall family ethernet-switching filter mf-user term
database-1 then loss-priority low


root@ex3# set firewall family ethernet-switching filter mf-user term
database-2 from protocol tcp

root@ex3# set firewall family ethernet-switching filter mf-user term
database-2 from destination-port 9876

root@ex3# set firewall family ethernet-switching filter mf-user term
database-2 then accept

root@ex3# set firewall family ethernet-switching filter mf-user term
database-2 then forwarding-class server-silver

root@ex3# set firewall family ethernet-switching filter mf-user term
database-2 then loss-priority low


root@ex3# set firewall family ethernet-switching filter mf-user term
all-other then accept

root@ex3# set firewall family ethernet-switching filter mf-user term
all-other then forwarding-class best-effort

root@ex3# set firewall family ethernet-switching filter mf-user term
all-other then loss-priority high

root@ex3# set firewall family ethernet-switching filter mf-user term
all-other then policer best-effort-limit
```

9) Apply the MF classifier on the interfaces where the user VLANs reside. Make sure that if new interfaces are added to the VLANs, the MF classifier will be applied without modifying the configuration:

We applied the previously configured firewall filter to the VLANs:

```
root@ex3# set vlans Omega filter input mf-user

root@ex3# set vlans Theta filter input mf-user

root@ex3# set vlans Theta-community filter input mf-user
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

10) Configure all voice enabled interfaces on ex1 and ex2 to advertise the forwarding-class 'voice' to the VoIP phones:

In order to have the switch advertise the forwarding class, we move to the [**ethernet-switching-options voip interface <*interface-name*> forwarding-class**] configuration stanza. Here, we can configure what is advertised:

```
root@ex1# set ethernet-switching-options voip interface ge-0/0/0.0
forwarding-class voice

root@ex1# set ethernet-switching-options voip interface ge-0/0/1.0
forwarding-class voice

root@ex1# set ethernet-switching-options voip interface ge-0/0/2.0
forwarding-class voice

root@ex1# set ethernet-switching-options voip interface ge-0/0/3.0
forwarding-class voice

root@ex1# set ethernet-switching-options voip interface ge-0/0/4.0
forwarding- class voice
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Chapter 7: A Full Day Lab Challenge 1

Task 1: Initial System Configuration

1) Load the baseline configuration:

```
root@device# load override terminal
```

Use Ctrl+D at the end, commit and reboot SRXs to change security mode to packet mode.

2) Configure a hostname on every device:

```
root@device# set system host-name srx1
```

3) Configure management interface on all of the devices:

```
root@srx1# set interfaces ge-0/0/0 unit 0 description "OoB management connection"

root@srx1# set interfaces ge-0/0/0 unit 0 family inet address 10.10.1.1/24

root@srx1# set interfaces ge-0/0/0 unit 0 family inet6 address fdb3:0309:419b:9ec4::1/64
```

4) Enable each device to accept SSHv2 and FTP. Make sure that the user **root** is not allowed to login using SSH. Allow no more than 3 simultaneous connections for SSH and FTP:

Enabling FTP and specifying a connection limit:

```
root@srx1# set system services ftp connection-limit 3
```

Enabling SSHv2, specifying a connection limit and disabling the root user from being able to login:

```
root@srx1# set system services ssh root-login deny

root@srx1# set system services ssh protocol-version v2

root@srx1# set system services ssh connection-limit 3
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

5) Make sure that the terminal session is disconnected when the console cable is disconnected:

```
root@srx1# set system ports console log-out-on-disconnect
```

6) Configure the static routes to the management network:

Two static routes need to be configured. To ensure that they will not be eligeble for redistribution, we configure both routes with the 'no-readvertise' option:

```
root@srx1# set routing-options rib inet6.0 static route
fd19:2f8d:105a:69bd::/64 next-hop fdb3:0309:419b:9ec4::fe

root@srx1# set routing-options rib inet6.0 static route
fd19:2f8d:105a:69bd::/64 no-readvertise


root@srx1# set routing-options static route 10.10.10.0/24 next-hop
10.10.1.254

root@srx1# set routing-options static route 10.10.10.0/24 no-
readvertise
```

7) Configure the S1 server as DNS server. Also configure a domain for the device and a static name mapping based on the task:

Configuring the domain of the device is done as follows:

```
root@srx1# set system domain-name jncie-lab.com
```

Static mappings are configured under the [**system static-host-maping**] stanza:

```
root@srx1# set system static-host-mapping s1.jncie-lab.com inet
10.10.10.1
```

We can configure the devices to use the S1 server as DNS server using the following configuration:

```
root@srx1# set system name-server 10.10.10.1
```

8) Set the time zone on all devices:

```
root@srx1# set system time-zone Europe/Amsterdam
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

9) Configure the S1 server to be the NTP server:

We use the '**boot-server**' configuration command to ensure synchronization at boot time. We configure the usages of an authentication-key in the [**system ntp server]** stanza to ensure that the NTP packets send by the local device are authenticated. Additionally, we configure the '**system ntp trusted-key'** to ensure that the local device will only accept authenticated packets using the configured key:

```
root@srx1# set system ntp boot-server 10.10.10.1

root@srx1# set system ntp authentication-key 1 type md5

root@srx1# set system ntp authentication-key 1 value workbook

root@srx1# set system ntp server 10.10.10.1 key 1

root@srx1# set system ntp server 10.10.10.1 version 4

root@srx1# set system ntp trusted-key 1
```

10) Configure archival of the configuration to an FTP server every 60 minutes:

```
root@srx1# set system archival configuration transfer-interval 60

root@srx1# set system archival configuration archive-sites
"ftp://lab@10.10.10.1" password lab123
```

11) Configure TACACS+ and allow local password authentication when the TACACS+ authentication failes:

```
root@srx1# set system authentication-order tacplus

root@srx1# set system authentication-order password

root@srx1# set system tacplus-server 10.10.10.1 secret workbook

root@srx1# set system tacplus-server 10.10.10.1 timeout 1
```

12) Configure user with super user privilege:

```
root@srx1# set system login user lab uid 2000

root@srx1# set system login user lab class super-user
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set system login user lab authentication plain-text-
password
```

13) Configure additional users and classes:

First, we configure a class for the 'noc1' user:

```
root@srx1# set system login class noc-daily allowed-days monday

root@srx1# set system login class noc-daily allowed-days tuesday

root@srx1# set system login class noc-daily allowed-days wednesday

root@srx1# set system login class noc-daily allowed-days thursday

root@srx1# set system login class noc-daily allowed-days friday

root@srx1# set system login class noc-daily access-start "08:30:00
+0200"

root@srx1# set system login class noc-daily access-end "18:00:00
+0200"

root@srx1# set system login class noc-daily permissions view

root@srx1# set system login class noc-daily permissions view-
configuration

root@srx1# set system login class noc-daily allow-commands "configure
private"
```

Then, we configure the 'noc1' user itself and reference the class:

```
root@srx1# set system login user noc1 uid 2001

root@srx1# set system login user noc1 class noc-daily

root@srx1# set system login user noc1 authentication plain-text-
password
```

Then we configure a class for 'noc2'. This class should not be allowed to delete file, but it should have all other permissions:

```
root@srx1# set system login class noc-duty permissions all

root@srx1# set system login class noc-duty deny-commands "file delete
*"
```

Lastly, we configure the 'noc2' user:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set system login user noc2 uid 2002

root@srx1# set system login user noc2 class noc-duty

root@srx1# set system login user noc2 authentication plain-text-
password
```

14) Make sure that the newly configured users from the previous task logout after they have been idle for 5 minutes:

We can only configure the **'idle-timeout'** on a single location. Therfore, we first configure a group and then apply that group to the relevant users:

```
root@srx1# set groups login-timeout system login class <*> idle-
timeout 5

root@srx1# set system login user noc1 apply-groups login-timeout

root@srx1# set system login user noc2 apply-groups login-timeout
```

15) Configure the number of login attempts before the user is disconnected to 5:

Retry-limits for all users can be configured under the [**system login retry-options**] stanza. Additionally, they can be configured per user under the [**system login user** <***username***>] stanza. Since both of these options are not allowed, we resort to a group configuration.

First we define a group:

```
root@srx1# set groups retry-limit system login retry-options tries-
before-disconnect 5
```

Then we apply the group:

```
root@srx1# set system login apply-groups retry-options
```

Technically, we did not configure any statements under the [**system login retry-options**] hierarchy. To verify the **apply-group** configuration in configuration mode, issue the ' **show system login | display inheritance**' command.

16) Configure a banner on every device:

```
root@srx1# set system login announcement "\nJNCIE-ENT Lab Rack
1\n\nMaster Junos CLI to perfection\n\n"
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

To insert a newline in the banner, we include the '**\n**' in our banner. This inserts a new line on the place when it is put inside the banner configuration.

17) All changes in the configuration should be logged to noc1 user console:

```
root@srx1# set system syslog user noc1 change-log any
```

18) Log all authentication and authorization messages to host 10.10.10.1. Make sure the logs have a prefix name "jncie-ent":

```
root@srx1# set system syslog host 10.10.10.1 authorization any
root@srx1# set system syslog host 10.10.10.1 log-prefix jncie-ent
```

19) Track all executed commands in a file called "interactive-commands":

```
root@srx1# set system syslog file interactive-commands interactive-
commands any
```

20) Configure the default log message:

First, we configure logging for **any** facility up to the “notice” severity level to be logged to the **messages** file:

```
root@srx1# set system syslog file messages any notice
```

Then, we specifically configure the **authorization** facility to log up to the "info" severity:

```
root@srx1# set system syslog file messages authorization info
```

21) Ensure that a maximum of 5 log files, 10M each, are stored :

```
root@srx1# set system syslog archive size 10m
root@srx1# set system syslog archive files 5
```

22) Configure SNMPv3 with the given parameters:

ET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set snmp v3 usm local-engine user lab authentication-md5
authentication-key workbook

root@srx1# set snmp v3 usm local-engine user lab privacy-aes128
privacy-key workbook

root@srx1# set snmp v3 vacm security-to-group security-model usm
security-name jncie group jncie-group

root@srx1# set snmp v3 vacm access group jncie-group default-context-
prefix security-model usm security-level privacy read-view default

root@srx1# set snmp engine-id use-mac-address

root@srx1# set snmp view default oid .1 include
```

23) Configure SNMPv3 notification based on the task:

```
root@srx1# set snmp v3 target-address s1 address 10.10.10.1

root@srx1# set snmp v3 target-address s1 target-parameters tp1

root@srx1# set snmp v3 target-parameters tp1 parameters message-
processing-model v3

root@srx1# set snmp v3 target-parameters tp1 parameters security-model
usm

root@srx1# set snmp v3 target-parameters tp1 parameters security-level
privacy

root@srx1# set snmp v3 target-parameters tp1 parameters security-name
jncie

root@srx1# set snmp v3 target-parameters tp1 notify-filter nf1

root@srx1# set snmp v3 notify n1 type trap

root@srx1# set snmp v3 notify n1 tag lab

root@srx1# set snmp v3 notify-filter nf1 oid 1.3.6.1.2.1.1.5 include
```

24) Configure an SNMP v1/v2c community named *workbook*. Assign the security name *jncie* to the community.

```
root@srx1# set snmp v3 snmp-community workbook community-name workbook
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set snmp v3 snmp-community workbook security-name jncie

root@srx1# set snmp v3 snmp-community workbook tag lab
```

25-30) Here, we combine the explanation of tasks 25 to 30:

For these taks we need to configure an IPv4 firewall filter that protects the control plane. We configure a term for every individual protocol. We also need to make sure that newly configured BGP neighbors are added to the firewall filter automatically.

Additionally, we rate-limit ICMP and traceroute traffic for both IPv4 as well as IPv6 on the srx. On the ex switches, we simply permit both of these protocols.

Lastly, we conclude the firewall filters with a term that discards the traffic and increments a named counter whenever a packet hits the term.

The policer for the ICMP and traceroute traffic on the srx:

```
root@srx1# set firewall policer rate-limit if-exceeding bandwidth-
limit 50k

root@srx1# set firewall policer rate-limit if-exceeding burst-size-
limit 12k

root@srx1# set firewall policer rate-limit then discard
```

To be able to automatically add BGP sessions to the firewall filter:

```
root@srx1# set policy-options prefix-list bgp-peers apply-path
"protocols bgp group <*> neighbor <*>"
```

The IPv4 firewall filter:

```
root@srx1# set firewall family inet filter protect-re-ipv4 term bfd
from protocol udp

root@srx1# set firewall family inet filter protect-re-ipv4 term bfd
from destination-port 3784

root@srx1# set firewall family inet filter protect-re-ipv4 term bfd
from destination-port 4784
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set firewall family inet filter protect-re-ipv4 term bfd
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term vrrp
from protocol vrrp

root@srx1# set firewall family inet filter protect-re-ipv4 term vrrp
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term ospf
from protocol ospf

root@srx1# set firewall family inet filter protect-re-ipv4 term ospf
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term bgp-
source from source-prefix-list bgp-peers

root@srx1# set firewall family inet filter protect-re-ipv4 term bgp-
source from protocol tcp

root@srx1# set firewall family inet filter protect-re-ipv4 term bgp-
source from source-port bgp

root@srx1# set firewall family inet filter protect-re-ipv4 term bgp-
source then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term bgp-
destination from protocol tcp

root@srx1# set firewall family inet filter protect-re-ipv4 term bgp-
destination from destination-port bgp

root@srx1# set firewall family inet filter protect-re-ipv4 term bgp-
destination then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term pim
from protocol pim

root@srx1# set firewall family inet filter protect-re-ipv4 term pim
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term igmp
from protocol igmp

root@srx1# set firewall family inet filter protect-re-ipv4 term igmp
then accept
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set firewall family inet filter protect-re-ipv4 term ntp
from source-address 10.10.10.0/24

root@srx1# set firewall family inet filter protect-re-ipv4 term ntp
from protocol udp

root@srx1# set firewall family inet filter protect-re-ipv4 term ntp
from source-port ntp

root@srx1# set firewall family inet filter protect-re-ipv4 term ntp
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term ntp-
local from protocol udp

root@srx1# set firewall family inet filter protect-re-ipv4 term ntp-
local from destination-port ntp

root@srx1# set firewall family inet filter protect-re-ipv4 term ntp-
local then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term snmp
from source-address 10.10.10.0/24

root@srx1# set firewall family inet filter protect-re-ipv4 term snmp
from protocol udp

root@srx1# set firewall family inet filter protect-re-ipv4 term snmp
from destination-port snmp

root@srx1# set firewall family inet filter protect-re-ipv4 term snmp
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term radius
from source-address 10.10.10.0/24

root@srx1# set firewall family inet filter protect-re-ipv4 term radius
from protocol udp

root@srx1# set firewall family inet filter protect-re-ipv4 term radius
from source-port radius

root@srx1# set firewall family inet filter protect-re-ipv4 term radius
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term tacacs
from source-address 10.10.10.0/24
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set firewall family inet filter protect-re-ipv4 term tacacs
from protocol tcp

root@srx1# set firewall family inet filter protect-re-ipv4 term tacacs
from source-port tacacs

root@srx1# set firewall family inet filter protect-re-ipv4 term tacacs
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term dns
from source-address 10.10.10.0/24

root@srx1# set firewall family inet filter protect-re-ipv4 term dns
from protocol udp

root@srx1# set firewall family inet filter protect-re-ipv4 term dns
from source-port domain

root@srx1# set firewall family inet filter protect-re-ipv4 term dns
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term ssh
from source-address 10.10.1.0/24

root@srx1# set firewall family inet filter protect-re-ipv4 term ssh
from protocol tcp

root@srx1# set firewall family inet filter protect-re-ipv4 term ssh
from destination-port ssh

root@srx1# set firewall family inet filter protect-re-ipv4 term ssh
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term ssh-
local from source-address 10.10.1.0/24

root@srx1# set firewall family inet filter protect-re-ipv4 term ssh-
local from protocol tcp

root@srx1# set firewall family inet filter protect-re-ipv4 term ssh-
local from source-port ssh

root@srx1# set firewall family inet filter protect-re-ipv4 term ssh-
local then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term ftp
from source-address 10.10.1.0/24
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set firewall family inet filter protect-re-ipv4 term ftp
from protocol tcp

root@srx1# set firewall family inet filter protect-re-ipv4 term ftp
from destination-port ftp

root@srx1# set firewall family inet filter protect-re-ipv4 term ftp
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term ftp-
data from source-address 10.10.1.0/24

root@srx1# set firewall family inet filter protect-re-ipv4 term ftp-
data from protocol tcp

root@srx1# set firewall family inet filter protect-re-ipv4 term ftp-
data from destination-port ftp-data

root@srx1# set firewall family inet filter protect-re-ipv4 term ftp-
data then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term icmp
from protocol icmp

root@srx1# set firewall family inet filter protect-re-ipv4 term icmp
then policer rate-limit

root@srx1# set firewall family inet filter protect-re-ipv4 term icmp
then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term
traceroute from protocol udp

root@srx1# set firewall family inet filter protect-re-ipv4 term
traceroute from destination-port 33434-33534

root@srx1# set firewall family inet filter protect-re-ipv4 term
traceroute then policer rate-limit

root@srx1# set firewall family inet filter protect-re-ipv4 term
traceroute then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term last
then count dropped-ipv4

root@srx1# set firewall family inet filter protect-re-ipv4 term last
then discard
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

31) Ensure that the firewall filter is used for RE protection by applying it to the loopback0 interface:

```
root@srx1# set interfaces lo0 unit 0 family inet filter input protect-re-ipv4
```

32-35) Here, we combine the explanation of tasks 32 to 35:

The objective here is to configure an IPv6 firewall filter that can be use to protect the control plane from IPv6 traffic.

We can use the same policer which we used to limit icmp and traceroute traffics for IPv4 on the srx. On the ex switches, we only need to accept traceroute and icmp.

IPv6 firewall filter:

```
root@srx1# set firewall family inet6 filter protect-re-ipv6 term ospf from next-header ospf

root@srx1# set firewall family inet6 filter protect-re-ipv6 term ospf then accept

root@srx1# set firewall family inet6 filter protect-re-ipv6 term bgp-source from next-header tcp

root@srx1# set firewall family inet6 filter protect-re-ipv6 term bgp-source from source-port bgp

root@srx1# set firewall family inet6 filter protect-re-ipv6 term bgp-source then accept

root@srx1# set firewall family inet6 filter protect-re-ipv6 term bgp-destination from next-header tcp

root@srx1# set firewall family inet6 filter protect-re-ipv6 term bgp-destination from destination-port bgp

root@srx1# set firewall family inet6 filter protect-re-ipv6 term bgp-destination then accept

root@srx1# set firewall family inet6 filter protect-re-ipv6 term vrrp from next-header vrrp

root@srx1# set firewall family inet6 filter protect-re-ipv6 term vrrp then accept
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set firewall family inet6 filter protect-re-ipv6 term pim
from next-header pim

root@srx1# set firewall family inet6 filter protect-re-ipv6 term pim
then accept

root@srx1# set firewall family inet6 filter protect-re-ipv6 term icmp
from next-header icmp6

root@srx1# set firewall family inet6 filter protect-re-ipv6 term icmp
then policer rate-limit

root@srx1# set firewall family inet6 filter protect-re-ipv6 term icmp
then accept

root@srx1# set firewall family inet6 filter protect-re-ipv6 term
traceroute from next-header udp

root@srx1# set firewall family inet6 filter protect-re-ipv6 term
traceroute from destination-port 33434-33534

root@srx1# set firewall family inet6 filter protect-re-ipv6 term
traceroute then policer rate-limit

root@srx1# set firewall family inet6 filter protect-re-ipv6 term
traceroute then accept

root@srx1# set firewall family inet6 filter protect-re-ipv6 term last
then count dropped-ipv6

root@srx1# set firewall family inet6 filter protect-re-ipv6 term last
then discard
```

The ICMP and traceroute terms for the EX switches:

```
root@ex1# set firewall family inet6 filter protect-re-ipv6 term icmp
from next-header icmp6

root@ex1# set firewall family inet6 filter protect-re-ipv6 term icmp
then accept

root@ex1# set firewall family inet6 filter protect-re-ipv6 term
traceroute from next-header udp

root@ex1# set firewall family inet6 filter protect-re-ipv6 term
traceroute from destination-port 33434-33534
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set firewall family inet6 filter protect-re-ipv6 term
traceroute then accept
```

36) Apply the firewall filter to the loopback 0, ensuring it is used for RE protection:

```
root@srx1# set interfaces lo0 unit 0 family inet6 filter input
protect-re-ipv6
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 2: L2 switching and security

1) Create a virtual chassis using ex3 and ex4. The ex4 switch should act as a line card and a new switch should not be able to automatically join this virtual chassis:

We specify the virtual-chassis mode to be preprovisioned:

```
root@ex3# set virtual-chassis preprovisioned
```

We add the member switches to the configuration and specify ex4 as the line-card:

```
root@ex3# set virtual-chassis member 0 role routing-engine
root@ex3# set virtual-chassis member 0 serial-number BM0210045284
root@ex3# set virtual-chassis member 1 role line-card
root@ex3# set virtual-chassis member 1 serial-number BM0210359941
```

We then enable the necessary ports as Virtual Chassis Ports:

```
root@ex3# request virtual-chassis vc-port set interface vcp-0
root@ex3# request virtual-chassis vc-port set interface vcp-1
root@ex4# request virtual-chassis vc-port set interface vcp-0
root@ex3# request virtual-chassis vc-port set interface vcp-1 member 1
```

Lastly, we change the hostname of the VC:

```
root@ex3# set system host-name ex3_ex4
```

```
root@ex3_ex4# run show virtual-chassis

Preprovisioned Virtual Chassis
Virtual Chassis ID: 3afd.deb3.ed7d
Virtual Chassis Mode: Enabled
                                        Mstr           Mixed
Neighbor List
Member ID  Status   Serial No    Model     prio  Role      Mode ID
Interface
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
0 (FPC 0)   Prsnt    BM0210045284 ex4200-24t 129  Master*      N  1
vcp-0
1 (FPC 1)   Prsnt    BM0210359941 ex4200-24t   0  Linecard     N  0
vcp-1
```

2) Disable split-detection:

```
root@ex3_ex4# set virtual-chassis no-split-detection
```

3) Configure the vme interface for the VC:

```
root@ex3_ex4# set interfaces vme unit 0 description "OoB management
connection"

root@ex3_ex4# set interfaces vme unit 0 family inet address
10.10.1.13/24

root@ex3_ex4# set interfaces vme unit 0 family inet6 address
fdb3:0309:419b:9ec4::13/64
```

4) Build L2 network and link aggregation based on the task, note that we combine task 6 with this task:

We need to configure link aggregation between two sets of devices. We configure link aggregation between the ex1 and srx1 devices as well as ex2 and srx2 devices. Under [**chassis aggregated-devices ethernet device-count**], we specify the number of aggregated Ethernet devices to be 2. This way, we can configure the interface as AE1.

srx1:

```
root@srx1# set chassis aggregated-devices ethernet device-count 2

root@srx1# set interfaces ge-0/0/6 description "Physical Link to EX1
No.1"

root@srx1# set interfaces ge-0/0/6 gigether-options 802.3ad ae1

root@srx1# set interfaces ge-0/0/9 description "Physical Link to EX1
No.2"

root@srx1# set interfaces ge-0/0/9 gigether-options 802.3ad ae1

root@srx1# set interfaces ae1 aggregated-ether-options lacp active

root@srx1# set interfaces ae1 unit 0 description "L2 Link to EX1"
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# run show lacp interfaces ae1
Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout
Activity
      ge-0/0/6       Actor    No    No   Yes  Yes  Yes   Yes     Fast
Passive
      ge-0/0/6     Partner    No    No   Yes  Yes  Yes   Yes     Fast
Active
      ge-0/0/9       Actor    No    No   Yes  Yes  Yes   Yes     Fast
Passive
      ge-0/0/9     Partner    No    No   Yes  Yes  Yes   Yes     Fast
Active
    LACP protocol:        Receive State  Transmit State          Mux
State
      ge-0/0/6                  Current   Fast periodic Collecting
distributing
      ge-0/0/9                  Current   Fast periodic Collecting
distributing

root@ex1# run show interfaces ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 579
  Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error:
None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth
needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Current address: 3e:94:d5:46:73:04, Hardware address:
3e:94:d5:46:73:04
  Last flapped   : 2014-12-30 10:26:15 CET (00:40:46 ago)
  Input rate     : 4648 bps (5 pps)
  Output rate    : 2048 bps (2 pps)
  Logical interface ae1.0 (Index 86) (SNMP ifIndex 583)
    Description: L2 Link to SRX1
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          1354          0        199070            0
        Output:            56          0         10584            0
    Protocol eth-switch
      Flags: Trunk-Mode
```

We configure the VLANs on all devices:

```
root@srx1# set vlans Acc-servers vlan-id 702
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set vlans Accounting vlan-id 700

root@srx1# set vlans Engineering vlan-id 500

root@srx1# set vlans Transported vlan-id 800

root@srx1# set vlans User vlan-id 600

root@srx1# set vlans default vlan-id 1
```

```
root@srx1# run show vlans brief
Name            Tag
Accounting      700
Engineering     500
Transported     800
User            600
Voice           601
default         1
```

Then we apply the individual interface configuration to every device.

ex1:

```
root@ex1# set interfaces ge-0/0/7 unit 0 description "Connection to
H1"

root@ex1# set interfaces ge-0/0/7 unit 0 family ethernet-switching
port-mode access

root@ex1# set interfaces ge-0/0/10 unit 0 description "L2 Link to
SRX2"

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
port-mode trunk

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 1

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 500

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 600

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 601
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 700

root@ex1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 800

root@ex1# set interfaces ge-0/0/12 unit 0 description "L2 Link to EX2
No.1"

root@ex1# set interfaces ge-0/0/12 unit 0 family ethernet-switching
port-mode trunk

root@ex1# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 1

root@ex1# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 500

root@ex1# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 600

root@ex1# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 601

root@ex1# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 700

root@ex1# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 800

root@ex1# set interfaces ge-0/0/13 description "L2 Link to EX2 No.2"

root@ex1# set interfaces ge-0/0/13 unit 0 family ethernet-switching
port-mode trunk

root@ex1# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 1

root@ex1# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 500

root@ex1# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 600

root@ex1# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 601

root@ex1# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 700
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 800

root@ex1# set interfaces ge-0/0/16 unit 0 description "L2 Link to
SRX3"

root@ex1# set interfaces ge-0/0/16 unit 0 family ethernet-switching
port-mode trunk

root@ex1# set interfaces ge-0/0/16 unit 0 family ethernet-switching
vlan members 1

root@ex1# set interfaces ge-0/0/16 unit 0 family ethernet-switching
vlan members 500

root@ex1# set interfaces ge-0/0/18 unit 0 description "L2 Link to
EX3_EX4"

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
port-mode trunk

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members 1

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members 500

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members 600

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members 601

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members 700

root@ex1# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members 800

root@ex1# set interfaces ae1 unit 0 family ethernet-switching port-
mode trunk

root@ex1# set interfaces ae1 unit 0 family ethernet-switching vlan
members Accounting

root@ex1# set interfaces ae1 unit 0 family ethernet-switching vlan
members Engineering
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set interfaces ae1 unit 0 family ethernet-switching vlan
members Transported

root@ex1# set interfaces ae1 unit 0 family ethernet-switching vlan
members User

root@ex1# set interfaces ae1 unit 0 family ethernet-switching vlan
members Voice
```

ex2:

```
root@ex2# set interfaces ge-0/0/7 unit 0 description "Connection to
H3"

root@ex2# set interfaces ge-0/0/7 unit 0 family ethernet-switching

root@ex2# set interfaces ge-0/0/9 description "Physical Link to SRX2
No.2"

root@ex2# set interfaces ge-0/0/10 unit 0 description "L2 Link to
SRX1"

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
port-mode trunk

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 1

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 500

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 600

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 601

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 700

root@ex2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members 800

root@ex2# set interfaces ge-0/0/12 unit 0 description "L2 Link to EX1
No.1"

root@ex2# set interfaces ge-0/0/12 unit 0 family ethernet-switching
port-mode trunk
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 1

root@ex2# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 500

root@ex2# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 600

root@ex2# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 601

root@ex2# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 700

root@ex2# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 800

root@ex2# set interfaces ge-0/0/13 unit 0 description "L2 Link to EX1
No.2"

root@ex2# set interfaces ge-0/0/13 unit 0 family ethernet-switching
port-mode trunk

root@ex2# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 1

root@ex2# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 500

root@ex2# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 600

root@ex2# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 601

root@ex2# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 700

root@ex2# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 800

root@ex2# set interfaces ge-0/0/17 unit 0 description "L2 Link to
SRX4"

root@ex2# set interfaces ge-0/0/17 unit 0 family ethernet-switching
port-mode trunk
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set interfaces ge-0/0/17 unit 0 family ethernet-switching
vlan members 1

root@ex2# set interfaces ge-0/0/17 unit 0 family ethernet-switching
vlan members 500

root@ex2# set interfaces ge-0/0/19 unit 0 description "L2 Link to
EX3_EX4"

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
port-mode trunk

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
vlan members 800

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
vlan members 1

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
vlan members 500

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
vlan members 600

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
vlan members 601

root@ex2# set interfaces ge-0/0/19 unit 0 family ethernet-switching
vlan members 700

root@ex2# set interfaces ae0 unit 0 family ethernet-switching port-
mode trunk

root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members Accounting

root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members Engineering

root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members Transported

root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members User

root@ex2# set interfaces ae0 unit 0 family ethernet-switching vlan
members Voice
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

ex3_ex4:

```
root@ex3_ex4# set interfaces ge-0/0/6 description "H5 connection"

root@ex3_ex4# set interfaces ge-0/0/6 unit 0 family ethernet-switching

root@ex3_ex4# set interfaces ge-0/0/7 description "H6 connection"

root@ex3_ex4# set interfaces ge-0/0/7 unit 0 family ethernet-switching

root@ex3_ex4# set interfaces ge-0/0/8 description "H7 connection"

root@ex3_ex4# set interfaces ge-0/0/8 unit 0 family ethernet-switching

root@ex3_ex4# set interfaces ge-0/0/9 unit 0 description "L2 Link to
SRX3"

root@ex3_ex4# set interfaces ge-0/0/9 unit 0 family ethernet-switching
port-mode trunk

root@ex3_ex4# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 1

root@ex3_ex4# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 500

root@ex3_ex4# set interfaces ge-0/0/10 unit 0 family ethernet-
switching

root@ex3_ex4# set interfaces ge-0/0/18 unit 0 description "L2 Link to
EX1"

root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching port-mode trunk

root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching vlan members 800

root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching vlan members 1

root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching vlan members 500

root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching vlan members 600

root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching vlan members 601
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set interfaces ge-0/0/18 unit 0 family ethernet-
switching vlan members 700

root@ex3_ex4# set interfaces ge-0/0/20 description "SPAN port"

root@ex3_ex4# set interfaces ge-0/0/20 unit 0 family ethernet-
switching

root@ex3_ex4# set interfaces ge-1/0/6 description "H10 connection"

root@ex3_ex4# set interfaces ge-1/0/6 unit 0 family ethernet-switching

root@ex3_ex4# set interfaces ge-1/0/7 description "H9 connection"

root@ex3_ex4# set interfaces ge-1/0/7 unit 0 family ethernet-switching

root@ex3_ex4# set interfaces ge-1/0/8 description "H8 connection"

root@ex3_ex4# set interfaces ge-1/0/8 unit 0 family ethernet-switching

root@ex3_ex4# set interfaces ge-1/0/9 unit 0 description "L2 Link to
SRX4"

root@ex3_ex4# set interfaces ge-1/0/9 unit 0 family ethernet-switching
port-mode trunk

root@ex3_ex4# set interfaces ge-1/0/9 unit 0 family ethernet-switching
vlan members 1

root@ex3_ex4# set interfaces ge-1/0/9 unit 0 family ethernet-switching
vlan members 500

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 description "L2 Link to
EX2"

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching port-mode trunk

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching vlan members 800

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching vlan members 1

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching vlan members 500

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching vlan members 600
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching vlan members 601

root@ex3_ex4# set interfaces ge-1/0/19 unit 0 family ethernet-
switching vlan members 700
```

srx1:

```
root@srx1# set interfaces ge-0/0/1 unit 0 description "L2 Link to SRX2
No.1"

root@srx1# set interfaces ge-0/0/1 unit 0 family ethernet-switching
port-mode trunk

root@srx1# set interfaces ge-0/0/1 unit 0 family ethernet-switching
vlan members all

root@srx1# set interfaces ge-0/0/1 unit 0 family ethernet-switching
native-vlan-id 1

root@srx1# set interfaces ge-0/0/2 unit 0 description "L2 Link to SRX2
No.2"

root@srx1# set interfaces ge-0/0/2 unit 0 family ethernet-switching
port-mode trunk

root@srx1# set interfaces ge-0/0/2 unit 0 family ethernet-switching
vlan members all

root@srx1# set interfaces ge-0/0/10 unit 0 description "L2 Link to
EX2"

root@srx1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
port-mode trunk

root@srx1# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members all

root@srx1# set interfaces ae1 unit 0 family ethernet-switching port-
mode trunk

root@srx1# set interfaces ae1 unit 0 family ethernet-switching vlan
members all
```

srx2:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set interfaces ge-0/0/1 unit 0 description "L2 Link to SRX1
No.1"

root@srx2# set interfaces ge-0/0/1 unit 0 family ethernet-switching
port-mode trunk

root@srx2# set interfaces ge-0/0/1 unit 0 family ethernet-switching
vlan members all

root@srx2# set interfaces ge-0/0/2 unit 0 description "L2 Link to SRX1
No.2"

root@srx2# set interfaces ge-0/0/2 unit 0 family ethernet-switching
port-mode trunk

root@srx2# set interfaces ge-0/0/2 unit 0 family ethernet-switching
vlan members all

root@srx2# set interfaces ge-0/0/10 unit 0 description "L2 Link to
EX1"

root@srx2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
port-mode trunk

root@srx2# set interfaces ge-0/0/10 unit 0 family ethernet-switching
vlan members all

root@srx2# set interfaces ae0 unit 0 description "L2 Link to EX2"

root@srx2# set interfaces ae0 unit 0 family ethernet-switching port-
mode trunk

root@srx2# set interfaces ae0 unit 0 family ethernet-switching vlan
members all
```

srx3:

```
root@srx3# set interfaces ge-0/0/9 unit 0 description "L2 Link to
EX3_EX4"

root@srx3# set interfaces ge-0/0/9 unit 0 family ethernet-switching
port-mode trunk

root@srx3# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 1

root@srx3# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 500
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set interfaces ge-0/0/12 unit 0 description "L2 Link to
EX1"

root@srx3# set interfaces ge-0/0/12 unit 0 family ethernet-switching
port-mode trunk

root@srx3# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 1

root@srx3# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members 500

root@srx3# set interfaces ge-0/0/14 unit 0 description "Connection to
H2"

root@srx3# set interfaces ge-0/0/14 unit 0 family ethernet-switching
port-mode access
```

srx4:

```
root@srx4# set interfaces ge-0/0/9 unit 0 description "L2 Link to
EX3_EX4"

root@srx4# set interfaces ge-0/0/9 unit 0 family ethernet-switching
port-mode trunk

root@srx4# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 1

root@srx4# set interfaces ge-0/0/9 unit 0 family ethernet-switching
vlan members 500

root@srx4# set interfaces ge-0/0/13 unit 0 description "L2 Link to
EX2"

root@srx4# set interfaces ge-0/0/13 unit 0 family ethernet-switching
port-mode trunk

root@srx4# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 1

root@srx4# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members 500

root@srx4# set interfaces ge-0/0/14 unit 0 description "Connection to
H4"

root@srx4# set interfaces ge-0/0/14 unit 0 family ethernet-switching
port-mode access
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

5) Decrease storm control on ex1 and ex2 to 30% on all interfaces, enable it for multicast and disable it for unknown unicast traffic:

By default, storm control is enabled for 80 percent of the available bandwidth and it is active for broadcast and unkown unicast traffic streams. When the combined ingress traffic streams reach 80 percent of the available bandwidth on an interface, packets for those streams are dropped.

To complete this task, we reduce the bandwidth on which storm-control is triggerd:

```
root@ex1# set ethernet-switching-options storm-control interface all
bandwidth 30000
```

Then, we disable storm-control for unknown-unicast:

```
root@ex1# set ethernet-switching-options storm-control interface all
no-unknown-unicast
```

Finally, we enable storm control for multicast:

```
root@ex1# set ethernet-switching-options storm-control interface all
multicast
```

6) This task was combined with task 4.

7) Configure H7 as an isolated host in vlan 700:

We configure the VLAN with an isolation-id and the host-facing itnerface:

```
root@ex3_ex4# set vlans Accounting isolation-id 701
root@ex3_ex4# set vlans Accounting interface ge-0/0/8.0
```

We also configure the uplinks of the switch to be a private VLAN trunk port for this VLAN. We can do this using the **pvlan-trunk** configuration command:

```
root@ex3_ex4# set vlans Accounting interface ge-0/0/18.0 pvlan-trunk
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set vlans Accounting interface ge-1/0/19.0 pvlan-trunk
```

8) Configure vlan 702 as private vlan of vlan 700 for H6 and H8 on ex3_ex4. H1 on ex1 should be part of this community:

On ex3_ex4:

```
root@ex3_ex4# set vlans Acc-servers vlan-id 702
root@ex3_ex4# set vlans Acc-servers interface ge-0/0/7.0
root@ex3_ex4# set vlans Acc-servers interface ge-1/0/8.0
root@ex3_ex4# set vlans Acc-servers primary-vlan Accounting
```

On ex1:

```
root@ex1# set vlans Acc-servers vlan-id 702
root@ex1# set vlans Acc-servers interface ge-0/0/7.0
root@ex1# set vlans Acc-servers primary-vlan Accounting

root@ex1# set vlans Accounting interface ge-0/0/12.0 pvlan-trunk
root@ex1# set vlans Accounting interface ge-0/0/13.0 pvlan-trunk
root@ex1# set vlans Accounting interface ge-0/0/18.0 pvlan-trunk
root@ex1# set vlans Accounting isolation-id 701
```

9) Configure native vlan 1 on all of the ports that operate as a trunk port:

On all trunk links of all the devices:

```
root@srx1# set interfaces ae1 unit 0 family ethernet-switching native-vlan-id 1
```

```
root@srx1# run show vlans 1 extensive
VLAN: default, Created at: Tue Dec 30 09:51:30 2014
802.1Q Tag: 1, Internal index: 2, Admin State: Enabled, Origin: Static
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 4 (Active = 4), Untagged  4 (Active = 4)
      ae1.0*, tagged, trunk
      ge-0/0/1.0*, tagged, trunk
      ge-0/0/2.0*, tagged, trunk
      ge-0/0/10.0*, tagged, trunk
      ae1.0*, untagged, trunk
      ge-0/0/1.0*, untagged, trunk
      ge-0/0/2.0*, untagged, trunk
      ge-0/0/10.0*, untagged, trunk
```

10) H5 is preconfigured with static IP address 192.168.0.1, configure it in VLAN800 (Transported). H5 should be able to communitcate with H3. By mistake, the H3 interface is configured to send and receive tagged frames in VLAN10. You are not allowed to use dot1q-tunneling and change the interface configuration:

In this task, we use **mapping** and **swap** configuration commands. We configure the srx1 to swap, or translate, VLAN tag 10 with VLAN tag 800:

```
root@srx1# set vlans Transported interface ge-0/0/7.0 mapping 10 swap
```

```
root@srx3# run ping routing-instance H5 192.168.0.2
PING 192.168.0.2 (192.168.0.2): 56 data bytes
64 bytes from 192.168.0.2: icmp_seq=0 ttl=64 time=2.748 ms
64 bytes from 192.168.0.2: icmp_seq=1 ttl=64 time=1.400 ms
64 bytes from 192.168.0.2: icmp_seq=2 ttl=64 time=1.264 ms
64 bytes from 192.168.0.2: icmp_seq=3 ttl=64 time=1.093 ms
```

11) On ex1 and ex2, enable interfaces ge-0/0/0 to ge-0/0/4 for VoiP. Make sure voice traffic is tagged in VLAN 601. Data traffic should use VLAN 600 and it should be untagged. Also enable the interfaces connected to the IP phones for POE.

We configure the interfaces that will be used for VoIP as an **interface-range**. The interface range is configured with the '**port-mode trunk**' command.

On ex1 and ex2:

```
root@ex1# set interfaces interface-range voice member-range ge-0/0/0
to ge-0/0/4
```

```
root@ex1# set interfaces interface-range voice unit 0 family ethernet-
switching port-mode trunk
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

We then add the two VLANs to the interface range, VLANs 600 and 601. In order to make sure packets in VLAN 600 are send without a tag, we use the '**native-vlan-id 600**' configuration command.

On ex1 and ex2:

```
root@ex1# set interfaces interface-range voice unit 0 family ethernet-switching vlan members 601

root@ex1# set interfaces interface-range voice unit 0 family ethernet-switching native-vlan-id 600
```

Lastly, we configure power over Ethernet for the interfaces connected to the IP phones:

On ex1 and ex2:

```
root@ex1# set poe interface ge-0/0/0

root@ex1# set poe interface ge-0/0/1

root@ex1# set poe interface ge-0/0/2

root@ex1# set poe interface ge-0/0/3

root@ex1# set poe interface ge-0/0/4
```

12) MSTP configuration on srx1, srx2, srx4 and ex2. Please review the theory section of chapter 2 for MSTP configuration to complete this task:

srx1:

```
root@srx1# set protocols mstp configuration-name east

root@srx1# set protocols mstp revision-level 1

root@srx1# set protocols mstp msti 1 bridge-priority 4k

root@srx1# set protocols mstp msti 1 vlan 600

root@srx1# set protocols mstp msti 1 vlan 700

root@srx1# set protocols mstp msti 2 vlan 601
```

srx2:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set protocols mstp configuration-name east
root@srx2# set protocols mstp revision-level 1
root@srx2# set protocols mstp msti 1 vlan 600
root@srx2# set protocols mstp msti 1 vlan 700
root@srx2# set protocols mstp msti 2 bridge-priority 4k
root@srx2# set protocols mstp msti 2 vlan 601
```

srx4:

```
root@srx4# set protocols mstp configuration-name east
root@srx4# set protocols mstp revision-level 1
root@srx4# set protocols mstp msti 1 vlan 600
root@srx4# set protocols mstp msti 1 vlan 700
root@srx4# set protocols mstp msti 2 vlan 601
```

ex2:

```
root@ex2# set protocols mstp configuration-name east
root@ex2# set protocols mstp revision-level 1
root@ex2# set protocols mstp bridge-priority 4k
root@ex2# set protocols mstp msti 1 vlan 600
root@ex2# set protocols mstp msti 1 vlan 700
root@ex2# set protocols mstp msti 2 vlan 601
```

Verification:

```
root@srx1# run show spanning-tree bridge

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
STP bridge parameters for CIST
  Root ID                              : 4096.2c:6b:f5:8f:87:c1
  Root cost                            : 0
  Root port                            : ge-0/0/10.0
  CIST regional root                   : 4096.2c:6b:f5:8f:87:c1
  CIST internal root cost              : 20000
  Hello time                           : 2 seconds
  Maximum age                          : 10 seconds
  Forward delay                        : 15 seconds
  Hop count                            : 4
  Message age                          : 0
  Number of topology changes           : 8
  Time since last topology change      : 13727 seconds
  Topology change initiator            : ge-0/0/10.0
  Topology change last recvd. from     : 78:19:f7:18:3d:81
  Local parameters
    Bridge ID                          : 32768.b0:a8:6e:b9:61:10
    Extended system ID                 : 0
    Internal instance ID               : 0

STP bridge parameters for MSTI 1
  MSTI regional root                   : 4097.b0:a8:6e:b9:61:10
  Hello time                           : 2 seconds
  Maximum age                          : 10 seconds
  Forward delay                        : 15 seconds
  Number of topology changes           : 5
  Topology change initiator            : ge-0/0/10.0
  Topology change last recvd. from     : 2c:6b:f5:8f:87:cd
  Local parameters
    Bridge ID                          : 4097.b0:a8:6e:b9:61:10
    Extended system ID                 : 0
    Internal instance ID               : 1

STP bridge parameters for MSTI 2
  MSTI regional root                   : 4098.78:19:f7:18:3d:90
  Root cost                            : 20000
  Root port                            : ge-0/0/1.0
  Hello time                           : 2 seconds
  Maximum age                          : 10 seconds
  Forward delay                        : 15 seconds
  Hop count                            : 4
  Number of topology changes           : 8
  Topology change initiator            : ge-0/0/1.0
  Topology change last recvd. from     : 78:19:f7:18:3d:81
  Local parameters
    Bridge ID                          : 32770.b0:a8:6e:b9:61:10
    Extended system ID                 : 0
    Internal instance ID               : 2

root@srx2# run show spanning-tree bridge
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 4096.2c:6b:f5:8f:87:c1
  Root cost                         : 0
  Root port                         : ae0.0
  CIST regional root                : 4096.2c:6b:f5:8f:87:c1
  CIST internal root cost           : 10000
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 4
  Message age                       : 0
  Number of topology changes        : 4
  Time since last topology change   : 13784 seconds
  Topology change initiator         : ae0.0
  Topology change last recvd. from  : 2e:6b:f5:8f:87:c3
  Local parameters
    Bridge ID                       : 32768.78:19:f7:18:3d:90
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 4097.b0:a8:6e:b9:61:10
  Root cost                         : 20000
  Root port                         : ge-0/0/1.0
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 4
  Number of topology changes        : 4
  Topology change initiator         : ge-0/0/1.0
  Topology change last recvd. from  : b0:a8:6e:b9:61:02
  Local parameters
    Bridge ID                       : 32769.78:19:f7:18:3d:90
    Extended system ID              : 0
    Internal instance ID            : 1

STP bridge parameters for MSTI 2
  MSTI regional root                : 4098.78:19:f7:18:3d:90
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 4
  Topology change initiator         : ae0.0
  Topology change last recvd. from  : 2e:6b:f5:8f:87:c3
  Local parameters
    Bridge ID                       : 4098.78:19:f7:18:3d:90
    Extended system ID              : 0
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Internal instance ID            : 2

root@srx4# run show spanning-tree bridge

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 4096.2c:6b:f5:8f:87:c1
  Root cost                         : 0
  Root port                         : ge-0/0/13.0
  CIST regional root                : 4096.2c:6b:f5:8f:87:c1
  CIST internal root cost           : 20000
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 4
  Message age                       : 0
  Number of topology changes        : 2
  Time since last topology change   : 13835 seconds
  Topology change initiator         : ge-0/0/13.0
  Topology change last recvd. from  : 2c:6b:f5:8f:87:d4
  Local parameters
    Bridge ID                       : 32768.3c:61:04:8d:be:90
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 32769.3c:61:04:8d:be:90
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 0
  Local parameters
    Bridge ID                       : 32769.3c:61:04:8d:be:90
    Extended system ID              : 0
    Internal instance ID            : 1

STP bridge parameters for MSTI 2
  MSTI regional root                : 32770.3c:61:04:8d:be:90
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 0
  Local parameters
    Bridge ID                       : 32770.3c:61:04:8d:be:90
    Extended system ID              : 0
    Internal instance ID            : 2

root@ex2# run show spanning-tree bridge
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 4096.2c:6b:f5:8f:87:c1
  CIST regional root                : 4096.2c:6b:f5:8f:87:c1
  CIST internal root cost           : 0
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 2
  Time since last topology change   : 13895 seconds
  Topology change initiator         : ge-0/0/12.0
  Topology change last recvd. from  : 3c:94:d5:46:73:0f
  Local parameters
    Bridge ID                       : 4096.2c:6b:f5:8f:87:c1
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 4097.b0:a8:6e:b9:61:10
  Root cost                         : 20000
  Root port                         : ge-0/0/10.0
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 4
  Number of topology changes        : 2
  Topology change initiator         : ge-0/0/12.0
  Topology change last recvd. from  : b0:a8:6e:b9:61:0a
  Local parameters
    Bridge ID                       : 32769.2c:6b:f5:8f:87:c1
    Extended system ID              : 0
    Internal instance ID            : 1

STP bridge parameters for MSTI 2
  MSTI regional root                : 4098.78:19:f7:18:3d:90
  Root cost                         : 10000
  Root port                         : ae0.0
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 4
  Number of topology changes        : 2
  Topology change initiator         : ge-0/0/12.0
  Topology change last recvd. from  : 78:19:f7:18:3e:00
  Local parameters
    Bridge ID                       : 32770.2c:6b:f5:8f:87:c1
    Extended system ID              : 0
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Internal instance ID             : 2
```

13) Configure ex1, ex3_ex4 and srx3 in a different MSTP region (*west*) with single instance 1 with VLAN 500. Regional root bridge should be ex1. Make sure that it will not become CIST root for the network:

ex1:

```
root@ex1# set protocols mstp configuration-name west
root@ex1# set protocols mstp revision-level 1
root@ex1# set protocols mstp bridge-priority 8k
root@ex1# set protocols mstp msti 1 bridge-priority 4k
root@ex1# set protocols mstp msti 1 vlan 500
```

ex3_ex4:

```
root@ex3_ex4# set protocols mstp configuration-name west
root@ex3_ex4# set protocols mstp revision-level 1
root@ex3_ex4# set protocols mstp msti 1 vlan 500
```

srx3:

```
root@srx3# set protocols mstp configuration-name west
root@srx3# set protocols mstp revision-level 1
root@srx3# set protocols mstp msti 1 vlan 500
```

Verification:

```
root@ex1# run show spanning-tree bridge

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 4096.2c:6b:f5:8f:87:c1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Root cost                          : 1000
  Root port                          : ge-0/0/12.0
  CIST regional root                 : 8192.3c:94:d5:46:73:01
  CIST internal root cost            : 0
  Hello time                         : 2 seconds
  Maximum age                        : 10 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 20
  Message age                        : 1
  Number of topology changes         : 3
  Time since last topology change    : 13876 seconds
  Topology change initiator          : ge-0/0/12.0
  Topology change last recvd. from   : 78:19:f7:18:3d:8a
  Local parameters
    Bridge ID                        : 8192.3c:94:d5:46:73:01
    Extended system ID               : 0
    Internal instance ID             : 0

STP bridge parameters for MSTI 1
  MSTI regional root                 : 4097.3c:94:d5:46:73:01
  Hello time                         : 2 seconds
  Maximum age                        : 10 seconds
  Forward delay                      : 15 seconds
  Number of topology changes         : 3
  Topology change initiator          : ge-0/0/12.0
  Topology change last recvd. from   : 2c:6b:f5:35:1b:95
  Local parameters
    Bridge ID                        : 4097.3c:94:d5:46:73:01
    Extended system ID               : 0
    Internal instance ID             : 1

  {master:0}[edit]

root@ex1# run show spanning-tree mstp configuration
MSTP information
Context identifier     : 0
Region name            : west
Revision               : 1
Configuration digest   : 0xef4bb67d5832114251b4c7d236f7bc89


MSTI     Member VLANs
   0 0-499,501-4094
   1 500

root@ex3_ex4# run show spanning-tree bridge

STP bridge parameters
Context ID                         : 0
Enabled protocol                   : MSTP
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
STP bridge parameters for CIST
  Root ID                              : 4096.2c:6b:f5:8f:87:c1
  Root cost                            : 1000
  Root port                            : ge-0/0/18.0
  CIST regional root                   : 8192.3c:94:d5:46:73:01
  CIST internal root cost              : 20000
  Hello time                           : 2 seconds
  Maximum age                          : 10 seconds
  Forward delay                        : 15 seconds
  Hop count                            : 19
  Message age                          : 1
  Number of topology changes           : 2
  Time since last topology change      : 14157 seconds
  Topology change initiator            : ge-0/0/18.0
  Topology change last recvd. from     : 3c:61:04:8d:be:89
  Local parameters
    Bridge ID                          : 32768.2c:6b:f5:35:1b:81
    Extended system ID                 : 0
    Internal instance ID               : 0

STP bridge parameters for MSTI 1
  MSTI regional root                   : 4097.3c:94:d5:46:73:01
  Root cost                            : 20000
  Root port                            : ge-0/0/18.0
  Hello time                           : 2 seconds
  Maximum age                          : 10 seconds
  Forward delay                        : 15 seconds
  Hop count                            : 19
  Number of topology changes           : 2
  Topology change initiator            : ge-0/0/18.0
  Topology change last recvd. from     : 3c:94:d5:46:73:15
  Local parameters
    Bridge ID                          : 32769.2c:6b:f5:35:1b:81
    Extended system ID                 : 0
    Internal instance ID               : 1

    {master:0}[edit]

root@ex3_ex4# run show spanning-tree mstp configuration
MSTP information
Context identifier     : 0
Region name            : west
Revision               : 1
Configuration digest   : 0xef4bb67d5832114251b4c7d236f7bc89


MSTI     Member VLANs
   0 0-499,501-4094
   1 500

root@srx3# run show spanning-tree bridge
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 4096.2c:6b:f5:8f:87:c1
  Root cost                         : 1000
  Root port                         : ge-0/0/12.0
  CIST regional root                : 8192.3c:94:d5:46:73:01
  CIST internal root cost           : 20000
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Message age                       : 1
  Number of topology changes        : 1
  Time since last topology change   : 14800 seconds
  Topology change initiator         : ge-0/0/12.0
  Topology change last recvd. from  : 2c:6b:f5:35:1b:8c
  Local parameters
    Bridge ID                       : 32768.3c:8a:b0:5c:7a:b0
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 4097.3c:94:d5:46:73:01
  Root cost                         : 20000
  Root port                         : ge-0/0/12.0
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Number of topology changes        : 1
  Topology change initiator         : ge-0/0/12.0
  Topology change last recvd. from  : 2c:6b:f5:35:1b:8c
  Local parameters
    Bridge ID                       : 32769.3c:8a:b0:5c:7a:b0
    Extended system ID              : 0
    Internal instance ID            : 1
```

14) Make sure that the best possible path to reach the CIST root bridge of the network on ex1 is the most optimal:

To achieve our objective, we simply make sure that the interface cost between ex1 and ex2 is less than default port cost (20000):

On both ex1 and ex2:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set protocols mstp interface ge-0/0/12.0 cost 1000

root@ex1# set protocols mstp interface ge-0/0/13.0 cost 1000

root@ex1# set protocols mstp msti 1 interface ge-0/0/12.0 cost 1000

root@ex1# set protocols mstp msti 1 interface ge-0/0/13.0 cost 1000
```

Verification:

```
root@ex1# run show spanning-tree bridge

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 4096.2c:6b:f5:8f:87:c1
  Root cost                         : 1000
  Root port                         : ge-0/0/12.0
  CIST regional root                : 8192.3c:94:d5:46:73:01
  CIST internal root cost           : 0
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 20
  Message age                       : 1
  Number of topology changes        : 3
  Time since last topology change   : 13876 seconds
  Topology change initiator         : ge-0/0/12.0
  Topology change last recvd. from  : 78:19:f7:18:3d:8a
  Local parameters
    Bridge ID                       : 8192.3c:94:d5:46:73:01
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 4097.3c:94:d5:46:73:01
  Hello time                        : 2 seconds
  Maximum age                       : 10 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 3
  Topology change initiator         : ge-0/0/12.0
  Topology change last recvd. from  : 2c:6b:f5:35:1b:95
  Local parameters
    Bridge ID                       : 4097.3c:94:d5:46:73:01
    Extended system ID              : 0
    Internal instance ID            : 1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# run show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port
State  Role
                        port ID         bridge ID          Cost
ae1.0            128:2        128:2  32768.b0a86eb96110     10000  BLK
ALT
ge-0/0/7.0     128:520      128:520   8192.3c94d5467301     20000  FWD
DESG
ge-0/0/10.0    128:523      128:523  32768.7819f7183d90     20000  BLK
ALT
ge-0/0/12.0    128:525      128:525   4096.2c6bf58f87c1      1000  FWD
ROOT
ge-0/0/13.0    128:526      128:526   4096.2c6bf58f87c1      1000  BLK
ALT
ge-0/0/16.0    128:529      128:529   8192.3c94d5467301     20000  FWD
DESG
ge-0/0/18.0    128:531      128:531   8192.3c94d5467301     20000  FWD
DESG

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port
State  Role
                        port ID         bridge ID          Cost
ae1.0            128:2        128:2   4097.3c94d5467301     10000  BLK
ALT
ge-0/0/10.0    128:523      128:523   4097.3c94d5467301     20000  BLK
ALT
ge-0/0/12.0    128:525      128:525   4097.3c94d5467301      1000  FWD
MSTR
ge-0/0/13.0    128:526      128:526   4097.3c94d5467301      1000  BLK
ALT
ge-0/0/16.0    128:529      128:529   4097.3c94d5467301     20000  FWD
DESG
ge-0/0/18.0    128:531      128:531   4097.3c94d5467301     20000  FWD
DESG
```

15) Ensure that interface ge-1/0/9 on ex3_ex4 is the worst path to the root bridge of instance 1.

We complete this task by assigning a high port cost to this interface:

```
root@ex3_ex4# set protocols mstp msti 1 interface ge-1/0/9.0 cost
2000000
```

Verification:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# run show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port
State  Role
                        port ID         bridge ID          Cost
ge-0/0/6.0     128:519      128:519  32768.2c6bf5351b81     20000  FWD
DESG
ge-0/0/7.0     128:520      128:520  32768.2c6bf5351b81     20000  FWD
DESG
ge-0/0/8.0     128:521      128:521  32768.2c6bf5351b81     20000  FWD
DESG
ge-0/0/9.0     128:522      128:522  32768.2c6bf5351b81     20000  FWD
DESG
ge-0/0/10.0    128:523      128:523  32768.2c6bf5351b81     20000  FWD
DESG
ge-0/0/18.0    128:531      128:531   8192.3c94d5467301     20000  FWD
ROOT
ge-1/0/6.0     128:631      128:631  32768.2c6bf5351b81     20000  FWD
DESG
ge-1/0/7.0     128:632      128:632  32768.2c6bf5351b81     20000  FWD
DESG
ge-1/0/8.0     128:633      128:633  32768.2c6bf5351b81     20000  FWD
DESG
ge-1/0/9.0     128:634      128:522  32768.3c61048dbe90     20000  BLK
ALT
ge-1/0/19.0    128:644      128:532   4096.2c6bf58f87c1     20000  BLK
ALT

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port
State  Role
                        port ID         bridge ID          Cost
ge-0/0/9.0     128:522      128:522  32769.2c6bf5351b81     20000  FWD
DESG
ge-0/0/18.0    128:531      128:531   4097.3c94d5467301     20000  FWD
ROOT
ge-1/0/6.0     128:631      128:631  32769.2c6bf5351b81     20000  FWD
DESG
ge-1/0/7.0     128:632      128:632  32769.2c6bf5351b81     20000  FWD
DESG
ge-1/0/9.0     128:634      128:634  32769.2c6bf5351b81   2000000  BLK
ALT
ge-1/0/19.0    128:644      128:644  32769.2c6bf5351b81     20000  BLK
ALT
```

16) Configuration of MSTP hop and BPDU hello timer on all devices of region east:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set protocols mstp max-hops 5

root@srx1# set protocols mstp max-age 10
```

17) Ensure tht if srx3 sneds superior BPDUs to ex1, interface ge-0/0/16 is blocked:

To prevent interface ge-0/0/16 on ex1 to become the root port, we configure it using the '**no-root-port**' configuration command under the [**protocols mstp interface**] stanza.

```
root@ex1# set protocols mstp interface ge-0/0/16.0 no-root-port
```

When superior BPDUs are received on interfaces configured with this command, the interface is blocked.

18) Configure interface ge-0/0/0 to ge-0/0/04 on ex1 and ex2 to be shutdown when BPDUs are received. Do not change the configuration in the [**protocols mstp**] stanza:

```
root@ex1# set ethernet-switching-options bpdu-block interface ge-0/0/0.0

root@ex1# set ethernet-switching-options bpdu-block interface ge-0/0/1.0

root@ex1# set ethernet-switching-options bpdu-block interface ge-0/0/2.0

root@ex1# set ethernet-switching-options bpdu-block interface ge-0/0/3.0

root@ex1# set ethernet-switching-options bpdu-block interface ge-0/0/4.0
```

19) Ensure that ports ge-0/0/0 to ge-0/0/4 and ports ge-1/0/0 to ge-1/0/4 on ex3_ex4 transition to the L2 forwarding state immediately. Do this without using the **edge** knob:

```
root@ex3_ex4# set interfaces interface-range no-stp member-range ge-0/0/0 to ge-0/0/4
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set interfaces interface-range no-stp member-range ge-1/0/0 to ge-1/0/4

root@ex3_ex4# set interfaces interface-range no-stp unit 0 family ethernet-switching

root@ex3_ex4# set protocols mstp interface no-stp disable
```

20) Configure 802.1X authentication on ex3_ex4:

```
root@ex3_ex4# set protocols dot1x authenticator authentication-profile-name dot1x

root@ex3_ex4# set protocols dot1x authenticator interface no-stp

root@ex3_ex4# set access radius-server 10.10.10.1 secret workbook

root@ex3_ex4# set access profile dot1x authentication-order radius

root@ex3_ex4# set access profile dot1x radius authentication-server 10.10.10.1
```

21) Make sure that non-802.1X enabled devices are authenticated by its MAC address on those ports:

```
root@ex3_ex4# set protocols dot1x authenticator interface no-stp mac-radius
```

22) Making sure that the host with MAC address 28:c0:da:dd:33:40 bypasses 802.1X authentication and becomes a member of VLAN 700:

```
root@ex3_ex4# set protocols dot1x authenticator static 28:c0:da:dd:33:40/48 vlan-assignment Accounting
```

23) Enable IP spoofing for the Engineering VLAN and configure static database entries for interfaces ge-1/0/6 and ge-1/0/7:

```
root@ex3_ex4# set ethernet-switching-options secure-access-port interface ge-1/0/6.0 static-ip 192.168.50.100 vlan 500
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4set ethernet-switching-options secure-access-port
interface ge-1/0/6.0 static-ip 192.168.50.100 mac 3c:61:04:8d:be:86

root@ex3_ex4set ethernet-switching-options secure-access-port
interface ge-1/0/7.0 static-ip 192.168.50.90 vlan 500

root@ex3_ex4set ethernet-switching-options secure-access-port
interface ge-1/0/7.0 static-ip 192.168.50.90 mac 3c:61:04:8d:be:87

root@ex3_ex4set ethernet-switching-options secure-access-port vlan
Engineering ip-source-guard
```

Verification:

```
root@ex3_ex4# run show dhcp snooping binding
DHCP Snooping Information:
MAC address        IP address        Lease (seconds)  Type     VLAN
Interface
3C:61:04:8D:BE:86  192.168.50.100                 -  static
Engineering ge-1/0/6.0
3C:61:04:8D:BE:87  192.168.50.90                  -  static
Engineering ge-1/0/7.0
```

24) Port mirroring configuration to copy incoming traffic on ge-0/0/18 and ge-0/0/19 to the ge-0/0/20 interface:

```
root@ex3_ex4# set ethernet-switching-options analyzer span input
ingress interface ge-0/0/18.0

root@ex3_ex4# set ethernet-switching-options analyzer span input
ingress interface ge-1/0/19.0

root@ex3_ex4# set ethernet-switching-options analyzer span output
interface ge-0/0/20.0
```

25-27) Here, we combine task 25 to 27 together.

We need to configure a VRRP group on VLAN500 between srx3 and srx4. We need to make sure that srx3 is the VRRP master, so we configure it with a priority higher then the default priority of 100. Another requirement is to make the router respond to ICMP echo requests that are send to the virtual-address. we achieve this using the '**accept-data**' keyword.

```
root@srx3# set interfaces vlan unit 500 family inet address
192.168.50.2/24 vrrp-group 1 virtual-address 192.168.50.1
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set interfaces vlan unit 500 family inet address
192.168.50.2/24 vrrp-group 1 priority 200

root@srx3# set interfaces vlan unit 500 family inet address
192.168.50.2/24 vrrp-group 1 accept-data

root@srx3# set vlans Engineering l3-interface vlan.500


root@srx4# set interfaces vlan unit 500 family inet address
192.168.50.3/24 vrrp-group 1 virtual-address 192.168.50.1

root@srx4# set interfaces vlan unit 500 family inet address
192.168.50.3/24 vrrp-group 1 accept-data

root@srx4# set vlans Engineering l3-interface vlan.500
```

```
root@srx4# run show vrrp
Interface     State       Group   VR state VR Mode   Timer    Type
Address
vlan.500      up              1   backup   Active      D  3.378 lcl
192.168.50.3
                                                             vip
192.168.50.1
                                                             mas
192.168.50.2
root@srx4# run ping routing-instance H9 192.168.50.1
PING 192.168.50.1 (192.168.50.1): 56 data bytes
64 bytes from 192.168.50.1: icmp_seq=0 ttl=64 time=8.514 ms
64 bytes from 192.168.50.1: icmp_seq=1 ttl=64 time=1.201 ms
64 bytes from 192.168.50.1: icmp_seq=2 ttl=64 time=1.203 ms
64 bytes from 192.168.50.1: icmp_seq=3 ttl=64 time=1.317 ms
```

We also configure another VRRP group on VLAN700 between srx1 and srx2. We ensure that srx1 is the VRRP master by configuring a higher priority and we make sure the that the routers respond to ICMP echo requests send to the virtual-address:

```
root@srx1# set interfaces vlan unit 700 family inet6 address
fd08:a131:17d8:ffff::2/64 vrrp-inet6-group 1 virtual-inet6-address
fd08:a131:17d8:ffff::1

root@srx1# set interfaces vlan unit 700 family inet6 address
fd08:a131:17d8:ffff::2/64 vrrp-inet6-group 1 virtual-link-local-
address fe80::1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set interfaces vlan unit 700 family inet6 address
fd08:a131:17d8:ffff::2/64 vrrp-inet6-group 1 priority 200

root@srx1# set interfaces vlan unit 700 family inet6 address
fd08:a131:17d8:ffff::2/64 vrrp-inet6-group 1 accept-data

root@srx1# set interfaces vlan unit 700 family inet6 address
fe80::2/64

root@srx1# set vlans Accounting l3-interface vlan.700


root@srx2# set interfaces vlan unit 700 family inet6 address
fd08:a131:17d8:ffff::3/64 vrrp-inet6-group 1 virtual-inet6-address
fd08:a131:17d8:ffff::1

root@srx2# set interfaces vlan unit 700 family inet6 address
fd08:a131:17d8:ffff::3/64 vrrp-inet6-group 1 virtual-link-local-
address fe80::1

root@srx2# set interfaces vlan unit 700 family inet6 address
fd08:a131:17d8:ffff::3/64 vrrp-inet6-group 1 accept-data

root@srx2# set interfaces vlan unit 700 family inet6 address
fe80::3/64

root@srx2# set vlans Accounting l3-interface vlan.700
```

```
root@srx2# run show vrrp
vlan.700      up              1   backup   Active      D  2.781 lcl
fd08:a131:17d8:ffff::3
                                                             vip
fe80::1
                                                             vip
fd08:a131:17d8:ffff::1
                                                             mas
fe80::2
```

28) Configure srx1 and srx2 for stateless autoconfiguration. Make sure that only the master router is sending router advertisements:

```
root@srx1# set protocols router-advertisement interface vlan.700
virtual-router-only

root@srx1# set protocols router-advertisement interface vlan.700
prefix fd08:a131:17d8:ffff::/64 valid-lifetime 43200
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 3: IGP Routing

1) The configuration of the virtual router SRX3a and SRX4b:

SRX3a:

```
root@srx3# set interfaces lo0 unit 1 family inet filter input protect-re-ipv4
root@srx3# set interfaces lo0 unit 1 family inet address 172.30.30.8/32
root@srx3# set interfaces lo0 unit 1 family inet6 filter input protect-re-ipv6
root@srx3# set interfaces lo0 unit 1 family inet6 address fd17:f0f4:f691:30::8/128
root@srx3# set routing-instances SRX3a instance-type virtual-router
root@srx3# set routing-instances SRX3a interface ge-0/0/5.58
root@srx3# set routing-instances SRX3a interface ge-0/0/5.68
root@srx3# set routing-instances SRX3a interface ge-0/0/5.81
root@srx3# set routing-instances SRX3a interface lo0.1
```

SRX4b:

```
root@srx4# set interfaces lo0 unit 1 family inet filter input protect-re-ipv4
root@srx4# set interfaces lo0 unit 1 family inet address 172.30.30.9/32
root@srx4# set interfaces lo0 unit 1 family inet6 filter input protect-re-ipv6
root@srx4# set interfaces lo0 unit 1 family inet6 address fd17:f0f4:f691:30::9/128
root@srx4# set routing-instances SRX4b instance-type virtual-router
root@srx4# set routing-instances SRX4b interface ge-0/0/5.49
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx4# set routing-instances SRX4b interface ge-0/0/5.69

root@srx4# set routing-instances SRX4b interface lo0.1
```

2) Configure the layer 3 interfaces based on table 29. We show srx1 here as an example:

```
root@srx1# set interfaces ge-0/0/4 vlan-tagging

root@srx1# set interfaces ge-0/0/4 unit 12 description "L3 connection
to SRX2"

root@srx1# set interfaces ge-0/0/4 unit 12 vlan-id 12

root@srx1# set interfaces ge-0/0/4 unit 12 family inet address
172.30.0.1/30

root@srx1# set interfaces ge-0/0/4 unit 12 family inet6 address
fd17:f0f4:f691::1/126

root@srx1# set interfaces ge-0/0/4 unit 13 description "L3 connection
to SRX3"

root@srx1# set interfaces ge-0/0/4 unit 13 vlan-id 13

root@srx1# set interfaces ge-0/0/4 unit 13 family inet address
172.30.0.5/30

root@srx1# set interfaces ge-0/0/4 unit 13 family inet6 address
fd17:f0f4:f691::5/126

root@srx1# set interfaces ge-0/0/4 unit 15 description "L3 connection
to EX1"

root@srx1# set interfaces ge-0/0/4 unit 15 vlan-id 15

root@srx1# set interfaces ge-0/0/4 unit 15 family inet address
172.30.0.9/30

root@srx1# set interfaces ge-0/0/4 unit 15 family inet6 address
fd17:f0f4:f691::9/126
```

In some cases, it can make sense to combine several tasks. When tasks are closely related to each other, this can speed things up significantly. In this case, we can combine tasks 3 to task 7 together. This is

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

because all the configuration tasks take place in the same stanza and they reveolve around the same topic: OSPF.

Let's briefly summarize task 3 to task 7 first:

3) We configure OSPFv2 for IPv4 and OSPFv3 for IPv6. Inside area 10, OSPFv3 is used for both IPv4 and IPv6 routing. Area 2 is configured to be an nssa and we ensure that all routers have a router-id that equals the loopback 0 IP address.
4) We ensure no hello messages are send on the loopback interfaces by configuring those interfaces with the **passive** option.
5) We disable the DR and BDR election on all links in Area 0 and Area 10 by configuring the links as point-to-point links. This is done by configuring the interfaces with the **interface-type p2p** option
6) This was achieved by configuring **priority 0** on the interfaces.
7) Inside Area 1, we make sure interfaces are configured with the '**md5 1 key workbook**' configuration in order to make sure that MD5 authentication is enabled.

Now we can start the configuration of all these things, touching every router only once:

srx1:

```
root@srx1# set routing-options router-id 172.30.30.1
```

```
root@srx1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.12
interface-type p2p

root@srx1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.13
interface-type p2p

root@srx1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.15
interface-type p2p

root@srx1# set protocols ospf area 0.0.0.0 interface lo0.0 passive

root@srx1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.40
passive
```

```
root@srx1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.12
interface-type p2p

root@srx1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.13
interface-type p2p
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.15
interface-type p2p

root@srx1# set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
```

srx2:

```
root@srx2# set routing-options router-id 172.30.30.2


root@srx2# set protocols ospf area 0.0.0.0 interface ge-0/0/4.12
interface-type p2p

root@srx2# set protocols ospf area 0.0.0.0 interface ge-0/0/4.24
interface-type p2p

root@srx2# set protocols ospf area 0.0.0.0 interface lo0.0 passive

root@srx2# set protocols ospf area 0.0.0.1 interface ge-0/0/4.25
authentication md5 1 key workbook

root@srx2# set protocols ospf area 0.0.0.1 interface ge-0/0/4.26
authentication md5 1 key workbook


root@srx2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.12
interface-type p2p

root@srx2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.24
interface-type p2p

root@srx2# set protocols ospf3 area 0.0.0.0 interface lo0.0 passive

root@srx2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.41
passive

root@srx2# set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.25

root@srx2# set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.26
```

srx3:

```
root@srx3# set routing-options router-id 172.30.30.3
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set protocols ospf area 0.0.0.0 interface ge-0/0/4.13
interface-type p2p

root@srx3# set protocols ospf area 0.0.0.0 interface ge-0/0/4.35
interface-type p2p

root@srx3# set protocols ospf area 0.0.0.0 interface lo0.0 passive

root@srx3# set protocols ospf area 0.0.0.0 interface ge-0/0/4.44
passive

root@srx3# set protocols ospf area 0.0.0.2 nssa

root@srx3# set protocols ospf area 0.0.0.2 interface ge-0/0/4.37


root@srx3# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.13
interface-type p2p

root@srx3# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.35
interface-type p2p

root@srx3# set protocols ospf3 area 0.0.0.0 interface lo0.0 passive

root@srx3# set protocols ospf3 area 0.0.0.2 nssa

root@srx3# set protocols ospf3 area 0.0.0.2 interface ge-0/0/4.37
```

srx4:

```
root@srx4# set routing-options router-id 172.30.30.4


root@srx4# set protocols ospf area 0.0.0.0 interface lo0.0 passive

root@srx4# set protocols ospf area 0.0.0.0 interface ge-0/0/4.24
interface-type p2p

root@srx4# set protocols ospf area 0.0.0.1 interface ge-0/0/4.64
authentication md5 1 key workbook


root@srx4# set protocols ospf3 realm ipv4-unicast area 0.0.0.10
interface ge-0/0/4.49 interface-type p2p
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx4# set protocols ospf3 area 0.0.0.0 interface lo0.0 passive

root@srx4# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.24
interface-type p2p

root@srx4# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.45
passive

root@srx4# set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.64

root@srx4# set protocols ospf3 area 0.0.0.10 interface ge-0/0/4.49
interface-type p2p
```

ex1:

```
root@ex1# set routing-options router-id 172.30.30.5


root@ex1# set protocols ospf area 0.0.0.0 interface ge-0/0/14.15
interface-type p2p

root@ex1# set protocols ospf area 0.0.0.0 interface ge-0/0/14.35
interface-type p2p

root@ex1# set protocols ospf area 0.0.0.0 interface lo0.0 passive

root@ex1# set protocols ospf area 0.0.0.1 interface ge-0/0/14.25
authentication md5 1 key workbook

root@ex1# set protocols ospf area 0.0.0.1 interface ge-0/0/14.56
authentication md5 1 key workbook

root@ex1# set protocols ospf area 0.0.0.1 interface ge-0/0/14.58
authentication md5 1 key workbook

root@ex1# set protocols ospf area 0.0.0.2 interface ge-0/0/14.57


root@ex1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/14.15
interface-type p2p

root@ex1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/14.35
interface-type p2p

root@ex1# set protocols ospf3 area 0.0.0.0 interface lo0.0 passive

root@ex1# set protocols ospf3 area 0.0.0.1 interface ge-0/0/14.25
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set protocols ospf3 area 0.0.0.1 interface ge-0/0/14.56

root@ex1# set protocols ospf3 area 0.0.0.1 interface ge-0/0/14.58

root@ex1# set protocols ospf3 area 0.0.0.2 interface ge-0/0/14.57
```

ex2:

```
root@ex2# set routing-options router-id 172.30.30.6


root@ex2# set protocols ospf area 0.0.0.1 interface ge-0/0/14.26
authentication md5 1 key workbook

root@ex2# set protocols ospf area 0.0.0.1 interface ge-0/0/14.64
authentication md5 1 key workbook

root@ex2# set protocols ospf area 0.0.0.1 interface ge-0/0/14.68
authentication md5 1 key workbook

root@ex2# set protocols ospf area 0.0.0.1 interface ge-0/0/14.56
authentication md5 1 key workbook

root@ex2# set protocols ospf area 0.0.0.1 interface lo0.0 passive


root@ex2# set protocols ospf3 realm ipv4-unicast area 0.0.0.10
interface ge-0/0/14.69 interface-type p2p

root@ex2# set protocols ospf3 area 0.0.0.1 interface ge-0/0/14.26

root@ex2# set protocols ospf3 area 0.0.0.1 interface ge-0/0/14.64

root@ex2# set protocols ospf3 area 0.0.0.1 interface ge-0/0/14.68

root@ex2# set protocols ospf3 area 0.0.0.1 interface ge-0/0/14.56

root@ex2# set protocols ospf3 area 0.0.0.1 interface lo0.0 passive

root@ex2# set protocols ospf3 area 0.0.0.10 interface ge-0/0/14.69
interface-type p2p
```

ex3_ex4:

```
root@ex3_ex4# set routing-options router-id 172.30.30.7
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set protocols ospf area 0.0.0.2 nssa

root@ex3_ex4# set protocols ospf area 0.0.0.2 interface ge-0/0/14.37
priority 0

root@ex3_ex4# set protocols ospf area 0.0.0.2 interface ge-0/0/14.57
priority 0

root@ex3_ex4# set protocols ospf area 0.0.0.2 interface lo0.0 passive

root@ex3_ex4# set protocols ospf area 0.0.0.2 interface ge-0/0/14.80
passive

root@ex3_ex4# set protocols ospf area 0.0.0.2 interface ge-0/0/14.43
passive


root@ex3_ex4# set protocols ospf3 area 0.0.0.2 nssa

root@ex3_ex4# set protocols ospf3 area 0.0.0.2 interface ge-0/0/14.37
priority 0

root@ex3_ex4# set protocols ospf3 area 0.0.0.2 interface ge-0/0/14.57
priority 0

root@ex3_ex4# set protocols ospf3 area 0.0.0.2 interface lo0.0 passive

root@ex3_ex4# set protocols ospf3 area 0.0.0.2 interface ge-0/0/14.80
passive
```

SRX3a:

```
root@srx3# set routing-instances SRX3a routing-options router-id
172.30.30.8


root@srx3# set routing-instances SRX3a protocols ospf area 0.0.0.1
interface ge-0/0/5.58 authentication md5 1 key workbook

root@srx3# set routing-instances SRX3a protocols ospf area 0.0.0.1
interface ge-0/0/5.68 authentication md5 1 key workbook

root@srx3# set routing-instances SRX3a protocols ospf area 0.0.0.1
interface ge-0/0/5.81 passive
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set routing-instances SRX3a protocols ospf area 0.0.0.1
interface lo0.1 passive


root@srx3# set routing-instances SRX3a protocols ospf3 area 0.0.0.1
interface ge-0/0/5.58

root@srx3# set routing-instances SRX3a protocols ospf3 area 0.0.0.1
interface ge-0/0/5.68

root@srx3# set routing-instances SRX3a protocols ospf3 area 0.0.0.1
interface ge-0/0/5.81 passive

root@srx3# set routing-instances SRX3a protocols ospf3 area 0.0.0.1
interface lo0.1 passive
```

SRX4b:

```
root@srx4# set routing-instances SRX4b protocols ospf3 realm ipv4-
unicast area 0.0.0.10 interface ge-0/0/5.49 interface-type p2p

root@srx4# set routing-instances SRX4b protocols ospf3 realm ipv4-
unicast area 0.0.0.10 interface ge-0/0/5.69 interface-type p2p

root@srx4# set routing-instances SRX4b protocols ospf3 realm ipv4-
unicast area 0.0.0.10 interface lo0.1 passive

root@srx4# set routing-instances SRX4b protocols ospf3 area 0.0.0.10
interface ge-0/0/5.49 interface-type p2p

root@srx4# set routing-instances SRX4b protocols ospf3 area 0.0.0.10
interface ge-0/0/5.69 interface-type p2p

root@srx4# set routing-instances SRX4b protocols ospf3 area 0.0.0.10
interface lo0.1 passive
```

Verification:

srx1:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.2       ge-0/0/4.12            Full      172.30.30.2      128
34
172.30.0.6       ge-0/0/4.13            Full      172.30.30.3      128
39
172.30.0.10      ge-0/0/4.15            Full      172.30.30.5      128
38
```

srx2:

```
root@srx2# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.1       ge-0/0/4.12            Full      172.30.30.1      128
36
172.30.0.14      ge-0/0/4.24            Full      172.30.30.4      128
33
172.30.0.66      ge-0/0/4.25            Full      172.30.30.5      128
39
172.30.0.70      ge-0/0/4.26            Full      172.30.30.6      128
37

root@srx2# run show ospf3 neighbor
ID               Interface              State     Pri   Dead
172.30.30.1      ge-0/0/4.12            Full      128     39
  Neighbor-address fe80::b2a8:6e00:cb9:6104
172.30.30.4      ge-0/0/4.24            Full      128     37
  Neighbor-address fe80::3e61:400:188d:be84
172.30.30.5      ge-0/0/4.25            Full      128     35
  Neighbor-address fe80::3e94:d500:1946:7311
172.30.30.6      ge-0/0/4.26            Full      128     38
  Neighbor-address fe80::2e6b:f500:1a8f:87d1
```

srx3:

```
root@srx3# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.5       ge-0/0/4.13            Full      172.30.30.1      128
33
172.30.0.17      ge-0/0/4.35            Full      172.30.30.5      128
37
172.30.0.133     ge-0/0/4.37            Full      172.30.30.7        0
32

root@srx3# run show ospf3 neighbor
ID               Interface              State     Pri   Dead
172.30.30.1      ge-0/0/4.13            Full      128     32
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Neighbor-address fe80::b2a8:6e00:db9:6104
172.30.30.5      ge-0/0/4.35            Full      128     30
  Neighbor-address fe80::3e94:d500:2346:7311
172.30.30.7      ge-0/0/4.37            Full        0     32
  Neighbor-address fe80::2e6b:f500:2535:1b91
```

srx4:

```
root@srx4# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.13      ge-0/0/4.24            Full      172.30.30.2      128
34
172.30.0.81      ge-0/0/4.64            Full      172.30.30.6      128
35

root@srx4# run show ospf3 neighbor
ID               Interface              State     Pri    Dead
172.30.30.2      ge-0/0/4.24            Full      128      32
  Neighbor-address fe80::7a19:f700:1818:3d84
172.30.30.6      ge-0/0/4.64            Full      128      30
  Neighbor-address fe80::2e6b:f500:408f:87d1
172.30.30.9      ge-0/0/4.49            Full        0      32
  Neighbor-address fe80::3e61:400:318d:be85

root@ex1# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.65      ge-0/0/14.25           ExStart   172.30.30.2      128
38
172.30.0.78      ge-0/0/14.58           ExStart   172.30.30.8      128
39
172.30.0.130     ge-0/0/14.57           2Way      172.30.30.7        0
38
```

To troubleshoot the neighbor relationship formation process on the ex1 switch, we activate traceoptions:

```
root@ex1# set protocols ospf traceoptions file ospf.log

root@ex1# set protocols ospf traceoptions flag all
```

```
root@ex1# run show log ospf.log | match mtu
Dec 30 14:57:17.445120 OSPF packet ignored: MTU mismatch from
172.30.0.9 on intf ge-0/0/14.15 area 0.0.0.0
Dec 30 14:57:17.446214 OSPF packet ignored: MTU mismatch from
172.30.0.18 on intf ge-0/0/14.35 area 0.0.0.0
Dec 30 14:57:18.616598   options 0x52, i 1, m 1, ms 1, r 0, seq
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
0xac1fa135, mtu 1500
Dec 30 14:57:18.616825 OSPF packet ignored: MTU mismatch from
172.30.0.74 on intf ge-0/0/14.56 area 0.0.0.1
Dec 30 14:57:18.703923   options 0x50, i 1, m 1, ms 1, r 0, seq
0xac1c1863, mtu 1382
Dec 30 14:57:19.133818   options 0x52, i 1, m 1, ms 1, r 0, seq
0xac1bac9f, mtu 1382
Dec 30 14:57:19.138539   options 0x52, i 0, m 0, ms 0, r 0, seq
0xac1bac9f, mtu 1500
Dec 30 14:57:19.142421 OSPF packet ignored: MTU mismatch from
172.30.0.65 on intf ge-0/0/14.25 area 0.0.0.1
Dec 30 14:57:20.026833   options 0x50, i 1, m 1, ms 1, r 0, seq
0xac1e3d9a, mtu 1500
Dec 30 14:57:20.027047 OSPF packet ignored: MTU mismatch from
172.30.0.130 on intf ge-0/0/14.57 area 0.0.0.2
```

We can very easily spot the MTU mismatch by examining the **ospf.log** file. To fix the issue, we remove the MTU statement that is causing the problem:

```
root@ex1# delete interfaces ge-0/0/14 mtu
```

ex1:

```
root@ex1# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.9       ge-0/0/14.15           Full      172.30.30.1      128
39
172.30.0.18      ge-0/0/14.35           Full      172.30.30.3      128
38
172.30.0.65      ge-0/0/14.25           Full      172.30.30.2      128
35
172.30.0.74      ge-0/0/14.56           Full      172.30.30.6      128
36
172.30.0.78      ge-0/0/14.58           Full      172.30.30.8      128
36
172.30.0.130     ge-0/0/14.57           Full      172.30.30.7        0
38

root@ex1# run show ospf3 neighbor
ID               Interface              State     Pri    Dead
172.30.30.1      ge-0/0/14.15           Full      128      31
  Neighbor-address fe80::b2a8:6e00:fb9:6104
172.30.30.3      ge-0/0/14.35           Full      128      37
  Neighbor-address fe80::3e8a:b000:235c:7aa4
172.30.30.2      ge-0/0/14.25           Full      128      37
  Neighbor-address fe80::7a19:f700:1918:3d84
172.30.30.6      ge-0/0/14.56           Full      128      30
  Neighbor-address fe80::2e6b:f500:388f:87d1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
172.30.30.8      ge-0/0/14.58            Full      128      35
  Neighbor-address fe80::3e8a:b000:3a5c:7aa5
172.30.30.7      ge-0/0/14.57            Full        0      30
  Neighbor-address fe80::2e6b:f500:3935:1b91
```

ex2:

```
root@ex2# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.69      ge-0/0/14.26           Full      172.30.30.2      128
32
172.30.0.73      ge-0/0/14.56           Full      172.30.30.5      128
32
172.30.0.82      ge-0/0/14.64           Full      172.30.30.4      128
31
172.30.0.86      ge-0/0/14.68           Full      172.30.30.8      128
35

root@ex2# run show ospf3 neighbor
ID               Interface              State     Pri    Dead
172.30.30.2      ge-0/0/14.26           Full      128      36
  Neighbor-address fe80::7a19:f700:1a18:3d84
172.30.30.5      ge-0/0/14.56           Full      128      30
  Neighbor-address fe80::3e94:d500:3846:7311
172.30.30.4      ge-0/0/14.64           Full      128      36
  Neighbor-address fe80::3e61:400:408d:be84
172.30.30.8      ge-0/0/14.68           Full      128      38
  Neighbor-address fe80::3e8a:b000:445c:7aa5
172.30.30.9      ge-0/0/14.69           Full      128      37
  Neighbor-address fe80::3e61:400:458d:be85
```

ex3_ex4:

```
root@ex3_ex4# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.134     ge-0/0/14.37           Full      172.30.30.3      128
32
172.30.0.129     ge-0/0/14.57           Full      172.30.30.5      128
34

root@ex3_ex4# run show ospf3 neighbor
ID               Interface              State     Pri    Dead
172.30.30.3      ge-0/0/14.37           Full      128      32
  Neighbor-address fe80::3e8a:b000:255c:7aa4
172.30.30.5      ge-0/0/14.57           Full      128      35
  Neighbor-address fe80::3e94:d500:3946:7311
```

SRX3a:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# run show ospf3 neighbor instance SRX3a
ID              Interface             State     Pri   Dead
172.30.30.5     ge-0/0/5.58           Full      128     38
  Neighbor-address fe80::3e94:d500:3a46:7311
172.30.30.6     ge-0/0/5.68           Full      128     36
  Neighbor-address fe80::2e6b:f500:448f:87d1
```

SRX4a:

```
root@srx4# run show ospf3 neighbor instance SRX4b
ID              Interface             State     Pri   Dead
172.30.30.4     ge-0/0/5.49           Full      128     37
  Neighbor-address fe80::3e61:400:318d:be84
172.30.30.6     ge-0/0/5.69           Full      128     36
  Neighbor-address fe80::2e6b:f500:458f:87d1
```

```
root@ex3_ex4# run show ospf interface
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/14.37        DRother 0.0.0.2         172.30.30.3     0.0.0.0
1
ge-0/0/14.43        DRother 0.0.0.2         0.0.0.0         0.0.0.0
0
ge-0/0/14.57        DRother 0.0.0.2         172.30.30.5     0.0.0.0
1
ge-0/0/14.80        DRother 0.0.0.2         0.0.0.0         0.0.0.0
0
lo0.0               DRother 0.0.0.2         0.0.0.0         0.0.0.0
0

{master:0}[edit]
root@ex3_ex4# run show ospf3 interface
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/14.37        DRother 0.0.0.2         172.30.30.3     0.0.0.0
1
ge-0/0/14.57        DRother 0.0.0.2         172.30.30.5     0.0.0.0
1
ge-0/0/14.80        DRother 0.0.0.2         0.0.0.0         0.0.0.0
0
lo0.0               DRother 0.0.0.2         0.0.0.0         0.0.0.0
0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

8) The best path to reach the srx2 IPv4 lookback from srx3a should be via ex2.

We first manipulate the interface metrics to make the following links more costly than others:

On srx3:

```
root@srx3# set routing-instances SRX3a protocols ospf area 0.0.0.1
interface ge-0/0/5.58 metric 100

root@srx3# set routing-instances SRX3a protocols ospf3 area 0.0.0.1
interface ge-0/0/5.58 metric 100
```

On ex1:

```
root@ex1# set protocols ospf area 0.0.0.1 interface ge-0/0/14.58
metric 100
```

Verification:

```
root@srx3# run show route table SRX3a.inet.0 172.30.30.2

SRX3a.inet.0: 538 destinations, 539 routes (538 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

172.30.30.2/32     *[OSPF/10] 00:00:46, metric 101
                    > to 172.30.0.77 via ge-0/0/5.58


root@srx3# run traceroute 172.30.30.2 routing-instance SRX3a no-
resolve
traceroute to 172.30.30.2 (172.30.30.2), 30 hops max, 40 byte packets
 1  172.30.0.77  3.034 ms  3.216 ms  2.438 ms
 2  172.30.30.2  3.143 ms  12.668 ms  11.131 ms
```

After examining the traceroute, we can see that the path is still via ex1. Therfore, we more to ex2 and have a closer look at the configuration.

As we can see, there is a group configuration active on ex2:

```
groups {
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    initial {
        protocols {
            ospf {
                overload;
            }
        }
    }
}
```

The overload option is altering the cost of thcauses this behavior. The overload mode making the device appear overloaded by increasing the metric for all of the links. We should deactivate this group to fix the issue.

```
root@ex2# deactivate protocols ospf apply-groups
```

After deactivating the group, we examine the routing-table and perform another traceroute:

```
root@srx3# run show route table SRX3a.inet.0 172.30.30.2

SRX3a.inet.0: 539 destinations, 540 routes (539 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

172.30.30.2/32     *[OSPF/10] 04:32:03, metric 2
                    > to 172.30.0.85 via ge-0/0/5.68




root@srx3# run traceroute 172.30.30.2 routing-instance SRX3a no-
resolve

traceroute to 172.30.30.2 (172.30.30.2), 30 hops max, 40 byte packets
 1  172.30.0.85  4.614 ms  1.793 ms  2.728 ms
 2  172.30.30.2  3.006 ms  2.864 ms  2.172 ms
```

9) Earlier on, we already configured Area 2 to be an **nssa**. This means that we will not encounter any Type 5 LSA in this area. In order to make sure that the Type 3 LSA is not seen either, we use the **no-summaries** option. To advertise a default gateway into the NSSA, we use the **default-metric** configuration command:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On srx3 and ex1:

```
root@srx3# set protocols ospf area 0.0.0.2 nssa default-lsa default-
metric 10

root@srx3# set protocols ospf area 0.0.0.2 nssa default-lsa type-7

root@srx3# set protocols ospf area 0.0.0.2 nssa no-summaries

root@srx3# set protocols ospf3 area 0.0.0.2 nssa default-lsa default-
metric 10

root@srx3# set protocols ospf3 area 0.0.0.2 nssa default-lsa type-7

root@srx3# set protocols ospf3 area 0.0.0.2 nssa no-summaries
```

Verification:

```
root@ex3_ex4# run show route
inet.0: 518 destinations, 518 routes (517 active, 0 holddown, 1
hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 00:29:03, metric 11, tag 0
                    > to 172.30.0.134 via ge-0/0/14.37
                      to 172.30.0.129 via ge-0/0/14.57
```

10) Summarize IPv4 and IPv6 routes to all non-backbone areas on all ABRs except on srx4. On srx4, only advertise its loopback IP address and suppress all other backbone routes:

The network summarization for area 0 is done by using the **area-range** configuration command.

On srx3, ex1 and ex2:

```
root@srx3# set protocols ospf area 0.0.0.0 area-range 172.31.0.0/20

root@srx3# set protocols ospf3 area 0.0.0.0 area-range
fd17:f0f4:f691:f::/64
```

To complete the objectives specified for srx4, we turn to the **network-summary-import** command:

On srx4:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx4# set policy-options policy-statement no.leak term 1 from
route-filter 172.30.30.4/32 exact

root@srx4# set policy-options policy-statement no.leak term 1 then
accept

root@srx4# set policy-options policy-statement no.leak term 2 then
reject

root@srx4# set protocols ospf area 0.0.0.0 network-summary-import
no.leak
```

On ex2:

```
root@ex2# run show ospf database netsummary advertising-router
172.30.30.4
  OSPF database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt
Cksum  Len
Summary  172.30.30.4      172.30.30.4      0x80000008   313  0x22
0xeb70  28
```

11) Make sure that no Area can become isolated for both OSPFv2 and OSPFv3:

We complete our objective by defining virtual-links from ex2 and srx4 to area 0:

```
root@srx2# set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id
172.30.30.2 transit-area 0.0.0.1

root@srx2# set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id
172.30.30.5 transit-area 0.0.0.1


root@srx4# set protocols ospf area 0.0.0.0 virtual-link neighbor-id
172.30.30.5 transit-area 0.0.0.1

root@srx4# set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id
172.30.30.5 transit-area 0.0.0.1


root@ex1# set protocols ospf area 0.0.0.0 virtual-link neighbor-id
172.30.30.4 transit-area 0.0.0.1

root@ex1# set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id
172.30.30.4 transit-area 0.0.0.1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id
172.30.30.6 transit-area 0.0.0.1


root@ex2# set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id
172.30.30.2 transit-area 0.0.0.1

root@ex2# set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id
172.30.30.5 transit-area 0.0.0.1
```

Verification:

```
root@ex2# run show ospf3 neighbor
ID              Interface              State     Pri   Dead
172.30.30.2     vl-172.30.30.2         Full        0     36
  Neighbor-address fd17:f0f4:f691:1::5
172.30.30.5     vl-172.30.30.5         Full        0     31
  Neighbor-address fd17:f0f4:f691:1::9
172.30.30.2     ge-0/0/14.26           Full      128     36
  Neighbor-address fe80::7a19:f700:1a18:3d84
172.30.30.5     ge-0/0/14.56           Full      128     30
  Neighbor-address fe80::3e94:d500:3846:7311
172.30.30.4     ge-0/0/14.64           Full      128     36
  Neighbor-address fe80::3e61:400:408d:be84
172.30.30.8     ge-0/0/14.68           Full      128     38
  Neighbor-address fe80::3e8a:b000:445c:7aa5
172.30.30.9     ge-0/0/14.69           Full      128     37
  Neighbor-address fe80::3e61:400:458d:be85

root@ex1# run show ospf neighbor
Address          Interface              State     ID               Pri
Dead
172.30.0.9       ge-0/0/14.15           Full      172.30.30.1      128
39
172.30.0.18      ge-0/0/14.35           Full      172.30.30.3      128
35
172.30.0.82      vl-172.30.30.4         Full      172.30.30.4        0
33
172.30.0.65      ge-0/0/14.25           Full      172.30.30.2      128
33
172.30.0.74      ge-0/0/14.56           Full      172.30.30.6      128
33
172.30.0.78      ge-0/0/14.58           Full      172.30.30.8      128
36
172.30.0.130     ge-0/0/14.57           Full      172.30.30.7        0
39
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
172.30.30.1     ge-0/0/14.15          Full      128     38
  Neighbor-address fe80::b2a8:6e00:fb9:6104
172.30.30.3     ge-0/0/14.35          Full      128     33
  Neighbor-address fe80::3e8a:b000:235c:7aa4
172.30.30.4     vl-172.30.30.4        Full        0     36
  Neighbor-address fd17:f0f4:f691:1::12
172.30.30.6     vl-172.30.30.6        Full        0     37
  Neighbor-address fd17:f0f4:f691:1::a
172.30.30.2     ge-0/0/14.25          Full      128     39
  Neighbor-address fe80::7a19:f700:1918:3d84
172.30.30.6     ge-0/0/14.56          Full      128     37
  Neighbor-address fe80::2e6b:f500:388f:87d1
172.30.30.8     ge-0/0/14.58          Full      128     32
  Neighbor-address fe80::3e8a:b000:3a5c:7aa5
172.30.30.7     ge-0/0/14.57          Full        0     32
  Neighbor-address fe80::2e6b:f500:3935:1b91
```

12) RIP and RIPng should be configured to DC1 on ex3_ex4 and srx3:

First, we make sure that the firewall filter applied to the loopback interface permits RIP. We configure a new term for the IPv4 as well as the IPv6 term:

On both devices:

```
root@ex3_ex4# set firewall family inet filter protect-re-ipv4 term rip
from protocol udp

root@ex3_ex4#set firewall family inet filter protect-re-ipv4 term rip
from destination-port 520

root@ex3_ex4#set firewall family inet filter protect-re-ipv4 term rip
then accept


root@ex3_ex4#set firewall family inet6 filter protect-re-ipv6 term
ripng from next-header udp

root@ex3_ex4#set firewall family inet6 filter protect-re-ipv6 term
ripng from destination-port 521

root@ex3_ex4#set firewall family inet6 filter protect-re-ipv6 term
ripng then accept
```

We then insert the term high enough to avoid any discards:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ ex3_ex4#insert firewall family inet filter protect-re-ipv4 term
rip before term icmp

root@ ex3_ex4#insert firewall family inet filter protect-re-ipv6 term
ripng before term icmp
```

Finally, we enable both RIP and RIPng.

On ex3_ex4:

```
root@ ex3_ex4#set protocols rip group DC1 neighbor ge-0/0/14.80

root@ ex3_ex4#set protocols ripng group DC1 neighbor ge-0/0/14.80
```

Verification:

```
root@ex3_ex4# run show rip neighbor
                    Local  Source           Destination      Send
Receive   In
Neighbor            State  Address          Address          Mode
Mode   Met
--------            -----  -------          -----------      ----
------  ---
ge-0/0/14.80          Up 172.30.80.1     224.0.0.9        mcast
both   1

root@ex3_ex4# run show ripng neighbor
                      Source                              Dest
In
Neighbor      State   Address                             Address
Send  Recv Met
--------      -----   -------                             -------
---- --- ---
ge-0/0/14.80     Up fe80::2e6b:f500:5035:1b91          ff02::9
yes yes  1
```

```
root@srx3#set routing-instances SRX3a protocols rip group DC1 neighbor
ge-0/0/5.81

root@srx3#set routing-instances SRX3a protocols ripng group DC1
neighbor ge-0/0/5.81
```

Verification:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# run show route protocol rip
172.31.0.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.1.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.2.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.3.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.4.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.5.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.6.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.7.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.8.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.9.0/24      *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.10.0/24     *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.11.0/24     *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.12.0/24     *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.13.0/24     *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.14.0/24     *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
172.31.15.0/24     *[RIP/100] 00:48:02, metric 3, tag 0
                    > to 172.30.81.2 via ge-0/0/5.81
224.0.0.9/32       *[RIP/100] 00:49:04, metric 1
                      MultiRecv

root@srx3# run show route protocol ripng

fd17:f0f4:f691:f::/80
                   *[RIPng/100] 00:48:12, metric 3, tag 0
                    > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:1::/80
                   *[RIPng/100] 00:48:12, metric 3, tag 0
                    > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:2::/80
                   *[RIPng/100] 00:48:12, metric 3, tag 0
                    > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:3::/80
                   *[RIPng/100] 00:48:12, metric 3, tag 0
                    > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:4::/80
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:5::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:6::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:7::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:8::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:9::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:a::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:b::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:c::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:d::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:e::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
fd17:f0f4:f691:f:f::/80
                    *[RIPng/100] 00:48:12, metric 3, tag 0
                     > to fe80::21f:1200:512e:6f81 via ge-0/0/5.81
ff02::9/128         *[RIPng/100] 00:49:14, metric 1
                       MultiRecv
```

13) In this task we should advertise default route to DC1. As we have already generated default route insde area 2, we can use this default route in our policy on ex3_ex4. On SRX3a, we use generate route.

To export the default route to RIP and RIPng, we create an apply an export policy to both RIP and RIPng:

ex3_ex4:

```
root@ex3_ex4# set policy-options policy-statement default.to.rip term
1 from protocol ospf
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set policy-options policy-statement default.to.rip term
1 from route-filter 0.0.0.0/0 exact

root@ex3_ex4# set policy-options policy-statement default.to.rip term
1 then accept

root@ex3_ex4# set policy-options policy-statement default.to.ripng
term 1 from protocol ospf3

root@ex3_ex4# set policy-options policy-statement default.to.ripng
term 1 from route-filter ::/0 exact

root@ex3_ex4# set policy-options policy-statement default.to.ripng
term 1 then accept

root@ex3_ex4# set protocols rip group DC1 export default.to.rip

root@ex3_ex4# set protocols ripng group DC1 export default.to.ripng
```

A caveat here is that the route preference for OSPF external routes (150) is higher then the route preference for RIP/RIPng (105). On ex3_ex4, we alter the external preference for both OSPFv2 as well as OSPFv3. This ensures optimal routing by preferring the OSPF route:

```
root@ex3_ex4# set protocols ospf external-preference 90

root@ex3_ex4# set protocols ospf3 external-preference 90
```

On srx3a, we lower the route-preference of the generate route to 90, making it more attractive than the RIP/RIPng route:

SRX3a:

```
root@srx3# set routing-instances SRX3a routing-options generate route
0.0.0.0/0 preference 90

root@srx3# set routing-instances SRX3a routing-options rib
SRX3a.inet6.0 generate route ::/0 preference 90

root@srx3# set policy-options policy-statement default.to.rip term 1
from protocol aggregate

root@srx3# set policy-options policy-statement default.to.rip term 1
from route-filter 0.0.0.0/0 exact

root@srx3# set policy-options policy-statement default.to.rip term 1
then accept
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set policy-options policy-statement default.to.ripng term 1
from protocol aggregate

root@srx3# set policy-options policy-statement default.to.ripng term 1
from route-filter ::/0 exact

root@srx3# set policy-options policy-statement default.to.ripng term 1
then accept

root@srx3# set routing-instances SRX3a protocols rip group DC1 export
default.to.rip

root@srx3# set routing-instances SRX3a protocols ripng group DC1
export default.to.ripng
```

Verification:

```
root@ex3_ex4# run show route

0.0.0.0/0          *[OSPF/90] 00:29:03, metric 11, tag 0
                    > to 172.30.0.134 via ge-0/0/14.37
                      to 172.30.0.129 via ge-0/0/14.57

root@srx3# run show route table SRX3a.inet.0

0.0.0.0/0          *[Aggregate/90] 05:14:23
                    > to 172.30.0.85 via ge-0/0/5.68
```

14) Advertise RIP and RIPng routes to Area 1 and Area 2. Assign an initial cost of 15 and make sure that routers adjust the cost to reflect the internal OSPF cost. ABRs in Area 1 and Area 2 should receive the more specific external routes.

We configure a policy where we match all RIP routes. We set the metric to 15, using 'then metric'. We ensure that the metric will be incremented further on by redistributing the RIP routes as an external type 1 route.

On ex3_ex4:

```
root@ex3_ex4# set policy-options policy-statement rip.to.ospf term 1
from protocol rip

root@ex3_ex4# set policy-options policy-statement rip.to.ospf term 1
then metric 15

root@ex3_ex4# set policy-options policy-statement rip.to.ospf term 1
then external type 1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set policy-options policy-statement rip.to.ospf term 1
then accept

root@ex3_ex4# set policy-options policy-statement ripng.to.ospf3
term 1 from protocol ripng

root@ex3_ex4# set policy-options policy-statement ripng.to.ospf3
term 1 then metric 15

root@ex3_ex4# set policy-options policy-statement ripng.to.ospf3
term 1 then external type 1

root@ex3_ex4# set policy-options policy-statement ripng.to.ospf3
term 1 then accept


root@ex3_ex4# set protocols ospf export rip.to.ospf

root@ex3_ex4# set protocols ospf3 export ripng.to.ospf3
```

Verification:

```
root@ex1# run show route 172.31.0.0/20

172.31.0.0/20      *[OSPF/150] 00:09:06, metric 17, tag 0
                    > to 172.30.0.74 via ge-0/0/14.56
172.31.0.0/24      *[OSPF/150] 00:37:22, metric 16, tag 0
                    > to 172.30.0.130 via ge-0/0/14.57
172.31.1.0/24      *[OSPF/150] 00:37:22, metric 16, tag 0
                    > to 172.30.0.130 via ge-0/0/14.57
172.31.2.0/24      *[OSPF/150] 00:37:22, metric 16, tag 0
                    > to 172.30.0.130 via ge-0/0/14.57
172.31.3.0/24      *[OSPF/150] 00:37:22, metric 16, tag 0
                    > to 172.30.0.130 via ge-0/0/14.57
172.31.4.0/24      *[OSPF/150] 00:37:22, metric 16, tag 0
                    > to 172.30.0.130 via ge-0/0/14.57
172.31.5.0/24      *[OSPF/150] 00:37:22, metric 16, tag 0
                    > to 172.30.0.130 via ge-0/0/14.57
```

We will configure **SRX3a** in the next task.

15) In this task we redistribute a RIP and RIPng summary route to OSPF on SRX3a. This summary route should be advertised to backbone area on ex1 but it should be restricted to backbone area on srx3.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

First, we create and redistribute an aggregate route for IPv4 as well as IPv6. The aggregate routes represent a summary route for RIP prefixes:

srx3:

```
root@srx3# set routing-instances SRX3a routing-options aggregate
route 172.31.0.0/20

root@srx3# set routing-instances SRX3a routing-options rib
SRX3a.inet6.0 aggregate route fd17:f0f4:f691:f::/64


root@srx3# set policy-options policy-statement rip.to.ospf term 1
from protocol aggregate

root@srx3# set policy-options policy-statement rip.to.ospf term 1
from route-filter 172.31.0.0/20 exact

root@srx3# set policy-options policy-statement rip.to.ospf term 1
then metric 15

root@srx3# set policy-options policy-statement rip.to.ospf term 1
then external type 1

root@srx3# set policy-options policy-statement rip.to.ospf term 1
then accept


root@srx3# set policy-options policy-statement ripng.to.ospf3 term 1
from protocol aggregate

root@srx3# set policy-options policy-statement ripng.to.ospf3 term 1
from route-filter fd17:f0f4:f691:f::/64 exact

root@srx3# set policy-options policy-statement ripng.to.ospf3 term 1
then metric 15

root@srx3# set policy-options policy-statement ripng.to.ospf3 term 1
then external type 1

root@srx3# set policy-options policy-statement ripng.to.ospf3 term 1
then accept


root@srx3# set routing-instances SRX3a protocols ospf export
rip.to.ospf
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set routing-instances SRX3a protocols ospf3 export
ripng.to.ospf3
```

To summarize the route on ex1:

ex1:

```
root@ex1# set protocols ospf area 0.0.0.2 nssa area-range
172.31.0.0/20

root@ex1# set protocols ospf3 area 0.0.0.2 nssa area-range
fd17:f0f4:f691:f::/64
```

To restrict the route on srx3:

srx3:

```
root@srx3# set protocols ospf area 0.0.0.2 nssa area-range
172.31.0.0/20 restrict

root@srx3# set protocols ospf3 area 0.0.0.2 nssa area-range
fd17:f0f4:f691:f::/64 restrict
```

Verification:

```
root@srx1# run show route 172.31.0.0/20


172.31.0.0/20      *[OSPF/150] 00:39:16, metric 17, tag 0
                    > to 172.30.0.10 via ge-0/0/4.15

root@srx1# run show route fd17:f0f4:f691:f::/64

fd17:f0f4:f691:f::/64
                   *[OSPF3/150] 00:40:42, metric 17, tag 0
                    > to fe80::3e94:d500:f46:7311 via ge-0/0/4.15
```

16) Ensure that all OSPF adjacencies are in the Full state and connectivity is provided between all of the routers' loopback interfaces. No routing loops may be present in the network:

OSPF adjacencies were verified earlier on as we were making our way through the different tasks.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

The last thing we need to do is redistribute OSPFv2 into the OSPFv3 IPv4 realm and vice versa. To achieve this, we configure the following:

On srx4 and ex2

```
root@srx4# set policy-options policy-statement ipv4.from.ospf term 1
from protocol ospf

root@srx4# set policy-options policy-statement ipv4.from.ospf term 1
then accept


root@srx4# set policy-options policy-statement ipv4.from.ospf3 term
1 from protocol ospf3

root@srx4# set policy-options policy-statement ipv4.from.ospf3 term
1 then accept


root@srx4# set protocols ospf export ipv4.from.ospf3

root@srx4# set protocols ospf3 realm ipv4-unicast export
ipv4.from.ospf
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 4: BGP Configuration

1) Additional interfaces to the external autonomous systems are configured (VRRP is configured for P1-facing interface):

srx1:

```
root@srx1# set interfaces ge-0/0/4 unit 301 description "L3 connection
to P1"

root@srx1# set interfaces ge-0/0/4 unit 301 vlan-id 301

root@srx1# set interfaces ge-0/0/4 unit 301 family inet address
172.16.31.1/24 vrrp-group 1 virtual-address 172.16.31.200

root@srx1# set interfaces ge-0/0/4 unit 301 family inet address
172.16.31.1/24 vrrp-group 1 priority 200

root@srx1# set interfaces ge-0/0/4 unit 301 family inet address
172.16.31.1/24 vrrp-group 1 accept-data

root@srx1# set interfaces ge-0/0/4 unit 400 description "L3 connection
to ISP1"

root@srx1# set interfaces ge-0/0/4 unit 400 vlan-id 400

root@srx1# set interfaces ge-0/0/4 unit 400 family inet address
172.16.40.2/30

root@srx1# set interfaces ge-0/0/4 unit 400 family inet6 address
2001:0:0:400::2/126

root@srx1# set interfaces ge-0/0/4 unit 401 description "L3 connection
to P2"

root@srx1# set interfaces ge-0/0/4 unit 401 vlan-id 401

root@srx1# set interfaces ge-0/0/4 unit 401 family inet address
172.16.41.1/30

root@srx1# set interfaces ge-0/0/4 unit 401 family inet6 address
::172.16.41.1/96
```

srx2:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set interfaces ge-0/0/4 unit 201 description "L3 connection
to ISP2 No.1"

root@srx2# set interfaces ge-0/0/4 unit 201 vlan-id 201

root@srx2# set interfaces ge-0/0/4 unit 201 family inet address
172.16.21.2/30

root@srx2# set interfaces ge-0/0/4 unit 201 family inet6 address
2001:aaaa:0:201::2/126

root@srx2# set interfaces ge-0/0/4 unit 202 description "L3 connection
to ISP2 No.2"

root@srx2# set interfaces ge-0/0/4 unit 202 vlan-id 202

root@srx2# set interfaces ge-0/0/4 unit 202 family inet address
172.16.22.2/30

root@srx2# set interfaces ge-0/0/4 unit 202 family inet6 address
2001:bbbb:0:202::2/126

root@srx2# set interfaces ge-0/0/4 unit 301 description "L3 connection
to P1"

root@srx2# set interfaces ge-0/0/4 unit 301 vlan-id 301

root@srx2# set interfaces ge-0/0/4 unit 301 family inet address
172.16.31.2/24 vrrp-group 1 virtual-address 172.16.31.200

root@srx2# set interfaces ge-0/0/4 unit 301 family inet address
172.16.31.2/24 vrrp-group 1 accept-data
```

2) We specify the AS on every device. Following that, we configure an IPv4 and an IPv6 group called 'internal4' and 'internal6' respectively. To make sure the loss of a single interface cannot impact any iBGP session, we source all sessions from the loopback interface:

On the all devices:

```
root@srx1# set routing-options autonomous-system 9876


root@srx1# set protocols bgp group internal4 type internal

root@srx1# set protocols bgp group internal4 local-address 172.30.30.1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set protocols bgp group internal4 neighbor 172.30.30.2
root@srx1# set protocols bgp group internal4 neighbor 172.30.30.3
root@srx1# set protocols bgp group internal4 neighbor 172.30.30.4
root@srx1# set protocols bgp group internal4 neighbor 172.30.30.5
root@srx1# set protocols bgp group internal4 neighbor 172.30.30.6
root@srx1# set protocols bgp group internal4 neighbor 172.30.30.7
root@srx1# set protocols bgp group internal4 neighbor 172.30.30.8
root@srx1# set protocols bgp group internal4 neighbor 172.30.30.9


root@srx1# set protocols bgp group internal6 type internal
root@srx1# set protocols bgp group internal6 local-address
fd17:f0f4:f691:30::1
root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:30::2
root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:30::3
root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:30::4
root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:30::5
root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:30::6
root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:30::7
root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:30::8
root@srx1# set protocols bgp group internal6 neighbor
fd17:f0f4:f691:30::9
```

Verification:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# run show bgp summary
Groups: 2 Peers: 16 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.0
                     493        483          0          0          0
0
inet6.0
                     109        104          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Accepted/Damped...
172.30.30.2            9876        430        452       0       0
16:26 Establ
  inet.0: 6/6/6/0
172.30.30.3            9876         31        422       0       0
14:00 Establ
  inet.0: 0/0/0/0
172.30.30.4            9876         29        419       0       0
12:58 Establ
  inet.0: 0/0/0/0
172.30.30.5            9876         34        424       0       0
14:57 Establ
  inet.0: 0/0/0/0
172.30.30.6            9876         30        420       0       0
13:12 Establ
  inet.0: 0/0/0/0
172.30.30.7            9876         16        407       0       0
7:21 Establ
  inet.0: 0/0/0/0
172.30.30.8            9876         32        422       0       0
14:04 Establ
  inet.0: 0/1/1/0
172.30.30.9            9876         28        419       0       0
12:45 Establ
  inet.0: 0/0/0/0
fd17:f0f4:f691:30::2        9876        117        115       0       0
15:46 Establ
  inet6.0: 4/4/4/0
fd17:f0f4:f691:30::3        9876         30        109       0       0
13:28 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:30::4        9876         28        108       0       0
12:57 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:30::5        9876         36        113       0       0
15:09 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:30::6        9876         29        109       0       0
13:09 Establ
  inet6.0: 0/0/0/0
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
fd17:f0f4:f691:30::7          9876          17          97          0          0
7:45 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:30::8          9876          32         110          0          0
13:48 Establ
  inet6.0: 0/1/1/0
fd17:f0f4:f691:30::9          9876          29         108          0          0
12:46 Establ
  inet6.0: 0/0/0/0
```

3) To make sure that BGP sessions use source port 179 from the ex3_ex4 point of view, we configure BGP with the **passive** configuration command on ex3_ex4:

Normally, when two devices are configured to engage in a peering relation, both peers try to set up a peering session. When we configure ex3_ex4 with the **passive** configuration command, we tell the device to wait for open messages send from the peer. The peers will set up the BGP session to destination port 179, which makes it the source port from the ex3_ex4 point of view.

By configuring the passive option directly under [**protocols bgp**], it applies to all configured sessions.

```
root@ex3_ex4# set protocols bgp passive
```

Verification:

```
root@ex3_ex4# run show bgp neighbor | match 179
Peer: 172.30.30.1+58881 AS 9876 Local: 172.30.30.7+179 AS 9876
Peer: 172.30.30.2+49374 AS 9876 Local: 172.30.30.7+179 AS 9876
Peer: 172.30.30.3+58375 AS 9876 Local: 172.30.30.7+179 AS 9876
Peer: 172.30.30.4+54125 AS 9876 Local: 172.30.30.7+179 AS 9876
Peer: 172.30.30.5+52269 AS 9876 Local: 172.30.30.7+179 AS 9876
Peer: 172.30.30.6+53553 AS 9876 Local: 172.30.30.7+179 AS 9876
Peer: 172.30.30.8+63709 AS 9876 Local: 172.30.30.7+179 AS 9876
Peer: 172.30.30.9+55213 AS 9876 Local: 172.30.30.7+179 AS 9876
Peer: fd17:f0f4:f691:30::1+60294 AS 9876 Local:
fd17:f0f4:f691:30::7+179 AS 9876
Peer: fd17:f0f4:f691:30::2+64487 AS 9876 Local:
fd17:f0f4:f691:30::7+179 AS 9876
Peer: fd17:f0f4:f691:30::3+54118 AS 9876 Local:
fd17:f0f4:f691:30::7+179 AS 9876
Peer: fd17:f0f4:f691:30::4+59158 AS 9876 Local:
fd17:f0f4:f691:30::7+179 AS 9876
Peer: fd17:f0f4:f691:30::5+64866 AS 9876 Local:
fd17:f0f4:f691:30::7+179 AS 9876
Peer: fd17:f0f4:f691:30::6+63289 AS 9876 Local:
fd17:f0f4:f691:30::7+179 AS 9876
Peer: fd17:f0f4:f691:30::8+60721 AS 9876 Local:
fd17:f0f4:f691:30::7+179 AS 9876
Peer: fd17:f0f4:f691:30::9+60769 AS 9876 Local:
fd17:f0f4:f691:30::7+179 AS 9876
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

4) Configure BFD liveness detection for the IPv4 BGP sessions:

```
root@srx1# set protocols bgp group internal4 bfd-liveness-detection
minimum-interval 600

root@srx1# set protocols bgp group internal4 bfd-liveness-detection
multiplier 3
```

Verification:

```
root@srx1# run show bfd session
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval
Multiplier
172.30.30.2              Up                       1.800     0.600
3
172.30.30.3              Up                       1.800     0.600
3
172.30.30.4              Up                       1.800     0.600
3
172.30.30.5              Up                       1.800     0.600
3
172.30.30.6              Up                       1.800     0.600
3
172.30.30.7              Up                       1.800     0.600
3
172.30.30.8              Up                       1.800     0.600
3
172.30.30.9              Up                       1.800     0.600
3
```

5) Exporting the aggregate routes that serve as RIP/RIPng summary routes to iBGP on SRX3a with community 9876:100:

```
root@srx3# set policy-options community rip_ripng members 9876:100

root@srx3# set policy-options policy-statement rip.to.bgp term 1 from
protocol aggregate

root@srx3# set policy-options policy-statement rip.to.bgp term 1 from
route-filter 172.31.0.0/20 exact
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx3# set policy-options policy-statement rip.to.bgp term 1 then
community add rip_ripng

root@srx3# set policy-options policy-statement rip.to.bgp term 1 then
accept


root@srx3# set policy-options policy-statement ripng.to.bgp term 1
from protocol aggregate

root@srx3# set policy-options policy-statement ripng.to.bgp term 1
from route-filter fd17:f0f4:f691:f::/64 exact

root@srx3# set policy-options policy-statement ripng.to.bgp term 1
then community add rip_ripng

root@srx3# set policy-options policy-statement ripng.to.bgp term 1
then accept


root@srx3# set routing-instances SRX3a protocols bgp group internal4
export rip.to.bgp

root@srx3# set routing-instances SRX3a protocols bgp group internal6
export ripng.to.bgp
```

Verification:

```
root@srx1# run show route 172.31/20

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/20       *[OSPF/150] 00:45:13, metric 17, tag 0
                     > to 172.30.0.10 via ge-0/0/4.15
                     [BGP/170] 00:51:18, localpref 100, from
172.30.30.8
                       AS path: I
                     > to 172.30.0.2 via ge-0/0/4.12
                       to 172.30.0.10 via ge-0/0/4.15
```

6) Configure srx1 and srx2 to peer with P1. Discover the IP address that P1 is using to try and set up a BGP session:

On srx1:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set protocols bgp group P1 neighbor 172.16.31.50 peer-as
65412
```

To ascertain the IP address that P1 is using, we use the **monitor traffic** command. With this command, we are able to see all RE generated traffic and all traffic that is directed to the RE. We issue the command an monitor interface ge-0/0/4.301 and wait for a BGP message that is send by P1:

```
root@srx1# run monitor traffic interface ge-0/0/4.301 no-resolve
verbose output suppressed, use <detail> or <extensive> for full
protocol decode
Address resolution is OFF.
Listening on ge-0/0/4.301, capture size 96 bytes

10:00:00.802163 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:01.749177 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:02.530078 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:03.420863 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:04.392010 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:05.149497 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:05.932359 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:06.719207 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:07.676414 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:08.534961 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:09.521186 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:10.400676 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:10.834032 Out IP truncated-ip - 4 bytes missing! 172.16.31.1 >
224.0.0.13: 172.16.31.1 > 224.0.0.13:PIMv2, Hello, length 34
10:00:11.300442 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:12.266714 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:13.090853 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:13.927992 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:14.856882 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:15.689086 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
20: vrid=1 prio=200 authtype=none intvl=1
10:00:16.504026 Out IP 172.16.31.1 > 224.0.0.18: VRRPv2-advertisement
20: vrid=1 prio=200 authtype=none intvl=1
10:00:17.205028  In IP 172.16.31.150.49626 > 172.16.31.200.179: S
3272866396:3272866396(0) win 16384 <mss 1460,nop,wscale
0,nop,nop,timestamp 4812080 0,sackOK,eol>
10:00:17.205157 Out IP truncated-ip - 14 bytes missing!
172.16.31.200.179 > 172.16.31.150.49626: S 3267961558:3267961558(0)
ack 3272866397 win 16384 <mss 1460,nop,wscale 0,nop,nop,[|tcp]>
10:00:17.207969  In IP 172.16.31.150.49626 > 172.16.31.200.179: . ack
1 win 17376 <nop,nop,timestamp 4812083 4800345>
10:00:17.208669 Out IP truncated-ip - 2 bytes missing!
172.16.31.200.179 > 172.16.31.150.49626: F 1:1(0) ack 1 win 17376
<nop,nop,timestamp 4800349[|tcp]>
10:00:17.208900  In IP 172.16.31.150.49626 > 172.16.31.200.179: P
1:60(59) ack 1 win 17376 <nop,nop,timestamp 4812084 4800345>: BGP,
length: 59
10:00:17.208972 Out IP 172.16.31.200.179 > 172.16.31.150.49626: R
3267961559:3267961559(0) win 0
10:00:17.214550  In IP 172.16.31.150.49626 > 172.16.31.200.179: . ack
2 win 17375 <nop,nop,timestamp 4812088 4800349>
```

From the output, we can see that P1 is sourcing the BGP session from 172.16.31.150. We proceed and configure this as the address to set up peering with P1 on both srx1 and srx2:

```
root@srx1# rename protocols bgp group P1 neighbor 172.16.31.50 to
neighbor 172.16.31.150
```

Verification:

```
root@srx1# run show bgp summary
Groups: 6 Peers: 20 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.0
                     493        483          0          0          0
0
inet6.0
                     109        104          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Accepted/Damped...
172.16.31.150          65412          6          6       0       0
54 Establ
  inet.0: 8/8/8/0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

7) In this task we configure a single EBGP session that supports IPv4 as well as IPv6 to P2. P2 was configured with a wrong peer-as number, we need to get things working without touching P2:

To complete this task, we use the **local-as** knob under the [**protocols bgp group P2**] stanza. By using this knob, we specify that the BGP session should simulate the AS number we specify. In this case, we specify AS 15.6789. We also configure the '**local-as private**' statement. By doing so, we can hide the configured **local-as** (AS 15.6789) in routing-updates learned from the peer.

By default, the inet family is enabled for BGP sessions. However, when another address family is specified, this default no longer applies. Since the session should be used for inet6, we make sure that we configure both the inet as well as the inet6 address family:

```
root@srx1# set protocols bgp group P2 family inet unicast

root@srx1# set protocols bgp group P2 family inet6 unicast

root@srx1# set protocols bgp group P2 local-as 15.6789

root@srx1# set protocols bgp group P2 local-as private

root@srx1# set protocols bgp group P2 neighbor 172.16.41.2 peer-as
365.18
```

Upon verification, we should see 16 IPv4 and IPv6 routes received from the neighbor.

Verification:

```
root@srx1# run show route receive-protocol bgp 172.16.41.2

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 195.24.32.0/24          172.16.41.2
365.18 I
* 195.25.33.0/24          172.16.41.2
365.18 I
* 195.26.34.0/24          172.16.41.2
365.18 I
* 195.27.35.0/24          172.16.41.2
365.18 I
* 195.28.36.0/24          172.16.41.2
365.18 I
* 195.29.37.0/24          172.16.41.2
365.18 I
* 195.30.38.0/24          172.16.41.2
365.18 I
* 195.31.39.0/24          172.16.41.2
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
365.18 I
* 195.32.40.0/24          172.16.41.2
365.18 I
* 195.33.41.0/24          172.16.41.2
365.18 I
* 195.34.42.0/24          172.16.41.2
365.18 I
* 195.35.43.0/24          172.16.41.2
365.18 I
* 195.36.44.0/24          172.16.41.2
365.18 I
* 195.36.45.0/24          172.16.41.2
365.18 I
* 195.36.46.0/24          172.16.41.2
365.18 I
* 195.36.47.0/24          172.16.41.2
365.18 I

inet6.0: 151 destinations, 162 routes (151 active, 0 holddown, 4
hidden)
```

As we can see, not a single IPv6 route is received. We can troubleshoot the issue by configuring traceoptions for the BGP protocol. We specify the **update detail** flag and examine the information. From the logs, we can learn that we are receiving the routes but they are ignored.

In the logs, we can see that the next-hop is listed as 'unexpectedly remote', so we add the **accept-remote-nexthop** option to the configuration of the BGP group:

```
root@srx1# set protocols bgp group P2 accept-remote-nexthop
```

Verification:

```
root@srx1# run show route receive-protocol bgp 172.16.41.2

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 195.24.32.0/24          172.16.41.2
365.18 I
* 195.25.33.0/24          172.16.41.2
365.18 I
* 195.26.34.0/24          172.16.41.2
365.18 I
* 195.27.35.0/24          172.16.41.2
365.18 I
* 195.28.36.0/24          172.16.41.2
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
365.18 I
* 195.29.37.0/24          172.16.41.2
365.18 I
* 195.30.38.0/24          172.16.41.2
365.18 I
* 195.31.39.0/24          172.16.41.2
365.18 I
* 195.32.40.0/24          172.16.41.2
365.18 I
* 195.33.41.0/24          172.16.41.2
365.18 I
* 195.34.42.0/24          172.16.41.2
365.18 I
* 195.35.43.0/24          172.16.41.2
365.18 I
* 195.36.44.0/24          172.16.41.2
365.18 I
* 195.36.45.0/24          172.16.41.2
365.18 I
* 195.36.46.0/24          172.16.41.2
365.18 I
* 195.36.47.0/24          172.16.41.2
365.18 I

inet6.0: 167 destinations, 178 routes (167 active, 0 holddown, 4
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
  fd17:f0f4:f691:d::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:1::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:2::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:3::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:4::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:5::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:6::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:7::/80
*                         ::ffff:172.16.41.2
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
365.18 I
  fd17:f0f4:f691:d:8::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:9::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:a::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:b::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:c::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:d::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:e::/80
*                         ::ffff:172.16.41.2
365.18 I
  fd17:f0f4:f691:d:f::/80
*                         ::ffff:172.16.41.2
365.18 I
```

8) We configure two EBGP sessions with ISP1. A session for IPv4 and a session for IPv6. Both are configured with MD5 authentication:

```
root@srx1# set protocols bgp group ISP1v4 authentication-key jncie

root@srx1# set protocols bgp group ISP1v4 neighbor 172.16.40.1 peer-as
456



root@srx1# set protocols bgp group ISP1v6 authentication-key jncie

root@srx1# set protocols bgp group ISP1v6 neighbor 2001:0:0:400::1
peer-as 456
```

Verification:

```
root@srx1# run show bgp summary
Groups: 6 Peers: 20 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Pending
inet.0
                     493        483          0          0          0
0
inet6.0
                     109        104          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Accepted/Damped...
172.16.31.150         65412         38         39       0       0
15:54 Establ
  inet.0: 8/8/8/0
172.16.40.1             456        459         49       0       0
16:53 Establ
  inet.0: 453/461/453/0
172.16.41.2          365.18         43        546       0       0
16:57 Establ
  inet.0: 16/17/16/0
  inet6.0: 16/16/16/0
172.30.30.2            9876        430        452       0       0
16:26 Establ
  inet.0: 6/6/6/0
172.30.30.3            9876         31        422       0       0
14:00 Establ
  inet.0: 0/0/0/0
172.30.30.4            9876         29        419       0       0
12:58 Establ
  inet.0: 0/0/0/0
172.30.30.5            9876         34        424       0       0
14:57 Establ
  inet.0: 0/0/0/0
172.30.30.6            9876         30        420       0       0
13:12 Establ
  inet.0: 0/0/0/0
172.30.30.7            9876         16        407       0       0
7:21 Establ
  inet.0: 0/0/0/0
172.30.30.8            9876         32        422       0       0
14:04 Establ
  inet.0: 0/1/1/0
172.30.30.9            9876         28        419       0       0
12:45 Establ
  inet.0: 0/0/0/0
2001:0:0:400::1         456        121         47       0       0
17:00 Establ
  inet6.0: 84/88/84/0
fd17:f0f4:f691:30::2        9876         117        115       0        0
15:46 Establ
  inet6.0: 4/4/4/0
fd17:f0f4:f691:30::3        9876          30        109       0        0
13:28 Establ
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  inet6.0: 0/0/0/0
fd17:f0f4:f691:30::4        9876          28         108        0          0
12:57 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:30::5        9876          36         113        0          0
15:09 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:30::6        9876          29         109        0          0
13:09 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:30::7        9876          17          97        0          0
7:45 Establ
  inet6.0: 0/0/0/0
fd17:f0f4:f691:30::8        9876          32         110        0          0
13:48 Establ
  inet6.0: 0/1/1/0
fd17:f0f4:f691:30::9        9876          29         108        0          0
12:46 Establ
  inet6.0: 0/0/0/0

root@srx1# run show route receive-protocol bgp 172.16.40.1

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 0.0.0.0/4               172.16.40.1                             456
1620 61671 I
* 0.0.0.0/5               172.16.40.1                             456
1620 61671 27075 I
* 1.64.0.0/10             172.16.40.1                             456
1620 61671 I
* 1.84.160.0/20           172.16.40.1                             456
1620 33112 I
* 1.96.0.0/11             172.16.40.1                             456
1620 33112 63164 40776 51777 I
* 1.161.192.0/21          172.16.40.1                             456
1620 33112 30404 32138 45045 I
```

9) In this task we configure two EBGP sessions with ISP2. Both should be enabled for the exchange of IPv4 as well as IPv6 address information. We need to enable load-balancing over both links with ISP2:

First, we enable the forwarding table for load balancing. We create a load-balancing policy:

```
root@srx2# set policy-options policy-statement load-balancing-policy
term 1 then load-balance per-packet
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

And we apply that policy to the forwarding table:

```
root@srx2# set routing-options forwarding-table export load-balancing-
policy
```

Then, we configure the EBGP sessions. To enable them for load-sharing, we configure the BGP group with the multipath knob:

```
root@srx2# set protocols bgp group ISP2v4 peer-as 6688
root@srx2# set protocols bgp group ISP2v4 multipath
root@srx2# set protocols bgp group ISP2v4 neighbor 172.16.21.1
root@srx2# set protocols bgp group ISP2v4 neighbor 172.16.22.1
root@srx2# set protocols bgp group ISP2v6 peer-as 6688
root@srx2# set protocols bgp group ISP2v6 multipath
root@srx2# set protocols bgp group ISP2v6 neighbor 2001:aaaa:0:201::1
root@srx2# set protocols bgp group ISP2v6 neighbor 2001:bbbb:0:202::1
```

Verification:

```
root@srx2# run show bgp summary
172.16.21.1            6688        459        456       0       0
17:12 Establ
  inet.0: 6/461/453/0
172.16.22.1            6688        433        456       0       0
17:09 Establ
2001:aaaa:0:201::1        6688        124        122       0       0
17:13 Establ
  inet6.0: 4/88/84/0
2001:bbbb:0:202::1        6688        120        122       0       0
17:09 Establ
  inet6.0: 0/0/0/0

root@srx2# run show route forwarding-table | match 172.16.21.1
18.0.0.0/7         user     0 172.16.21.1        ucst   628     6 ge-
0/0/4.201
172.16.21.1/32     dest     1 0:1f:12:2e:6f:81   ucst   628     6 ge-
0/0/4.201
172.160.0.0/12     user     0 172.16.21.1        ucst   628     6 ge-
0/0/4.201
177.180.0.0/14     user     0 172.16.21.1        ucst   628     6 ge-
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
0/0/4.201

root@srx2# run show route forwarding-table | match 172.16.22.1
28.4.0.0/15        user     0 172.16.22.1        ucst   704     6 ge-
0/0/4.202
137.160.175.128/28 user     0 172.16.22.1        ucst   704     6 ge-
0/0/4.202
142.180.0.0/15     user     0 172.16.22.1        ucst   704     6 ge-
0/0/4.202
172.16.22.1/32     dest     1 0:1f:12:2e:6f:81   ucst   704     6 ge-
0/0/4.202
```

10) We configure an RPM probe called P1 with the relevant characteristics:

```
root@srx2# set services rpm probe P1 test P1-icmp target address
172.16.31.200

root@srx2# set services rpm probe P1 test P1-icmp probe-count 3

root@srx2# set services rpm probe P1 test P1-icmp probe-interval 1

root@srx2# set services rpm probe P1 test P1-icmp test-interval 60
```

Verification:

```
root@srx2# run show services rpm probe-results
    Owner: P1, Test: P1-icmp
    Target address: 172.16.31.200, Probe type: icmp-ping, Test size: 3
probes
    Probe results:
      Response received, Fri Jan  2 13:51:54 2015, No hardware
timestamps
      Rtt: 1102 usec
    Results over current test:
      Probes sent: 3, Probes received: 3, Loss percentage: 0
      Measurement: Round trip time
        Samples: 3, Minimum: 1102 usec, Maximum: 1287 usec, Average:
1201 usec,
        Peak to peak: 185 usec, Stddev: 76 usec, Sum: 3602 usec
    Results over last test:
      Probes sent: 3, Probes received: 3, Loss percentage: 0
      Test completed on Fri Jan  2 13:51:54 2015
      Measurement: Round trip time
        Samples: 3, Minimum: 1102 usec, Maximum: 1287 usec, Average:
1201 usec,
        Peak to peak: 185 usec, Stddev: 76 usec, Sum: 3602 usec
    Results over all tests:
      Probes sent: 891, Probes received: 885, Loss percentage: 0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
      Measurement: Round trip time
        Samples: 885, Minimum: 160 usec, Maximum: 32023 usec,
        Average: 1836 usec, Peak to peak: 31863 usec, Stddev: 2178
usec,
        Sum: 1625099 usec

[edit]

root@srx2# run show services rpm history-results
    Owner, Test                 Probe received              Round trip
time
    P1, P1-icmp                 Fri Jan  2 13:35:20 2015
1206 usec
    P1, P1-icmp                 Fri Jan  2 13:35:21 2015
1130 usec
    P1, P1-icmp                 Fri Jan  2 13:36:21 2015
1027 usec
```

11) Advertise a RIP summary route from iBGP to EBGP:

When an IGP route supersedes a BGP learned route, the Junos OS will not advertise the BGP route. The IGP route will be the active route and the BGP route will be inactive. By default, inactive routes are not readvertised.

To accomplish our goal, we turn to the '**advertise-inactive**' knob. What this command does is it allows you to export the best BGP learned route, even though that BGP route is not active in the routing-table.

On srx1 and srx2:

```
root@srx1# set protocols bgp advertise-inactive
```

Verification:

```
root@srx1# run show route 172.31/20

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/20      *[OSPF/150] 04:56:53, metric 17, tag 0
                    > to 172.30.0.10 via ge-0/0/4.15
                    [BGP/170] 00:02:53, localpref 100, from
172.30.30.8
                      AS path: I
                    > to 172.30.0.2 via ge-0/0/4.12
                      to 172.30.0.10 via ge-0/0/4.15
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit]
root@srx1# run show route advertising-protocol bgp 172.16.41.2
172.31/20

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
  172.31.0.0/20           Self                                    9876
I
```

12) To make sure srx1 serves as the exit point for traffic to ISP1 and P1, we use the local-preference attribute:

On srx1, we configure import policies that alter the local-preference for all routes learned from the P1 and ISP1 neighbor. By setting the local-preference to a higher value than the default of 100, we ensure that the routes learned on srx1 will be the preferred routes in our AS:

```
root@srx1# set policy-options policy-statement isp1.in term
local.preference then local-preference 120


root@srx1# set policy-options policy-statement isp1.in-v6 term
local.preference then local-preference 120


root@srx1# set policy-options policy-statement p1.in term local-
preference then local-preference 150


root@srx1# set protocols bgp group P1 import p1.in

root@srx1# set protocols bgp group ISP1v4 import isp1.in

root@srx1# set protocols bgp group ISP1v6 import isp1.in-v6
```

13) Routes advertised from P1 and P2 should be preferred over the same routes received from ISP1 and ISP2:

To complete this task, we consider a local-preference higher than the 120 that we used for ISP1. We set the local-preference for P2 to a value of 150, which is equal to local preference that was used for P1 learned routes:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set policy-options policy-statement p2.in term local-
preference then local-preference 150

root@srx1# set protocols bgp group P2 import p2.in
```

Verification:

```
root@srx4# run show route 172.31.64.0/21

inet.0: 526 destinations, 551 routes (526 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

172.31.64.0/24     *[BGP/170] 00:07:42, localpref 150, from
172.30.30.1
                      AS path: 65412 I
                    > to 172.30.0.13 via ge-0/0/4.24
172.31.65.0/24     *[BGP/170] 00:07:42, localpref 150, from
172.30.30.1
                      AS path: 65412 I
                    > to 172.30.0.13 via ge-0/0/4.24
172.31.66.0/24     *[BGP/170] 00:07:42, localpref 150, from
172.30.30.1
                      AS path: 65412 I
```

14) In this task, we export a summary route of our AS internal IPv4 and IPv6 addresses to all EBGP except P1:

We do this through the creation of a policy that matches the relevant aggregate route and we apply the policy on export under the proper peer:

On srx1 and srx2:

```
root@srx1# set policy-options policy-statement isp1.out term
local.routes from protocol aggregate

root@srx1# set policy-options policy-statement isp1.out term
local.routes from route-filter 172.30.0.0/16 exact

root@srx1# set policy-options policy-statement isp1.out term
local.routes then accept


root@srx1# set policy-options policy-statement isp1.out-v6 term
local.ipv6.routes from protocol aggregate
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set policy-options policy-statement isp1.out-v6 term
local.ipv6.routes from route-filter fd17:f0f4:f691::/48 exact

root@srx1# set policy-options policy-statement isp1.out-v6 term
local.ipv6.routes then accept


root@srx1# set policy-options policy-statement p2.out term
local.routes from protocol aggregate

root@srx1# set policy-options policy-statement p2.out term
local.routes from route-filter 172.30.0.0/16 exact

root@srx1# set policy-options policy-statement p2.out term
local.routes then accept


root@srx1# set protocols bgp group ISP1v4 export isp1.out

root@srx1# set protocols bgp group ISP1v6 export isp1.out-v6

root@srx1# set protocols bgp group P2 export p2.out
```

Verification:

```
root@srx2# show route advertising-protocol bgp 172.16.22.1 172.30/16
exact

inet.0: 536 destinations, 1451 routes (534 active, 0 holddown, 16
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 172.30.0.0/16           Self                 100                9876
9876 9876 [9876] I

[edit]

root@srx2# show route advertising-protocol bgp 2001:aaaa:0:201::1
fd17:f0f4:f691::/48 exact

inet6.0: 168 destinations, 349 routes (168 active, 0 holddown, 8
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* fd17:f0f4:f691::/48     Self                                    9876
9876 986 [9876] {365.18} I
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

15) To P1, we should advertise only a default route and the routes local to our own AS:

To complete this task, we create a policy with two terms. In the first term, we match a static default route and in the second term we match routes local to our own AS.

We are able to match the routes local to our AS without any comprehensive route-filtering scheme. We can use an **as-path** that matches only routes with an AS-path that is 'NULL'. The AS number is prepended whenever a route is advertised across an EBGP peering session. Hence, a NULL AS-path means that the route has to be locally generated.

First, we create an as-path that matches only a NULL path:

```
root@srx1# set policy-options as-path null "()"
```

Then we configure and apply the policy:

```
root@srx1# set policy-options policy-statement p1.out term 1 from
protocol static

root@srx1# set policy-options policy-statement p1.out term 1 from
route-filter 0.0.0.0/0 exact

root@srx1# set policy-options policy-statement p1.out term 1 then
accept


root@srx1# set policy-options policy-statement p1.out term 2 from as-
path null

root@srx1# set policy-options policy-statement p1.out term 2 then next
policy

root@srx1# set policy-options policy-statement p1.out term 3 then
reject

root@srx1# set protocols bgp group P1 export p1.out
```

Verification:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# run show route advertising-protocol bgp 172.16.31.150

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 0.0.0.0/0               Self                                    I
  172.31.0.0/20           Self                                    I
```

16) We need to advertise P1 routes to all iBGP and eBGP peers. Additionally, we need to make sure that private AS numbers are not leaked outside your AS without using the **remove-private** knob:

The P1 routes are advertised to iBGP and eBGP neighbors by default. What we do need to realize is making sure that private AS numbers are not.

What we want is to advertise a summary route. Because an aggregate route would have a next-hop that is null, we turn to the use of a generate route. This way, the next-hop can be inherited from the contributing routes.

```
root@srx1# set routing-options generate route 172.31.64.0/21

root@srx1# set policy-options policy-statement isp1.out term p1.routes
from route-filter 172.31.64.0/21 exact
```

We then alter the import policy we configured for the P1 neighbor. We attach the **no-export** community to all of the routes we learn from the peer:

```
root@srx1# set policy-options community no.export members no-export

root@srx1# set policy-options policy-statement p1.in term no.export
then community add no.export
```

By tagging the routes with the well-known no-export community, we ensure that the routes are not advertised accross any of the eBGP sessions we have.

Verification:

```
root@srx1# run show route advertising-protocol bgp 172.16.40.1

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
path
* 18.0.0.0/7              Self                                    6688
54591 33112 36621 ?
* 28.4.0.0/15             Self                                    6688
54591 33003 ?
* 137.160.175.128/28      Self                                    6688
54591 33003 ?
* 142.180.0.0/15          Self                                    6688
54591 33112 22508 22395 534 51038 38631 I
* 172.30.0.0/16           Self                                    I
  172.31.0.0/20           Self                                    I
* 172.31.64.0/21          Self                                    I
```

17) Accept only routes originated from within the P2 AS and mark them with community 9876:200:

What we need to do here is configure an as-path filter in the [**policy as-path**] stanza. The as-path filter can be used to make sure we only match routes that originated from the P2 AS (AS365.18). We escape the special character so that we math 365.18 exactly:

```
root@srx1# set policy-options as-path AS365.18 "365\.18"
```

We configure the following community to be able to tag the routes:

```
root@srx1# set policy-options community AS365.18 members 9876:200
```

We already applied an import policy to this neighbor earlier on. So instead of creating a new one, we simply add to the already existing route-policy. We add three terms. One to add the community, a second to allow routes permitted by the as-path filter and a third term to reject the rest:

```
root@srx1# set policy-options policy-statement p2.in term
community.add then community add AS365.18

root@srx1# set policy-options policy-statement p2.in term only.p2.as
from as-path AS365.18

root@srx1# set policy-options policy-statement p2.in term only.p2.as
then next policy

root@srx1# set policy-options policy-statement p2.in term reject then
reject
```

Verification:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# run show route receive-protocol bgp 172.16.41.2

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 195.24.32.0/24          172.16.41.2
365.18 I
* 195.25.33.0/24          172.16.41.2
365.18 I
* 195.26.34.0/24          172.16.41.2
365.18 I
* 195.27.35.0/24          172.16.41.2
365.18 I
* 195.28.36.0/24          172.16.41.2
365.18 I
* 195.29.37.0/24          172.16.41.2
365.18 I
* 195.30.38.0/24          172.16.41.2
365.18 I
* 195.31.39.0/24          172.16.41.2
365.18 I
* 195.32.40.0/24          172.16.41.2
365.18 I
* 195.33.41.0/24          172.16.41.2
365.18 I
* 195.34.42.0/24          172.16.41.2
365.18 I
* 195.35.43.0/24          172.16.41.2
365.18 I

root@srx1# run show route receive-protocol bgp 172.16.41.2 hidden

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
  17.17.0.0/16            172.16.41.2
365.18 4444 5555 6666 I

inet6.0: 167 destinations, 178 routes (167 active, 0 holddown, 4
hidden)
```

18) Make sure that srx1 is the preferred entry point for traffic coming from the Internet:

To realize this our objective, we turn to the export policy on srx2. We create new terms for the policies we used on export and we prepend our own AS to the routes several times. By doing so, the peers will see the route through srx2 as a less attractive route when compared to the routes learned from srx1.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Also, we make sure that we move the newly created terms to the top of the policy using the **insert** command:

```
root@srx2# set policy-options policy-statement isp2.out term prepend
then as-path-prepend "9876 9876 9876"

root@srx2# set policy-options policy-statement isp2.out-v6 term
prepend then as-path-prepend "9876 9876 986"

root@srx2# insert policy-options policy-statement isp2.out term
prepend before term p1.routes

root@srx2# insert policy-options policy-statement isp2.out-v6 term
prepend before term local.ipv6.routes
```

Verification:

```
root@srx2# run show route advertising-protocol bgp 172.16.21.1

inet.0: 536 destinations, 1451 routes (534 active, 0 holddown, 16
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 0.0.0.0/4               Self                                    9876
9876 9876 [9876] 456 1620 61671 I
* 0.0.0.0/5               Self                                    9876
9876 9876 [9876] 456 1620 61671 27075 I
* 1.64.0.0/10             Self                                    9876
9876 9876 [9876] 456 1620 61671 I
* 1.84.160.0/20           Self                                    9876
9876 9876 [9876] 456 1620 33112 I
* 1.96.0.0/11             Self                                    9876
9876 9876 [9876] 456 1620 33112 63164 40776 51777 I
* 1.161.192.0/21          Self                                    9876
9876 9876 [9876] 456 1620 33112 30404 32138 45045 I
```

19) Inbound traffic from ISP2 to srx2 should be routes over interface ge-0/0/4.201:

The Multi Exit Discriminator, or MED, is an optional non-transitive attribute that can be used to influence the preferred path into the AS when there are multiple links between the two ASs.

When the MED is absent, the Junos OS treats it as though it was set to 0. To have the VR send traffic to our AS using the ge-0/0/4.201 interface, we increase the MED for the routes advertised across the session that uses the other interface.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

The MED can be set in a policy using the **'then metric'** statement. We can also configure the MED inside the [**protocols bgp group <*group-name*> neighbor <*neighbor-address*>**] stanza, like this:

```
root@srx2# set protocols bgp group ISP2v4 neighbor 172.16.22.1 metric-
out 100

root@srx2# set protocols bgp group ISP2v6 neighbor 2001:bbbb:0:202::1
metric-out 100
```

Verification:

```
root@srx2# run show route advertising-protocol bgp 172.16.21.1

inet.0: 536 destinations, 1451 routes (534 active, 0 holddown, 16
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 0.0.0.0/4               Self                                    9876
9876 9876 [9876] 456 1620 61671 I
* 0.0.0.0/5               Self                                    9876
9876 9876 [9876] 456 1620 61671 27075 I
* 1.64.0.0/10             Self                                    9876
9876 9876 [9876] 456 1620 61671 I
* 1.84.160.0/20           Self                                    9876
9876 9876 [9876] 456 1620 33112 I

root@srx2# run show route advertising-protocol bgp 172.16.22.1

inet.0: 536 destinations, 1451 routes (534 active, 0 holddown, 16
hidden)
  Prefix                  Nexthop              MED     Lclpref    AS
path
* 0.0.0.0/4               Self                 100                9876
9876 9876 [9876] 456 1620 61671 I
* 0.0.0.0/5               Self                 100                9876
9876 9876 [9876] 456 1620 61671 27075 I
* 1.64.0.0/10             Self                 100                9876
9876 9876 [9876] 456 1620 61671 I
```

20) Do not accept any of the RFC1918 mentioned routes from ISP1 or ISP2. Also reject community 456:1234 from ISP1 and the community 6688:1234 from ISP2:

Normally, a community is configured to match a route that is tagged with that community. If there is a match for a configured member, Junos OS returns a TRUE result. When there is no match, Junos OS returns a FALSE result.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

When we configure a community with the **invert-match** option, we turn things around. Junos OS will return a TRUE as long as there is no match. When there is a match, Junos OS returns a FALSE.

To make things easier, we can use the **invert-match** option here:

srx1:

First, we take care of the RFC1918 routes, for which we use a separate term. We prepare for any mask length by using the '**orlonger**' keyword:

```
root@srx1# set policy-options policy-statement isp1.in term rfc1918
from route-filter 10.0.0.0/8 orlonger reject

root@srx1# set policy-options policy-statement isp1.in term rfc1918
from route-filter 172.16.0.0/12 orlonger reject

root@srx1# set policy-options policy-statement isp1.in term rfc1918
from route-filter 192.168.0.0/16 orlonger reject
```

The following community will match as long as the 456:1234 community is absent:

```
root@srx1# set policy-options community except.isp1 invert-match

root@srx1# set policy-options community except.isp1 members 456:1234
```

The community is used in the IPv4 as well as the IPv6 policy:

```
root@srx1# set policy-options policy-statement isp1.in term
only.community.isp1 from community except.isp1

root@srx1# set policy-options policy-statement isp1.in term
only.community.isp1 then reject


root@srx1# set policy-options policy-statement isp1.in-v6 term
only.community.isp1 from community except.isp1

root@srx1# set policy-options policy-statement isp1.in-v6 term
only.community.isp1 then reject
```

srx2:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx2# set policy-options policy-statement isp2.in term rfc1918
from route-filter 10.0.0.0/8 orlonger reject

root@srx2# set policy-options policy-statement isp2.in term rfc1918
from route-filter 172.16.0.0/12 orlonger reject

root@srx2# set policy-options policy-statement isp2.in term rfc1918
from route-filter 192.168.0.0/16 orlonger reject


root@srx2# set policy-options community except.isp2 invert-match

root@srx2# set policy-options community except.isp2 members 6688:1234


root@srx2# set policy-options policy-statement isp2.in term
only.community.isp2 from community except.isp2

root@srx2# set policy-options policy-statement isp2.in term
only.community.isp2 then reject


root@srx2# set policy-options policy-statement isp2.in-v6 term
only.community.isp2 from community except.isp2

root@srx2# set policy-options policy-statement isp2.in-v6 term
only.community.isp2 then reject


root@srx2# set protocols bgp group ISP2v4 import isp2.in

root@srx2# set protocols bgp group ISP2v6 import isp2.in-v6
```

Verification:

To verify our policy, we check the hidden routes to see what prefixes have been filtered:

```
root@srx1# run show route hidden extensive

inet.0: 538 destinations, 545 routes (535 active, 0 holddown, 9
hidden)
10.0.0.0/10 (1 entry, 0 announced)
         BGP                 /-121
                Next hop type: Router, Next hop index: 606
                Address: 0x155c6e4
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                Next-hop reference count: 1367
                Source: 172.16.40.1
                Next hop: 172.16.40.1 via ge-0/0/4.400, selected
                State: <Hidden Ext>
                Local AS:  9876 Peer AS:   456
                Age: 28:09
                Task: BGP_456.172.16.40.1+179
                AS path: 456 1620 33112 60177 54491 19593 62822 59959
22793 22889 I
                Communities: 456:1234
                Localpref: 120
                Router ID: 172.16.40.1
```

21) To have all the devices display 4-byte AS in AS Dot format, we issue the following configuration command on all of the devices:

```
root@srx1# set routing-options autonomous-system asdot-notation
```

Verification:

```
root@srx4# run show route 195.36.46.0/24

inet.0: 526 destinations, 551 routes (526 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

195.36.46.0/24     *[BGP/170] 01:20:54, localpref 150, from
172.30.30.1
                      AS path: 365.18 I
                    > to 172.30.0.13 via ge-0/0/4.24

SRX4b.inet.0: 518 destinations, 519 routes (518 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

195.36.46.0/24     *[BGP/170] 01:20:41, localpref 150, from
172.30.30.1
                      AS path: 365.18 I
                      to 172.30.10.5 via ge-0/0/5.49
                    > to 172.30.10.1 via ge-0/0/5.69
```

22) You cannot have any blackholes or suboptimal routing:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

We need to make sure that all of our devices will be able to resolve the next-hop. We can easily achieve this by configuring the following policy:

```
root@srx1# set policy-options policy-statement internal.out term nhs
then next-hop self

root@srx1# set policy-options policy-statement internal.out-v6 term
nhs then next-hop self


root@srx1# set protocols bgp group internal4 export internal.out

root@srx1# set protocols bgp group internal6 export internal.out-v6
```

Verification:

```
root@ex1# run show route

inet.0: 543 destinations, 544 routes (543 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/4          *[BGP/170] 00:41:24, localpref 120, from
172.30.30.1
                     AS path: 456 1620 61671 I
                    > to 172.30.0.9 via ge-0/0/14.15
0.0.0.0/5          *[BGP/170] 00:41:24, localpref 120, from
172.30.30.1
                     AS path: 456 1620 61671 27075 I
                    > to 172.30.0.9 via ge-0/0/14.15
```

Further verification can be done using the '**show route resolution unresolved**' command.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 5: Multicast Routing and Switching

1) We add the new interfaces listed in table 31:

```
root@srx1# set interfaces ge-0/0/4 unit 40 description "L3 connection
to Src1"

root@srx1# set interfaces ge-0/0/4 unit 40 vlan-id 40

root@srx1# set interfaces ge-0/0/4 unit 40 family inet address
192.168.40.254/24
```

2) We inject the newly configured subnets into the IGP and we make sure that no adjacency will form across the link:

```
root@srx1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.40
passive
```

3) Configure PIM version 2 sparse mode on all devices. Enable PIM for both IPv4 and IPv6. Enable PIM on the interfaces to DC1 and make sure PIM is not enabled on the OoB interfaces:

The quickest way to achieve our objective is to enable PIM for all of the interfaces and then explicitly disable it for the OoB interface. The PIM mode and version as well as the **inet** and the **inet6** family can be configured under the [**protocols pim interface all**] stanza.

On all devices:

```
root@srx1# set protocols pim interface all family inet

root@srx1# set protocols pim interface all family inet6

root@srx1# set protocols pim interface all mode sparse

root@srx1# set protocols pim interface all version 2

root@srx1# set protocols pim interface ge-0/0/0.0 disable
```

Verification:

```
root@srx1# run show pim neighbors
Instance: PIM.master
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Interface           IP V Mode        Option      Uptime Neighbor addr
ge-0/0/4.12          4 2             HPLGT       00:10:48 172.30.0.2
ge-0/0/4.13          4 2             HPLGT       00:10:51 172.30.0.6
ge-0/0/4.15          4 2             HPLGT       00:10:51 172.30.0.10
ge-0/0/4.301         4 2             HPLGT       00:10:48 172.16.31.2
ge-0/0/4.12          6 2             HPLGT       00:10:44
fe80::7a19:f700:c18:3d84
ge-0/0/4.13          6 2             HPLGT       00:10:44
fe80::3e8a:b000:d5c:7aa4
ge-0/0/4.15          6 2             HPLGT       00:10:44
fe80::3e94:d500:f46:7311
vlan.700             6 2             HPLGT       00:20:47 fe80::3
```

4) Used MSDP to enable an Anycast RP configuration on srx1 and ex2. Configure the other devices with the static RP address (172.30.30.254):

srx1:

First, we configure the Anycast RP address on the loopback interface. We specify the currently active IP address as the primary address:

```
root@srx1# set interfaces lo0 unit 0 family inet address
172.30.30.1/32 primary

root@srx1# set interfaces lo0 unit 0 family inet address
172.30.30.254/32
```

We configure the device to use newly configured address as an RP address:

```
root@srx1# set protocols pim rp local address 172.30.30.254
```

We then move to the MSDP configuration. Configure neighbor relationship between srx1 and ex2 loopback addresses:

```
root@srx1# set protocols msdp group anycast mode mesh-group

root@srx1# set protocols msdp group anycast local-address 172.30.30.1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set protocols msdp group anycast peer 172.30.30.6
```

We then configure the firewall filter that is protecting the RE to allow MSDP:

```
root@srx1# set firewall family inet filter protect-re-ipv4 term msdp-source from protocol tcp

root@srx1# set firewall family inet filter protect-re-ipv4 term msdp-source from source-port msdp

root@srx1# set firewall family inet filter protect-re-ipv4 term msdp-source then accept

root@srx1# set firewall family inet filter protect-re-ipv4 term msdp-destination from protocol tcp

root@srx1# set firewall family inet filter protect-re-ipv4 term msdp-destination from destination-port msdp

root@srx1# set firewall family inet filter protect-re-ipv4 term msdp-destination then accept
```

We move the terms to a place prior to the discards term:

```
root@srx1# insert firewall family inet filter protect-re-ipv4 term msdp-source before term icmp

root@srx1# insert firewall family inet filter protect-re-ipv4 term msdp-destination before term icmp
```

The same steps are repeated on ex2:

ex2:

```
root@ex2# set interfaces lo0 unit 0 family inet address 172.30.30.6/32 primary

root@ex2# set interfaces lo0 unit 0 family inet address 172.30.30.254/32


root@ex2# set protocols pim rp local address 172.30.30.254
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex2# set protocols msdp group anycast mode mesh-group

root@ex2# set protocols msdp group anycast local-address 172.30.30.6

root@ex2# set protocols msdp group anycast peer 172.30.30.1


root@ex2# set firewall family inet filter protect-re-ipv4 term msdp-
source from protocol tcp

root@ex2# set firewall family inet filter protect-re-ipv4 term msdp-
source from source-port msdp

root@ex2# set firewall family inet filter protect-re-ipv4 term msdp-
source then accept

root@ex2# set firewall family inet filter protect-re-ipv4 term msdp-
destination from protocol tcp

root@ex2# set firewall family inet filter protect-re-ipv4 term msdp-
destination from destination-port msdp

root@ex2# set firewall family inet filter protect-re-ipv4 term msdp-
destination then accept


root@ex2# insert firewall family inet filter protect-re-ipv4 term
msdp-source before term icmp

root@ex2# insert firewall family inet filter protect-re-ipv4 term
msdp-destination before term icmp
```

Last step is to configure all of the other devices to use the Anycast-RP address.

All other devices:

```
root@srx2# set protocols pim rp static address 172.30.30.254
```

Verification:

```
root@srx1# run show msdp
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Peer address    Local address   State       Last up/down Peer-Group
SA Count
172.30.30.6     172.30.30.1     Established     00:11:40 anycast
0/0


root@srx1# run show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.2.127.254
    Source: *
    RP: 172.30.30.254
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:04:58
    Downstream neighbors:
        Interface: ge-0/0/4.15
            172.30.0.10 State: Join Flags: SRW Timeout: 152
            Uptime: 00:04:58 Time since last Join: 00:00:58

Group: 239.255.255.2
    Source: *
    RP: 172.30.30.254
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:04:58
    Downstream neighbors:
        Interface: ge-0/0/4.15
            172.30.0.10 State: Join Flags: SRW Timeout: 152
            Uptime: 00:04:58 Time since last Join: 00:00:58

Group: 239.255.255.3
    Source: *
    RP: 172.30.30.254
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:04:58
    Downstream neighbors:
        Interface: ge-0/0/4.15
            172.30.0.10 State: Join Flags: SRW Timeout: 152
            Uptime: 00:04:58 Time since last Join: 00:00:58

Instance: PIM.master Family: INET6
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

5) Configure BSR for IPv6 RP discovery. The RP candidates, servicing ff1e::2010/128, should be srx1 and srx2. Configure srx2 to be the primary BSR and configure srx4 to be the backup:

srx1:

```
root@srx1# set protocols pim rp local family inet6 address
fd17:f0f4:f691:30::1

root@srx1# set protocols pim rp local family inet6 group-ranges
ff1e::2010/128
```

srx2:

```
root@srx2# set protocols pim rp bootstrap family inet6 priority 200

root@srx2# set protocols pim rp local family inet6 address
fd17:f0f4:f691:30::2

root@srx2# set protocols pim rp local family inet6 group-ranges
ff1e::2010/128
```

srx4:

```
root@srx4# set protocols pim rp bootstrap family inet6 priority 100
```

Verification:

```
root@srx4# run show pim rps
Instance: PIM.master
Address family INET
RP address                Type        Holdtime Timeout Groups Group
prefixes
172.30.30.254             static             0    None      0
224.0.0.0/4

Address family INET6
RP address                Type        Holdtime Timeout Groups Group
prefixes
fd17:f0f4:f691:30::1      bootstrap         150     141      0
ff1e::2010/128
fd17:f0f4:f691:30::2      bootstrap         150     141      0
ff1e::2010/128
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

6) Multicast traffic for group 239.255.255.2 to Rcv1 should stay on the RPT:

The SPT to RPT switchover can be prevented by configuring an **spt-threshold** of infinity. We first configure a routing-policy that matches the relevant traffic. In this case, we match the multicast group address 239.255.255.2 and the Src1 address 192.168.40.100. We end the policy with a second term that rejects all other traffic:

```
root@srx3#set policy-options policy-statement spt.infinity term 1 from
route-filter 239.255.255.2/32 exact

root@srx3#set policy-options policy-statement spt.infinity term 1 from
source-address-filter 192.168.40.100/32 exact

root@srx3#set policy-options policy-statement spt.infinity term 1 then
accept

root@srx3#set policy-options policy-statement spt.infinity term 2 then
reject
```

The policy is applied under the [**protocols pim spt-threshold infinity**] stanza. The matched traffic will stay on the RPT forever:

```
root@srx3#set protocols pim spt-threshold infinity spt.infinity
```

Verification:

```
root@srx3# run show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.255.255.2
    Source: *
    Flags: sparse,rptree,wildcard
    Upstream interface: ge-0/0/4.13
    Upstream neighbor: 172.30.0.5
    Upstream state: Join to RP
    Keepalive timeout: 357
    Uptime: 00:00:32
    Downstream neighbors:
        Interface: ge-0/0/4.44
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
            192.168.44.254 State: Join Flags: S Timeout: 178
            Uptime: 00:00:32 Time since last Join: 00:00:32

Group: 239.255.255.3
    Source: 192.168.40.100
    Flags: sparse,spt
    Upstream interface: ge-0/0/4.13
    Upstream neighbor: 172.30.0.5
    Upstream state: Join to Source
    Keepalive timeout: 237
    Uptime: 00:02:50
    Downstream neighbors:
        Interface: ge-0/0/4.44
            192.168.44.10 State: Join Flags: S Timeout: 159
            Uptime: 00:02:50 Time since last Join: 00:00:50

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

7) On ex3_ex4 and srx3a, we use an import policy to block PIM from signaling traffic for certain groups. On ex3_ex4, we use a multicast scope-policy to block multicast traffic:

On ex3_ex4 and SRX3a:

```
root@ex3_ex4# set policy-options policy-statement block.pim.control
term 1 from interface ge-0/0/14.80

root@ex3_ex4# set policy-options policy-statement block.pim.control
term 1 from route-filter 239.255.255.2/32 exact

root@ex3_ex4# set policy-options policy-statement block.pim.control
term 1 then reject


root@ex3_ex4# set policy-options policy-statement block.pim.control
term 2 from interface ge-0/0/14.80

root@ex3_ex4# set policy-options policy-statement block.pim.control
term 2 from route-filter 239.255.255.3/32 exact

root@ex3_ex4# set policy-options policy-statement block.pim.control
term 2 then reject
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set policy-options policy-statement block.pim.control
term 3 from interface ge-0/0/14.80

root@ex3_ex4# set policy-options policy-statement block.pim.control
term 3 from route-filter ff1e::2010/128 exact

root@ex3_ex4# set policy-options policy-statement block.pim.control
term 3 then reject


root@ex3_ex4# set protocols pim import block.pim.control


root@ex3_ex4# set policy-options policy-statement dc1.scope term 1
from interface ge-0/0/14.80

root@ex3_ex4# set policy-options policy-statement dc1.scope term 1
from route-filter 239.255.255.2/32 exact

root@ex3_ex4# set policy-options policy-statement dc1.scope term 1
from route-filter 239.255.255.3/32 exact

root@ex3_ex4# set policy-options policy-statement dc1.scope term 1
then reject


root@ex3_ex4# set routing-options multicast scope-policy dc1.scope
```

8) To filter SA message from DC1, we create and apply a policy to MSDP in the import direction:

On srx1 and ex2:

```
root@ex3_ex4# set policy-options policy-statement block.sa.msdp term 1
from neighbor 172.31.0.1

root@ex3_ex4# set policy-options policy-statement block.sa.msdp term 1
from route-filter 239.255.255.2/32 exact

root@ex3_ex4# set policy-options policy-statement block.sa.msdp term 1
from source-address-filter 172.31.10.10/32 exact

root@ex3_ex4# set policy-options policy-statement block.sa.msdp term 1
then reject
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex3_ex4# set policy-options policy-statement block.sa.msdp term 2
then accept


root@ex3_ex4# set protocols msdp group anycast import block.sa.msdp
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 6: Class of Service

1) Configure IPv4 and IPv6 code-point aliases on every device:

```
root@srx1# set class-of-service code-point-aliases dscp control-
traffic 111000

root@srx1# set class-of-service code-point-aliases dscp voice-traffic
101110

root@srx1# set class-of-service code-point-aliases dscp data-traffic
010010

root@srx1# set class-of-service code-point-aliases dscp effort-traffic
000000


root@srx1# set class-of-service code-point-aliases dscp-ipv6 control-
traffic 111000

root@srx1# set class-of-service code-point-aliases dscp-ipv6 voice-
traffic 101110

root@srx1# set class-of-service code-point-aliases dscp-ipv6 data-
traffic 010010

root@srx1# set class-of-service code-point-aliases dscp-ipv6 effort-
traffic 000000
```

Verification:

```
root@srx1# run show class-of-service code-point-aliases
Code point type: dscp
  Alias               Bit pattern
  af11                001010
  af12                001100
  af13                001110
  af21                010010
  af22                010100
  af23                010110
  af31                011010
  af32                011100
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  af33                  011110
  af41                  100010
  af42                  100100
  af43                  100110
  be                    000000
  control-traffic       111000
  cs1                   001000
  cs2                   010000
  cs3                   011000
  cs4                   100000
  cs5                   101000
  cs6                   110000
  cs7                   111000
  data-traffic          010010
  ef                    101110
  effort-traffic        000000
  nc1                   110000
  nc2                   111000
  voice-traffic         101110

Code point type: dscp-ipv6
  Alias                 Bit pattern
  af11                  001010
  af12                  001100
  af13                  001110
  af21                  010010
  af22                  010100
  af23                  010110
  af31                  011010
  af32                  011100
  af33                  011110
  af41                  100010
  af42                  100100
  af43                  100110
  be                    000000
  control-traffic       111000
  cs1                   001000
  cs2                   010000
  cs3                   011000
  cs4                   100000
  cs5                   101000
  cs6                   110000
  cs7                   111000
  data-traffic          010010
  ef                    101110
  effort-traffic        000000
  nc1                   110000
  nc2                   111000
  voice-traffic         101110
```

2) Configure forwarding classes and map them to the output queues on all of the devices:

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set class-of-service forwarding-classes queue 7 control

root@srx1# set class-of-service forwarding-classes queue 5 voice

root@srx1# set class-of-service forwarding-classes queue 3 data

root@srx1# set class-of-service forwarding-classes queue 0 best-effort
```

Verification:

```
root@srx1# run show class-of-service forwarding-class
Forwarding class                  ID                        Queue
Policing priority    SPU priority
  best-effort                      0                          0
normal           low
  expedited-forwarding             1                          1
normal           low
  assured-forwarding               2                          2
normal           low
  data                             3                          3
normal           low
  voice                            5                          5
normal           low
  control                          7                          7
normal           low
```

3) Configuring the schedulers on SRX devices according to the parameters given in the task description:

```
root@srx1# set class-of-service schedulers control-sched shaping-rate
percent 7

root@srx1# set class-of-service schedulers control-sched buffer-size
percent 10

root@srx1# set class-of-service schedulers control-sched priority high


root@srx1# set class-of-service schedulers voice-sched shaping-rate
percent 15

root@srx1# set class-of-service schedulers voice-sched buffer-size
percent 5
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set class-of-service schedulers voice-sched buffer-size
exact

root@srx1# set class-of-service schedulers voice-sched priority
strict-high


root@srx1# set class-of-service schedulers data-sched transmit-rate
percent 60

root@srx1# set class-of-service schedulers data-sched buffer-size
remainder

root@srx1# set class-of-service schedulers data-sched priority medium-
low


root@srx1# set class-of-service schedulers best-effort-sched transmit-
rate remainder

root@srx1# set class-of-service schedulers best-effort-sched buffer-
size remainder

root@srx1# set class-of-service schedulers best-effort-sched priority
low
```

4) Configuring the schedulers on EX devices according to the parameters given in the task description:

```
root@ex1# set class-of-service schedulers control-sched shaping-rate
percent 7

root@ex1# set class-of-service schedulers control-sched buffer-size
percent 10

root@ex1# set class-of-service schedulers control-sched priority
strict-high


root@ex1# set class-of-service schedulers voice-sched shaping-rate
percent 15
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set class-of-service schedulers voice-sched buffer-size
percent 5

root@ex1# set class-of-service schedulers voice-sched buffer-size
exact

root@ex1# set class-of-service schedulers voice-sched priority strict-
high


root@ex1# set class-of-service schedulers data-sched transmit-rate
percent 60

root@ex1# set class-of-service schedulers data-sched buffer-size
remainder

root@ex1# set class-of-service schedulers data-sched priority low


root@ex1# set class-of-service schedulers best-effort-sched transmit-
rate remainder

root@ex1# set class-of-service schedulers best-effort-sched buffer-
size remainder

root@ex1# set class-of-service schedulers best-effort-sched priority
low
```

5) Applying the schedulers to all but the OoB interfaces:

To be able to apply the schedulers, we start out configuring a scheduler-map:

On SRXs:

```
root@srx1# set class-of-service scheduler-maps srx.map forwarding-
class control scheduler control-sched

root@srx1# set class-of-service scheduler-maps srx.map forwarding-
class voice scheduler voice-sched

root@srx1# set class-of-service scheduler-maps srx.map forwarding-
class data scheduler data-sched

root@srx1# set class-of-service scheduler-maps srx.map forwarding-
class best-effort scheduler best-effort-sched
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

We apply the scheduler-map to all interfaces under the [**class-of-service interface all**] stanza:

```
root@srx1# set class-of-service interfaces all scheduler-map srx.map
```

We configure the OoB interface under the same stanza. By doing so, we nullify the configuration statement 'interface all' for the OoB interface. It does not matter what the configuration is, as long as the interface is configured more specifically.

In this case, we simply apply the default classifier to the interface (something that is needed in a future task anyway):

```
root@srx1# set class-of-service interfaces ge-0/0/0 unit 0 classifiers
dscp-ipv6 default
```

```
root@srx1# set class-of-service interfaces ge-0/0/0 unit 0 classifiers
inet-precedence default
```

```
root@srx1# run show class-of-service interface
Physical interface: ae0, Index: 128
Queues supported: 8, Queues in use: 6
  Scheduler map: srx.map, Index: 39898
  Congestion-notification: Disabled

Physical interface: ae1, Index: 129
Queues supported: 8, Queues in use: 6
  Scheduler map: srx.map, Index: 39898
  Congestion-notification: Disabled

  Logical interface: ae1.0, Index: 70

Physical interface: ge-0/0/0, Index: 136
Queues supported: 8, Queues in use: 6
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: ge-0/0/0.0, Index: 74
    Object                  Name                    Type
Index
    Classifier              dscp-ipv6-default       dscp-ipv6
8
    Classifier              ipprec-default          ip
12

Physical interface: ge-0/0/1, Index: 137
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Queues supported: 8, Queues in use: 6
  Scheduler map: srx.map, Index: 39898
  Congestion-notification: Disabled

  Logical interface: ge-0/0/1.0, Index: 75

Physical interface: ge-0/0/10, Index: 146
Queues supported: 8, Queues in use: 6
  Scheduler map: srx.map, Index: 39898
  Congestion-notification: Disabled

  Logical interface: ge-0/0/10.0, Index: 88

Physical interface: ge-0/0/11, Index: 147
Queues supported: 8, Queues in use: 6
  Scheduler map: srx.map, Index: 39898
  Congestion-notification: Disabled

Physical interface: ge-0/0/12, Index: 148
Queues supported: 8, Queues in use: 6
  Scheduler map: srx.map, Index: 39898
  Congestion-notification: Disabled
```

We do the same thing on the ex switches.

On all EX devices:

```
root@srx1# set class-of-service scheduler-maps ex.map forwarding-class
control scheduler control-sched

root@srx1# set class-of-service scheduler-maps ex.map forwarding-class
voice scheduler voice-sched

root@srx1# set class-of-service scheduler-maps ex.map forwarding-class
data scheduler data-sched

root@srx1# set class-of-service scheduler-maps ex.map forwarding-class
best-effort scheduler best-effort-sched


root@srx1# set class-of-service interfaces all scheduler-map ex.map
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx1# set class-of-service interfaces me0 unit 0 classifiers
dscp-ipv6 default

root@srx1# set class-of-service interfaces me0 unit 0 classifiers
inet-precedence default
```

6) Configure the BA classifiers based on table 38. We import the values of the default classifiers as well:

```
root@srx1# set class-of-service classifiers dscp main.classifier
import default

root@srx1# set class-of-service classifiers dscp main.classifier
forwarding-class control loss-priority low code-points control-traffic

root@srx1# set class-of-service classifiers dscp main.classifier
forwarding-class voice loss-priority low code-points voice-traffic

root@srx1# set class-of-service classifiers dscp main.classifier
forwarding-class data loss-priority low code-points data-traffic

root@srx1# set class-of-service classifiers dscp main.classifier
forwarding-class best-effort loss-priority high code-points effort-
traffic
```

```
root@srx1# set class-of-service classifiers dscp-ipv6
main.classifier.ipv6 forwarding-class control loss-priority low code-
points control-traffic

root@srx1# set class-of-service classifiers dscp-ipv6
main.classifier.ipv6 forwarding-class voice loss-priority low code-
points voice-traffic

root@srx1# set class-of-service classifiers dscp-ipv6
main.classifier.ipv6 forwarding-class data loss-priority low code-
points data-traffic

root@srx1# set class-of-service classifiers dscp-ipv6
main.classifier.ipv6 forwarding-class best-effort loss-priority high
code-points effort-traffic
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification:

```
root@srx1# run show class-of-service classifier name main.classifier
Classifier: main.classifier, Code point type: dscp, Index: 27271
  Code point         Forwarding class                  Loss priority
  000000             best-effort                       high
  000001             best-effort                       low
  000010             best-effort                       low
  000011             best-effort                       low
  000100             best-effort                       low
  000101             best-effort                       low
  000110             best-effort                       low
  000111             best-effort                       low
  001000             best-effort                       low
  001001             best-effort                       low
  001010             assured-forwarding                low
  001011             best-effort                       low
  001100             assured-forwarding                high
  001101             best-effort                       low
  001110             assured-forwarding                high
  001111             best-effort                       low
  010000             best-effort                       low
  010001             best-effort                       low
  010010             data                              low
  010011             best-effort                       low
  010100             best-effort                       low
  010101             best-effort                       low
  010110             best-effort                       low
  010111             best-effort                       low
  011000             best-effort                       low
  011001             best-effort                       low
  011010             best-effort                       low
  011011             best-effort                       low
  011100             best-effort                       low
  011101             best-effort                       low
  011110             best-effort                       low
  011111             best-effort                       low
  100000             best-effort                       low
  100001             best-effort                       low
  100010             best-effort                       low
  100011             best-effort                       low
  100100             best-effort                       low
  100101             best-effort                       low
  100110             best-effort                       low
  100111             best-effort                       low
  101000             best-effort                       low
  101001             best-effort                       low
  101010             best-effort                       low
  101011             best-effort                       low
  101100             best-effort                       low
  101101             best-effort                       low
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
101110          voice                   low
101111          best-effort             low
110000          data                    low
110001          best-effort             low
110010          best-effort             low
110011          best-effort             low
110100          best-effort             low
110101          best-effort             low
110110          best-effort             low
110111          best-effort             low
111000          control                 low
111001          best-effort             low
111010          best-effort             low
111011          best-effort             low
111100          best-effort             low
111101          best-effort             low
111110          best-effort             low
111111          best-effort             low
```

7) Apply the classifiers to all physical interface units. Make sure that the OoB interface is not affected:

To configure all units with the classifier, we issue the following configuration command:

```
root@srx1# set class-of-service interfaces all unit * classifiers dscp
main.classifier

root@srx1# set class-of-service interfaces all unit * classifiers
dscp-ipv6 main.classifier.ipv6
```

We already configured the OoB with the default classifier earlier on. This was to nullify the 'interfaces all' configuration that was used to apply the scheduler-map. The configuration we applied there was the following:

```
root@srx1# set class-of-service interfaces ge-0/0/0 unit 0 classifiers
dscp-ipv6 default

root@srx1# set class-of-service interfaces ge-0/0/0 unit 0 classifiers
inet-precedence default
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

8) Configure IPv4 and IPv6 rewrite rules to mark outgoing traffic:

```
root@srx1# set class-of-service rewrite-rules dscp main.rewrite
forwarding-class control loss-priority low code-point control-traffic

root@srx1# set class-of-service rewrite-rules dscp main.rewrite
forwarding-class voice loss-priority low code-point voice-traffic

root@srx1# set class-of-service rewrite-rules dscp main.rewrite
forwarding-class data loss-priority low code-point data-traffic

root@srx1# set class-of-service rewrite-rules dscp main.rewrite
forwarding-class best-effort loss-priority low code-point effort-
traffic

root@srx1# set class-of-service rewrite-rules dscp-ipv6
main.rewrite.ipv6 forwarding-class control loss-priority low code-
point control-traffic

root@srx1# set class-of-service rewrite-rules dscp-ipv6
main.rewrite.ipv6 forwarding-class voice loss-priority low code-point
voice-traffic

root@srx1# set class-of-service rewrite-rules dscp-ipv6
main.rewrite.ipv6 forwarding-class data loss-priority low code-point
data-traffic

root@srx1# set class-of-service rewrite-rules dscp-ipv6
main.rewrite.ipv6 forwarding-class best-effort loss-priority low code-
point effort-traffic
```

Verification:

```
root@srx1# run show class-of-service rewrite-rule

Rewrite rule: main.rewrite, Code point type: dscp, Index: 64649
  Forwarding class                Loss priority       Code point
  best-effort                     low                 000000
  data                            low                 010010
  voice                           low                 101110
  control                         low                 111000

Rewrite rule: main.rewrite.ipv6, Code point type: dscp-ipv6, Index:
59113
  Forwarding class                Loss priority       Code point
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  best-effort                    low             000000
  data                           low             010010
  voice                          low             101110
  control                        low             111000
```

9) Applying the rewrite rules to all of the devices:

```
root@srx1# set class-of-service interfaces all unit * rewrite-rules
dscp main.rewrite

root@srx1# set class-of-service interfaces all unit * rewrite-rules
dscp-ipv6 main.rewrite.ipv6
```

Verification:

```
root@srx1# run show class-of-service interface

  Logical interface: ge-0/0/4.12, Index: 77
    Object                 Name                   Type
Index
    Rewrite                main.rewrite           dscp
64649
    Rewrite                main.rewrite.ipv6      dscp-ipv6
59113
    Classifier             main.classifier        dscp
27271
    Classifier             main.classifier.ipv6   dscp-ipv6
24132
```

10) In this task, we use a multi field firewall filter to classify traffic and to assign it the proper loss priority. We apply the MFCs on input in both the voice as well as the data VLAN on ex1 and ex2:

```
root@ex1# set firewall family ethernet-switching filter vlan601 term
1 then forwarding-class voice

root@ex1# set firewall family ethernet-switching filter vlan601 term
1 then loss-priority low

root@ex1# set firewall family ethernet-switching filter vlan600 term
1 then forwarding-class data
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex1# set firewall family ethernet-switching filter vlan600 term
1 then loss-priority low

root@ex1# set vlans User filter input vlan600

root@ex1# set vlans Voice filter input vlan601
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Chapter 8: A Full Day Lab Challenge 2

Task 1: Initial System Configuration

1) Load the latest workbook **baseline** configuration for this lab on every device (including the access-switch and the vr-device). Use the command "**load override terminal**". You will need to reboot the SRX and VR router. Please do **not** change the root password on our devices to prevent unnecessary password recovery actions.

```
load override terminal
```

2) Configure the hostname and management address for each device as indicated below:

Device	Hostname	Domain	Management IP
SRX 1	srx-1	inetzero.com	10.10.0.1/24
SRX 2	srx-2	inetzero.com	10.10.0.2/24
SRX 3	Srx-3	inetzero.com	10.10.0.3/24
SRX 4	srx-4	inetzero.com	10.10.0.4/24
Ex4200 1	Ex1	inetzero.com	10.10.0.5/24
Ex4200 2	Ex2	inetzero.com	10.10.0.6/24
Ex4200 3/4	Ex_VC	inetzero.com	10.10.0.7/24

```
[edit system]
+   host-name srx-1;
+   domain-name inetzero.com;
```

3) Add a route to the management network 10.10.10.0/24 and use 10.10.1.254 as the next-hop for this route.

```
[edit]
root@srx-1# show | compare
[edit]
+  routing-options {
+      static {
+          route 10.10.10.0/24 next-hop 10.10.1.254;
+      }
+  }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

4) Users have complained that when they logon to either SRX1 or SRX2, they get the following error:

```
warning: user "radius-template" does not have a valid login class

error: Unable to authenticate: bad auth parameter.
Login as root and 'commit' the configuration.
admin@srx-1-re-0>
```

They are unable to issue any commands. Fix this issue.

```
[edit]
root@srx-1-re-0# show | compare
[edit system login class super-user-admin-class]
+    permissions all;
```

5) You need to enforce the company password policy on every device. Make sure all users have a minimum password consisting of 8 characters. Also ensure there is a minimum of 1 change of case, digits or punctuation in their password.

```
[edit system]
+   login {
+       password {
+           minimum-length 8;
+           minimum-changes 1;
+       }
+   }
```

6) Upon login, users should have the following message displayed:

"UNAUTHORIZED USE IS PROHIBITED! By logging on you agree to our terms and conditions."

For security reasons, users should not be able to retry their password more than 3 times. In the event that they fail on the last try, they should be locked out of the device for 30 minutes.

```
[edit system login]
+     message "UNAUTHORIZED USE IS PROHIBITED! By logging on you agree to our
terms and conditions."
+     retry-options {
+         tries-before-disconnect 3;
+         lockout-period 30;
+     }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

7) To authenticate users, the devices should first try the radius server with IP address 10.10.10.1. The local user accounts should be available as a backup. No other method should be used. The radius server should be configured with the password “radiuspassword” and the device should retry contacting the radius server 2 times. The radius server should timeout after 10 seconds. Ensure that the user 'radius-template' that exists on SRX1 and SRX2 is configured on every device.

```
[edit system]
+  authentication-order [ radius password ];
+  radius-server {
+      10.10.10.1 {
+          secret "$9$69E6C0IEhrMWxp07VY2UD69Apu1hSeWX-vW4ZDjf5369"; ##
SECRET-DATA
+          timeout 10;
+          retry 2;
+      }
+  }

[edit system]
+   login {
+       class super-user-admin-class {
+           permissions all;
+           idle-timeout 20;
+       }
+       user radius-template {
+           class super-user-admin-class;
+       }
+   }
```

On SRX1 and SRX2

```
[edit]
root@srx-1-re-0# show | compare
[edit system]
-  authentication-order [ password tacplus radius ];
+  authentication-order [ radius password ];
```

8) UTC should be the default time-zone on every device. Use 10.10.10.1 as DNS server so names can be resolved.

```
[edit system]
+  time-zone UTC;
+  name-server {
+      10.10.10.1;
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+   }
```

9) Set the syslog server to 10.10.10.1 and log all events up to the informational level. Ensure that the hostname of the device appears in the logs. The syslog messages should be sourced from the management address.

```
[edit system]
+   syslog {
+       host 10.10.10.1 {
+           any info;
+           log-prefix srx1;
+           source-address 10.10.0.1;
+       }
+   }
```

10) Make sure that logs have a maximum of 5 archived files which are 2 megabyte in size. All users should be able to access the logs. Ensure that the year is displayed as part of the messages.

```
[edit system]
+   syslog {
+       archive size 2m files 5 world-readable;
+       time-format year;
+   }
```

11) Set the system to rate-limit ICMP to 500 packets per 5 seconds. Do this for both IPv4 as well as IPv6.

```
[edit system]
+   internet-options {
+       icmpv4-rate-limit packet-rate 500 bucket-size 5;
+       icmpv6-rate-limit packet-rate 500 bucket-size 5;
+   }
```

12) The auxiliary port is not used. For security reasons, this port should be disabled. Also make sure that USB ports are disabled on every SRX device.

```
[edit system]

+  ports {
+      auxiliary disable;
+  }

 [edit chassis]
+   usb {
+       storage {
+           disable;
+       }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+    }
```

13) The devices needs to accept SSH for management. It should use version 2. The maximum number of connections should be limited to 3 and it should only accept 3 connections per minute.

```
[edit system]
+   services {
+       ssh {
+           protocol-version v2;
+           connection-limit 3;
+           rate-limit 3;
+       }
```

14) Enable netconf over SSH. Use port 1026 for this.

```
set system services netconf ssh port 1026
```

15) Configure NTP to use the server s1. Use the password “workbook”. NTP should also sync when the device is booting.

```
[edit system]
+   ntp {
+       boot-server 10.10.10.1;
+       authentication-key 1 type md5 value "$9$qf39yrv8xdIESeWxwsqmfznC"; ##
SECRET-DATA
+       server 10.10.10.1 key 1 prefer; ## SECRET-DATA
+       trusted-key 1;
+   }
```

16) Set up SNMPv2 using the community 'jncie' as a read-only community. Use a list and add server s1 as well as 20.20.20.1. SNMP traps need to be send to server s1. Set the SNMP location to “inetzero dc”.

```
[edit]
+   snmp {
+       client-list list {
+           10.10.10.1/32;
+       }
+       community jncie {
+           authorization read-only;
+           client-list-name list;
+       }
+       trap-group s1-traps {
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

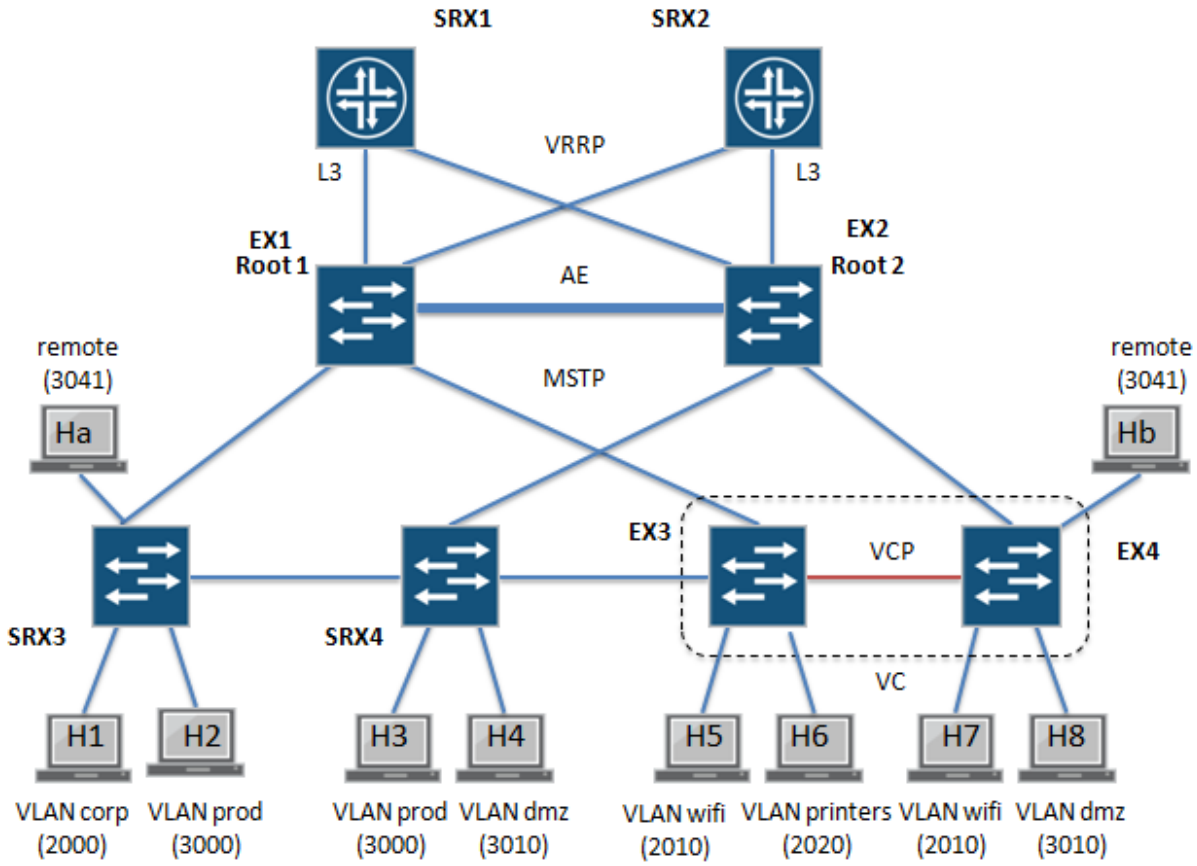
https://t.me/learningnets

```
+           targets {
+               10.10.10.1;
+           }
+       }
+   }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 2: Layer2 Configuration



Device Name	VLAN Name	VLAN ID	Interface	Mode	Subnet/ B End
EX1	-	All	AE0	Trunk	EX2
	-		Ge-0/0/6	Trunk	SRX1
	-		Ge-0/0/10	Trunk	SRX2
	-	All	Ge-0/0/16	Trunk	SRX3
	-	All	Ge-0/0/18	Trunk	EX_VC
		99	Ge-0/0/17	Access	SRX4
EX2	-	All	AE0	Trunk	EX1
	-	All	Ge-0/0/6	Trunk	SRX2
	-		Ge-0/0/10	Trunk	SRX1
	-	All	Ge-0/0/17	Trunk	SRX4
	-	All	Ge-0/0/19	Trunk	EX_VC
EX_VC	-	All	Ge-0/0/11	Trunk	SRX4
	-	All	Ge-0/0/18	Trunk	EX1
	-	All	Ge-1/0/19	Trunk	EX2

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

	guest	3041	Ge-1/0/23	Access	Hb
	DMZ	3010	Ge-0/0/23	Access	H8
	Printers	2020	Ge-0/0/8	Access	H6
	Wi-Fi	2010	Ge-1/0/11	Access	H5
SRX1	-	per-task	Ge-0/0/6	Tagged	Ex1
	-	per-task	Ge-0/0/10	Tagged	Ex2
SRX2	-	per-task	Ge-0/0/6	Tagged	Ex2
	-	per-task	Ge-0/0/10	Tagged	Ex1
SRX3	-	All	Ge-0/0/1	Trunk	SRX4
	-	All	Ge-0/0/12	Trunk	EX1
	guest	3041	Ge-0/0/15	Access	Ha
	Prod	3000	Ge-0/0/7	Access	H2
SRX4	-	All	Ge-0/0/1	Trunk	SRX3
	-	All	Ge-0/0/11	Trunk	EX_VC
	-	All	Ge-0/0/13	Trunk	EX2
	Prod	3000	Ge-0/0/7	Access	H3
	DMZ	3010	Ge-0/0/8	Access	H4
All	Corp-old	1000	Vlan.1000	-	21.10.0.0/24
All	Corp	2000	Vlan.2000	-	21.20.0.0/24
All	Wi-Fi	2010	Vlan.2010	-	21.20.10.0/24
All	Printers	2020	Vlan.2020	-	21.20.20.0/24
All	Mobile	2030	Vlan.2030	-	21.20.30.0/24
All	VoIP	2040	Vlan.2040	-	21.20.40.0/24
All	Prod	3000	Vlan.3000	-	21.30.0.0/24
All	DMZ	3010	Vlan.3010	-	21.30.10.0/24
All	DMZ 2	3020	Vlan.3020	-	21.30.20.0/24
All	Servers	3040	Vlan.3040	-	21.30.40.0/24
	OSPF	99	Vlan.99	-	

Table 2.1

1) EX3 and EX4 need to be set up as a Virtual Chassis. This should be set up in such a way that new switches are not able to join automatically. EX3 should be the Master and EX4 should be the Backup. They should be configured as member 0 and member 1 respectfully.

On both EX3 and EX4;

```
request virtual-chassis vc-port set interface vcp-0 local
request virtual-chassis vc-port set interface vcp-1 local
```

On EX3;

```
[edit]
+  virtual-chassis {
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+      preprovisioned;
+      no-split-detection;
+      member 0 {
+          role routing-engine;
+          serial-number BM0210045284;
+      }
+      member 1 {
+          role routing-engine;
+          serial-number BM0210359941;
+      }
+  }
```

Verification

```
{master:0}
root@ex_vc> show virtual-chassis status

Preprovisioned Virtual Chassis
Virtual Chassis ID: 5412.ad87.47a6
Virtual Chassis Mode: Enabled
                                                Mstr           Mixed Neighbor List
Member ID  Status   Serial No    Model     prio  Role      Mode ID  Interface
0 (FPC 0)  Prsnt    BM0210045284 ex4200-24t 129  Master*      N  1  vcp-0
1 (FPC 1)  Prsnt    BM0210359941 ex4200-24t 129  Backup       N  0  vcp-1
```

2) Users active on VLAN 3010 have complained about the network not working correctly. Look into the issue and make sure that the H8_VR is able to ping 21.30.10.4.

```
{master:0}
root@ex_vc> show VLANs DMZ
Name            Tag     Interfaces
DMZ             3010
                        None

{master:0}
root@ex_vc> show ethernet-switching interfaces ge-0/0/23
Interface     State  VLAN members         Tag   Tagging  Blocking
ge-0/0/23.0   up     DMZ-2                3020  untagged unblocked


{master:0}[edit]
root@ex_vc# show | compare
[edit interfaces ge-0/0/11 unit 0 family ethernet-switching VLAN]
-       members [ VoIP corp corp-old mobile printers prod servers Wi-Fi DMZ-2
];
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+       members [ VoIP corp corp-old mobile printers prod servers Wi-Fi DMZ-2
DMZ ];
[edit interfaces ge-0/0/18 unit 0 family ethernet-switching VLAN]
-       members [ VoIP corp corp-old mobile printers prod servers Wi-Fi DMZ-2
];
+       members [ VoIP corp corp-old mobile printers prod servers Wi-Fi DMZ-2
DMZ ];
[edit interfaces ge-0/0/23 unit 0 family ethernet-switching VLAN]
-       members DMZ-2;
+       members DMZ;
[edit interfaces ge-1/0/19 unit 0 family ethernet-switching VLAN]
-       members [ VoIP corp corp-old mobile printers prod servers Wi-Fi DMZ-2
];
+       members [ VoIP corp corp-old mobile printers prod servers Wi-Fi DMZ-2
DMZ ];

{master:0}

root@ex_vc> ping 21.30.10.4 routing-instance H8_VR rapid
PING 21.30.10.4 (21.30.10.4): 56 data bytes
!!!!!!!!!!!!!!!
--- 21.30.10.4 ping statistics ---
5 packets transmitted, 5 packets received, +10 duplicates, 0% packet loss
round-trip min/avg/max/stddev = 1.577/7.977/16.045/4.706 ms

{master:0}
root@ex_vc> show spanning-tree interface
Spanning-tree is not enabled at global level


{master:0}[edit]
root@ex_vc# show | compare
[edit]
+  protocols {
+      rstp;
+  }
```

Verification

```
{master:0}
root@ex_vc> ping 21.30.10.4 routing-instance H8_VR rapid
PING 21.30.10.4 (21.30.10.4): 56 data bytes
!!!!!
--- 21.30.10.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.119/2.789/3.915/1.303 ms
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
{master:0}
root@ex_vc> show ethernet-switching interfaces ge-0/0/23
Interface     State  VLAN members         Tag   Tagging  Blocking
ge-0/0/23.0   up     DMZ                  3010  untagged unblocked
```

3) One of the engineers has configured an aggregated interface between EX1 and EX2. The details of the aggregated interfaces are listed below:

Device	Interface
EX1	Ge-0/0/12, Ge-0/0/13
EX2	Ge-0/0/12, Ge-0/0/13

Table 2.2

Even though the interfaces is displayed as 'up', it is not working. The engineer has provided the following output:

```
root@ex1> show interfaces terse ae0
Interface               Admin Link Proto    Local                 Remote
ae0                     up    up
ae0.0                   up    up   eth-switch
```

Fix the issues with the aggregated interface.

{master:0}

```
root@ex-1> show interfaces ae0
error: device ae0 not found

{master:0}

root@ex-1> show lacp interfaces

{master:0}

root@ex-1> show interfaces ge-0/0/12 | match AE
    Protocol aenet, AE bundle: ae0.0

{master:0}

root@ex-1> show interfaces ge-0/0/13 | match AE
    Protocol aenet, AE bundle: ae0.0

{master:0}

root@ex-1> show configuration chassis
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
{master:0}[edit]
root@ex-1# show | compare
[edit]
+  chassis {
+      aggregated-devices {
+          ethernet {
+              device-count 4;
+          }
+      }
+  }


{master:0}[edit]

root@ex-2# show | compare
[edit]
+  chassis {
+      aggregated-devices {
+          ethernet {
+              device-count 4;
+          }
+      }
+  }
```

Verification

```
{master:0}[edit]
root@ex-1# run show lacp interfaces
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout
Activity
      ge-0/0/12      Actor    No    No   Yes  Yes  Yes   Yes     Fast
Active
      ge-0/0/12    Partner    No    No   Yes  Yes  Yes   Yes     Fast
Active
      ge-0/0/13      Actor    No    No   Yes  Yes  Yes   Yes     Fast
Active
      ge-0/0/13    Partner    No    No   Yes  Yes  Yes   Yes     Fast
Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-0/0/12                 Current   Fast periodic Collecting
distributing
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
      ge-0/0/13                 Current   Fast periodic Collecting
distributing

{master:0}
root@ex-1> show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 596
  Description: Link to EX 2
  Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Current address: 3c:94:d5:46:73:03, Hardware address: 3c:94:d5:46:73:03
  Last flapped   : 2016-03-03 17:48:45 UTC (00:00:25 ago)
  Input rate     : 2048 bps (2 pps)
  Output rate    : 2048 bps (2 pps)

  Logical interface ae0.0 (Index 76) (SNMP ifIndex 503)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol eth-switch
      Flags: Trunk-Mode
```

4) RSTP is configured on all of the switches. This needs to be migrated to MSTP using the information given in table 2.3. EX1 should be the root for region 1 and EX 2 should be the root for region 2.

MST Region	VLANS
1	2000,2010,2020,2030,2040
2	3000,3010,3020,3040

Table 2.3

On EX 1

```
{master:0}[edit]
root@ex-1# show | compare
[edit protocols]
-   rstp {
-       interface all {
-           edge;
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
-       }
-   }
+   mstp {
+       configuration-name JNCIE;
+       revision-level 1;
+       interface ge-0/0/16.0 {
+           mode point-to-point;
+       }
+       interface ge-0/0/18.0 {
+           mode point-to-point;
+       }
+       interface ae0.0 {
+           mode point-to-point;
+       }
+       interface all {
+           edge;
+       }
+       msti 1 {
+           bridge-priority 0;
+           VLAN [ 2000 2010 2020 2030 ];
+       }
+       msti 2 {
+           bridge-priority 4k;
+           VLAN [ 3000 3010 3020 3040 ];
+       }
+   }
```

On EX 2

```
{master:0}[edit]
root@ex-2# show | compare
[edit protocols]
-   rstp {
-       interface all {
-           edge;
-       }
-   }
+   mstp {
+       configuration-name JNCIE;
+       revision-level 1;
+       interface ge-0/0/17.0 {
+           mode point-to-point;
+       }
+       interface ge-0/0/19.0 {
+           mode point-to-point;
+       }
+       interface ae0.0 {
+           mode point-to-point;
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+       }
+       interface all {
+           edge;
+       }
+       msti 1 {
+           bridge-priority 4k;
+           VLAN [ 2000 2010 2020 2030 ];
+       }
+       msti 2 {
+           bridge-priority 0;
+           VLAN [ 3000 3010 3020 3040 ];
+       }
+   }
```

On ex_vc

```
{master:0}[edit]
root@ex_vc# show | compare
[edit protocols]
-   rstp;
+   mstp {
+       configuration-name JNCIE;
+       revision-level 1;
+       interface ge-0/0/11.0 {
+           mode point-to-point;
+       }
+       interface ge-0/0/18.0 {
+           mode point-to-point;
+       }
+       interface ge-1/0/19.0 {
+           mode point-to-point;
+       }
+       interface all {
+           edge;
+       }
+       msti 1 {
+           VLAN [ 2000 2010 2020 2030 ];
+       }
+       msti 2 {
+           VLAN [ 3000 3010 3020 3040 ];
+       }
+   }
```

On srx 3

```
[edit]
root@srx-3# load patch terminal
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[Type ^D at a new line to end input]
[edit protocols]
+   mstp {
+       configuration-name JNCIE;
+       revision-level 1;
+       interface ge-0/0/1.0 {
+           mode point-to-point;
+       }
+       interface ge-0/0/12.0 {
+           mode point-to-point;
+       }
+       interface all {
+           edge;
+       }
+       msti 1 {
+           VLAN [ 2000 2010 2020 2030 ];
+       }
+       msti 2 {
+           VLAN [ 3000 3010 3020 3040 ];
+       }
+   }
```

On srx 4

```
[edit]
root@srx-4# show | compare
[edit protocols]
-   rstp;
+   mstp {
+       configuration-name JNCIE;
+       revision-level 1;
+       interface ge-0/0/1.0 {
+           mode point-to-point;
+       }
+       interface ge-0/0/11.0 {
+           mode point-to-point;
+       }
+       interface ge-0/0/13.0 {
+           mode point-to-point;
+       }
+       interface all {
+           edge;
+       }
+       msti 1 {
+           VLAN [ 2000 2010 2020 2030 ];
+       }
+       msti 2 {
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+           VLAN [ 3000 3010 3020 3040 ];
+       }
+   }
```

Verification

```
{master:0}
root@ex-1> show spanning-tree bridge msti 1

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for MSTI 1
  MSTI regional root                : 1.3c:94:d5:46:73:01
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 5
  Topology change initiator         : ge-0/0/16.0
  Topology change last recvd. from  : 3c:8a:b0:5c:7a:ac
  Local parameters
    Bridge ID                       : 1.3c:94:d5:46:73:01
    Extended system ID              : 0
    Internal instance ID            : 1


{master:0}
root@ex-1> show spanning-tree bridge msti 2

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for MSTI 2
  MSTI regional root                : 2.2c:6b:f5:8f:87:c1
  Root cost                         : 10000
  Root port                         : ae0.0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Number of topology changes        : 6
  Topology change initiator         : ge-0/0/16.0
  Topology change last recvd. from  : 3c:8a:b0:5c:7a:ac
  Local parameters
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Bridge ID                       : 4098.3c:94:d5:46:73:01
    Extended system ID              : 0
    Internal instance ID            : 2

{master:0}
root@ex-2> show spanning-tree bridge msti 2

STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for MSTI 2
  MSTI regional root                : 2.2c:6b:f5:8f:87:c1
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 3
  Topology change initiator         : ae0.0
  Topology change last recvd. from  : 3c:94:d5:46:73:03
  Local parameters
    Bridge ID                       : 2.2c:6b:f5:8f:87:c1
    Extended system ID              : 0
    Internal instance ID            : 2
```

5) Configure the edge ports to block BPDU's on every EX switch. Whenever EX1 or EX2 receives a superior BPDU on a port, that port should be shut down.

On ex 1

```
{master:0}[edit]
root@ex-1# show | compare
[edit protocols mstp interface ge-0/0/16.0]
+    no-root-port;
[edit protocols mstp interface ge-0/0/18.0]
+    no-root-port;
[edit protocols mstp]
+   bpdu-block-on-edge;
```

On ex 2

```
{master:0}[edit]
root@ex-2# show | compare
[edit protocols mstp interface ge-0/0/17.0]
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+    no-root-port;
[edit protocols mstp interface ge-0/0/19.0]
+    no-root-port;
[edit protocols mstp]
+   bpdu-block-on-edge;
```

EX_VC

```
{master:0}[edit]
root@ex_vc# show | compare
[edit protocols mstp]
+   bpdu-block-on-edge;
```

Verification

```
{master:0}
root@ex-1> show spanning-tree interface ge-0/0/16 detail

Spanning tree interface parameters for instance 0

Interface name                 : ge-0/0/16.0
Port identifier                : 128.529
Designated port ID             : 128.525
Port cost                      : 20000
Port state                     : Blocking
Designated bridge ID           : 32768.3c:8a:b0:5c:7a:b0
Port role                      : Alternate (Root-Inconsistent)
Link type                      : Pt-Pt/NONEDGE
Boundary port                  : No
Edge delay while expiry count  : 0
Rcvd info while expiry count   : 1

root@ex-1# set protocols mstp bridge-priority 0

{master:0}
root@ex-1> show spanning-tree interface ge-0/0/16 detail

Spanning tree interface parameters for instance 0

Interface name                 : ge-0/0/16.0
Port identifier                : 128.529
Designated port ID             : 128.529
Port cost                      : 20000
Port state                     : Forwarding
Designated bridge ID           : 0.3c:94:d5:46:73:01
Port role                      : Designated
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Link type                        : Pt-Pt/NONEDGE
Boundary port                    : No
Edge delay while expiry count    : 0
Rcvd info while expiry count     : 2
```

6) SRX 4 should prefer the path via the EX_VC as the most attractive path to the root bridge for instance 1 and 2.

```
root@srx-4> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/1.0     128:514      128:514  32768.3c8ab05c7ab0     20000  BLK    ALT
ge-0/0/7.0     128:520      128:520  32768.3c61048dbe90     20000  FWD    DESG
ge-0/0/8.0     128:521      128:521  32768.3c61048dbe90     20000  FWD    DESG
ge-0/0/11.0    128:524      128:524  32768.2c6bf5351b81     20000  BLK    ALT
ge-0/0/13.0    128:526      128:530  32768.2c6bf58f87c1     20000  FWD    ROOT

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/1.0     128:514      128:514  32769.3c8ab05c7ab0     20000  BLK    ALT
ge-0/0/11.0    128:524      128:524  32769.2c6bf5351b81     20000  BLK    ALT
ge-0/0/13.0    128:526      128:530   4097.2c6bf58f87c1     20000  FWD    ROOT

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/1.0     128:514      128:514  32770.3c61048dbe90     20000  FWD    DESG
ge-0/0/7.0     128:520      128:520  32770.3c61048dbe90     20000  FWD    DESG
ge-0/0/8.0     128:521      128:521  32770.3c61048dbe90     20000  FWD    DESG
ge-0/0/11.0    128:524      128:524  32770.2c6bf5351b81     20000  BLK    ALT
ge-0/0/13.0    128:526      128:530      2.2c6bf58f87c1     20000  FWD    ROOT

[edit]
root@srx-4# show | compare
[edit protocols mstp msti 1]
+     interface ge-0/0/11.0 {
+         cost 1000;
+     }
+     interface ge-0/0/13.0 {
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+         cost 200000;
+     }
[edit protocols mstp msti 2]
+     interface ge-0/0/11.0 {
+         cost 1000;
+     }
+     interface ge-0/0/13.0 {
+         cost 200000;
+     }
```

Verification

```
[edit]
root@srx-4# run show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/1.0     128:514      128:514  32768.3c8ab05c7ab0     20000  BLK    ALT
ge-0/0/7.0     128:520      128:520  32768.3c61048dbe90     20000  FWD    DESG
ge-0/0/8.0     128:521      128:521  32768.3c61048dbe90     20000  FWD    DESG
ge-0/0/11.0    128:524      128:524  32768.2c6bf5351b81     20000  BLK    ALT
ge-0/0/13.0    128:526      128:530  32768.2c6bf58f87c1     20000  FWD    ROOT

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/1.0     128:514      128:514  32769.3c8ab05c7ab0     20000  BLK    ALT
ge-0/0/11.0    128:524      128:524  32769.2c6bf5351b81      1000  FWD    ROOT
ge-0/0/13.0    128:526      128:530   4097.2c6bf58f87c1    200000  BLK    ALT

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/1.0     128:514      128:514  32770.3c61048dbe90     20000  FWD    DESG
ge-0/0/7.0     128:520      128:520  32770.3c61048dbe90     20000  FWD    DESG
ge-0/0/8.0     128:521      128:521  32770.3c61048dbe90     20000  FWD    DESG
ge-0/0/11.0    128:524      128:524  32770.2c6bf5351b81      1000  FWD    ROOT
ge-0/0/13.0    128:526      128:530      2.2c6bf58f87c1    200000  BLK    ALT
```

7) Ensure SRX3 uses interface ge-0/0/12 to reach the CIST root bridge.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit]
root@srx-3# show | compare
[edit protocols mstp interface ge-0/0/12.0]
+    cost 2000;
```

Verification

```
[edit]
root@srx-3# run show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/1.0     128:514      128:514  32768.3c8ab05c7ab0     20000  FWD    DESG
ge-0/0/7.0     128:520      128:520  32768.3c8ab05c7ab0     20000  FWD    DESG
ge-0/0/12.0    128:525      128:529      0.3c94d5467301      2000  FWD    ROOT
ge-0/0/15.0    128:528      128:528  32768.3c8ab05c7ab0     20000  FWD    DESG
```

8) Configure VLAN 4010. Because the VLAN is used only for the firewall cluster, it only needs to be active on EX1 and EX2. Make sure that if ae0 is not working, the will VLAN still offer end-to-end reachability via the EX_VC. You are not allowed to add the VLANs to EX_VC. Create a new msti instance 3 for the VLAN on EX1, EX2 and the EX_VC. Make sure that EX1 is the root.

On ex 1

{master:0}[edit]

```
root@ex-1# show | compare
[edit protocols mstp]
+    msti 3 {
+        VLAN 4010;
+    }
[edit protocols]
+   mvrp {
+       interface ge-0/0/18.0;
+   }
[edit VLANs]
+   firewall {
+       VLAN-id 4010;
+   }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On ex 2

```
{master:0}[edit]
root@ex-2# show | compare
[edit protocols mstp]
+    msti 3 {
+        VLAN 4010;
+    }
[edit protocols]
+   mvrp {
+       interface ge-0/0/19.0;
+   }
[edit VLANs]
+   firewall {
+       VLAN-id 4010;
+   }
```

On ex_vc

```
{master:0}[edit]
root@ex_vc# show | compare
        [edit protocols]
+   mvrp {
+       interface ge-0/0/18.0;
+       interface ge-1/0/19.0;
+   }
```

Verification

```
{master:0}
root@ex_vc> show mvrp dynamic-VLAN-memberships
MVRP dynamic VLANs for routing instance 'default-switch'
 (s) static VLAN, (f) fixed registration


    VLAN ID         Interfaces
    99              ge-0/0/18.0
    1000(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
    2000(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
    2010(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    2020(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
    2030(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
    2040(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
    3000(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
    3010(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
    3020(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
    3040(s)         ge-0/0/18.0(f)
                    ge-1/0/19.0(f)
    4010            ge-0/0/18.0
                    ge-1/0/19.0
```

```
root@ex-1# run show spanning-tree interface msti 3

Spanning tree interface parameters for instance 3

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae0.0            128:1        128:1  32771.2c6bf58f87c1     10000  FWD    ROOT
ge-0/0/18.0    128:531      128:531  32771.3c94d5467301     20000  FWD    DESG

{master:0}

root@ex-2> show spanning-tree interface msti 3

Spanning tree interface parameters for instance 3

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae0.0            128:1        128:1  32771.2c6bf58f87c1     10000  FWD    DESG
ge-0/0/19.0    128:532      128:532  32771.2c6bf58f87c1     20000  FWD    DESG
```

9) Hosts Ha and Hb are using VLAN 3041. They need to be able to communicate using VLAN 3020. Do not make any changes to the VR devices.

On ex_vc

```
{master:0}[edit]
root@ex_vc# show | compare
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit interfaces ge-1/0/23]
+    unit 0 {
+        family ethernet-switching {
+            port-mode access;
+            VLAN {
+                members DMZ-2;
+            }
+        }
+    }
[edit]
+  ethernet-switching-options {
+      dot1q-tunneling {
+          ether-type 0x8100;
+      }
+  }
[edit VLANs DMZ-2]
+    dot1q-tunneling;
```

On srx-3

```
[edit]
root@srx-3# show | compare
[edit interfaces ge-0/0/15 unit 0 family ethernet-switching]
+      port-mode access;
+      VLAN {
+          members DMZ-2;
+      }
[edit]
+  ethernet-switching-options {
+      dot1q-tunneling {
+          ether-type 0x8100;
+      }
+  }
[edit VLANs DMZ-2]
+    dot1q-tunneling;
```

Verification

```
{master:0}[edit]
root@ex_vc# run ping 21.30.41.1 routing-instance Hb_VR rapid
PING 21.30.41.1 (21.30.41.1): 56 data bytes
!!!!!
--- 21.30.41.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
round-trip min/avg/max/stddev = 0.674/6.172/27.619/10.724 ms

{master:0}
root@ex_vc> show ethernet-switching interfaces ge-1/0/23 detail
Interface: ge-1/0/23.0, Index: 85, State: up, Port mode: Access
Ether type for the interface: 0x8100
VLAN membership:
    DMZ-2, 802.1Q Tag: 3020, dot1q-tunneled, untagged, msti-id: 2, unblocked
Number of MACs learned on IFL: 1
```

10) VLAN 3000 needs to offer high availability in case one of the core switches fails. Configure a FHRP using 21.30.0.254 as the VIP. The master (EX1) router should fail over when the link to the EX_VC is down. Users should be able to ping the VIP. For security reasons, use the password "jnciefhrp". Make sure that the password is not send in plain text.

On ex 1

{master:0}[edit]

```
root@ex-1# show | compare
[edit interfaces VLAN unit 3000 family inet address 21.30.0.253/24]
+        vrrp-group 30 {
+            virtual-address 21.30.0.254;
+            priority 200;
+            preempt;
+            accept-data;
+            authentication-type md5;
+            authentication-key "$9$vIF8xdDjqf5F24UH.P3n0BIhlMWLN"; ## SECRET-
DATA
+            track {
+                interface ge-0/0/18 {
+                    priority-cost 60;
+                }
+            }
+        }
```

On ex2

```
{master:0}[edit]
root@ex-2# show | compare
[edit interfaces VLAN unit 3000 family inet address 21.30.0.252/24]
+        vrrp-group 30 {
+            virtual-address 21.30.0.254;
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+            priority 150;
+            preempt;
+            accept-data;
+            authentication-type md5;
+            authentication-key "$9$vIF8xdDjqf5F24UH.P3n0BIhlMWLN"; ## SECRET-
DATA
+            track {
+                interface ge-0/0/19{
+                    priority-cost 60;
+                }
+            }
+        }
```

Verification

```
{master:0}
root@ex-2> show vrrp brief
Interface     State       Group   VR state VR Mode   Timer    Type   Address
vlan.3000     up             30   backup   Active      D  3.185 lcl
21.30.0.252
                                                                 vip
21.30.0.254
                                                                 mas
21.30.0.253

{master:0}
root@ex-2> show vrrp detail
Physical interface: VLAN, Unit: 3000, Address: 21.30.0.252/24
  Index: 79, SNMP ifIndex: 731, VRRP-Traps: disabled
  Interface state: up, Group: 30, State: backup, VRRP Mode: Active
  Priority: 150, Advertisement interval: 1, Authentication type: md5
  Advertisement threshold: 3, Delay threshold: 100, Computed send rate: 0
  Preempt: yes, Accept-data mode: yes, VIP count: 1, VIP: 21.30.0.254
  Dead timer: 2.968s, Master priority: 200, Master router: 21.30.0.253
  Virtual router uptime: 00:01:29
  Tracking: enabled
    Current priority: 150, Configured priority: 150
    Priority hold time: disabled
    Interface tracking: enabled, Interface count: 1
      Interface     Int state   Int speed   Incurred priority cost
      ge-0/0/19.0   up                1g                         0
    Route tracking: disabled
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On failure

```
{master:0}
root@ex-2> show vrrp detail
Physical interface: VLAN, Unit: 3000, Address: 21.30.0.252/24
  Index: 79, SNMP ifIndex: 731, VRRP-Traps: disabled
  Interface state: up, Group: 30, State: backup, VRRP Mode: Active
  Priority: 90, Advertisement interval: 1, Authentication type: md5
  Advertisement threshold: 3, Delay threshold: 100, Computed send rate: 0
  Preempt: yes, Accept-data mode: yes, VIP count: 1, VIP: 21.30.0.254
  Dead timer: 2.922s, Master priority: 200, Master router: 21.30.0.253
  Virtual router uptime: 00:00:37
  Tracking: enabled
    Current priority: 90, Configured priority: 150
    Priority hold time: disabled
    Interface tracking: enabled, Interface count: 1
      Interface     Int state   Int speed   Incurred priority cost
      ge-0/0/19.0   unknown             0                      60
    Route tracking: disabled
```

11) H5, a Wi-Fi access-point, is connected to port ge-1/0/11 on the EX_VC. You need to set up dot1x to make sure that the port is only used by the Wi-Fi access-point. This should authenticate the port based on the MAC address only. It should also allow other users to authenticate as well. For authentication, use server s1 with the key “jncie”. Traffic on the interface should be prevented when authentication fails. Add MAC address “40:60:10:ee:ff:00” to a white list.

```
{master:0}[edit]
root@ex_vc# show | compare
[edit interfaces]
+   ge-1/0/11 {
+       apply-groups-except interface-defaults;
+       description Wi-Fi;
+       unit 0 {
+           family ethernet-switching {
+               port-mode access;
+               VLAN {
+                   members Wi-Fi;
+               }
+           }
+       }
+   }
[edit protocols]
+   dot1x {
+       authenticator {
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+           authentication-profile-name dot1x;
+           static {
+               40:60:10:ee:ff:00/48 {
+                   interface ge-1/0/11.0;
+               }
+           }
+           interface {
+               ge-1/0/11.0 {
+                   supplicant multiple;
+                   MAC-radius {
+                       restrict;
+                   }
+                   server-fail deny;
+               }
+           }
+       }
+   }
[edit]
+  access {
+      radius-server {
+          10.10.10.1 secret "$9$W-vXNbiHmTQn4ajq"; ## SECRET-DATA
+      }
+      profile dot1x {
+          authentication-order radius;
+          radius {
+              authentication-server 10.10.10.1;
+          }
+      }
+  }
```

Verification

```
[edit]
root@srx-3# run ping 21.20.10.253 routing-instance H5_VR rapid
PING 21.20.10.253 (21.20.10.253): 56 data bytes
!!!!!
--- 21.20.10.253 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 2.111/4.294/8.993/2.522 ms

{master:0}
root@ex_vc> show dot1x static-MAC-address
MAC address/prefix          VLAN-Assignment          Interface
40:60:10:ee:ff:00/48                                 ge-1/0/11.0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

12) H6 is a printer using MAC address 30:20:10:aa:bb:cc. This MAC address should only be allowed to talk to 00:00:5e:00:01:20. Traffic from the printer to any other MAC address should be denied. Once this task is completed, you should be able to ping the gateway 21.20.10.254 from H6.

20 HEX = (2 x 16^1) + (0 x 16^0) = vrrp group 32

On EX1

```
root@ex-1# show | compare
[edit interfaces VLAN unit 2020 family inet address 21.20.20.253/24]
+        vrrp-group 32 {
+            virtual-address 21.20.20.254;
+            priority 200;
+            accept-data;
+        }
```

On EX 2

```
root@ex-2# show | compare
[edit interfaces VLAN unit 2020 family inet address 21.20.20.252/24]
+        vrrp-group 32 {
+            virtual-address 21.20.20.254;
+            priority 150;
+            accept-data;
+        }
```

On EX_VC

```
{master:0}[edit]
root@ex_vc# show | compare
[edit]
+  firewall {
+      family ethernet-switching {
+          filter printer-VLAN {
+              term 1 {
+                  from {
+                      source-MAC-address {
+                          30:20:10:aa:bb:cc/48;
+                      }
+                      destination-MAC-address {
+                          00:00:5e:00:01:20/48;
+                      }
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+                 }
+                 then accept;
+             }
+             term 2 {
+                 then discard;
+             }
+         }
+     }
+ }
[edit interfaces ge-0/0/8 unit 0 family ethernet-switching filter]
+       input printer-VLAN;
```

Verification

Before

```
[edit]
root@srx-3# run ping 21.20.20.253 routing-instance H6_VR
PING 21.20.20.253 (21.20.20.253): 56 data bytes
64 bytes from 21.20.20.253: icmp_seq=0 ttl=64 time=3.365 ms
64 bytes from 21.20.20.253: icmp_seq=1 ttl=64 time=2.076 ms
64 bytes from 21.20.20.253: icmp_seq=2 ttl=64 time=2.128 ms
64 bytes from 21.20.20.253: icmp_seq=3 ttl=64 time=1.695 ms
```

After

```
[edit]
root@srx-3# run ping 21.20.20.253 routing-instance H6_VR
PING 21.20.20.253 (21.20.20.253): 56 data bytes
^C
--- 21.20.20.253 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss

[edit]
root@srx-3# run show arp | match 21.20.20.253
3c:94:d5:46:73:01 21.20.20.253    21.20.20.253              ge-0/0/8.0
none

[edit]
root@srx-3# run ping 21.20.20.254 routing-instance H6_VR
PING 21.20.20.254 (21.20.20.254): 56 data bytes
64 bytes from 21.20.20.254: icmp_seq=0 ttl=64 time=2.197 ms
64 bytes from 21.20.20.254: icmp_seq=1 ttl=64 time=1.787 ms
64 bytes from 21.20.20.254: icmp_seq=2 ttl=64 time=1.827 ms
64 bytes from 21.20.20.254: icmp_seq=3 ttl=64 time=3.934 ms
64 bytes from 21.20.20.254: icmp_seq=4 ttl=64 time=1.730 ms
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
^C
--- 21.20.20.254 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.730/2.295/3.934/0.836 ms

[edit]
root@srx-3# run show arp | match 21.20.20.254
00:00:5e:00:01:20 21.20.20.254    21.20.20.254             ge-0/0/8.0
none
```

13) Configure sflow to sample 1 out of every 100 packets on EX1. Configure server s1 as the collector and use port 7000. Add interfaces ge-0/0/18 to sflow.

```
root@ex-1# show | compare
[edit protocols]
+   sflow {
+       sample-rate 100;
+       collector 10.10.10.1 {
+           udp-port 7000;
+       }
+       interfaces ge-0/0/18.0;
+   }
```

Verification

```
{master:0}
root@ex-1> show sflow collector
Collector         Udp-port    No. of samples
 address
10.10.10.1        7000        1

{master:0}
root@ex-1> show sflow interface
Interface         Status              Sample              Adapted
Polling
                                       rate             sample rate
interval

                Egress  Ingress     Egress   Ingress      Egress   Ingress
ge-0/0/18.0      Disabled Enabled   2048     100          2048     100
20
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

14) Make sure that the EX_VC synchronizes all layer 2 protocols between the members. Also make sure that the master RE switches over gracefully when the Back-up becomes master.

```
{master:0}[edit]
root@ex_vc# show | compare
[edit system]
+    commit synchronize;
[edit chassis]
+    redundancy {
+        graceful-switchover;
+    }
[edit ethernet-switching-options]
+   nonstop-bridging;
```

Verification

```
{backup:1}
root@ex3> show system switchover
fpc1:
--------------------------------------------------------------------------
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 3: Protocol Independent Routing

1) Create an ICMP ping probe that sends 3 probes. One probe should be send every 10 seconds and the test should run every 60 seconds. You should ping 8.8.8.8 using source address 8.254.157.2.

```
 [edit]
+  services {
+      rpm {
+          probe T3 {
+              test ping-T3 {
+                  probe-type ICMP-ping;
+                  target address 8.8.8.8;
+                  probe-count 3;
+                  probe-interval 10;
+                  test-interval 60;
+                  source-address 8.254.157.2;
+                  thresholds {
+                      successive-loss 1;
+                  }
+              }
+          }
+      }
+  }
```

Verification

```
[edit]
root@srx-1-re-0# run show services rpm probe-results
    Owner: T3, Test: ping-T3
    Target address: 8.8.8.8, Source address: 8.254.157.2, Probe type: icmp-
ping,
    Test size: 3 probes
    Probe results:
      Response received, Sat May 14 00:05:19 2016, No hardware timestamps
      Rtt: 971 usec
    Results over current test:
      Probes sent: 3, Probes received: 3, Loss percentage: 0
      Measurement: Round trip time
        Samples: 3, Minimum: 971 usec, Maximum: 1248 usec, Average: 1138 usec,
        Peak to peak: 277 usec, Stddev: 120 usec, Sum: 3414 usec
    Results over last test:
      Probes sent: 3, Probes received: 3, Loss percentage: 0
      Test completed on Sat May 14 00:05:19 2016
      Measurement: Round trip time
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
        Samples: 3, Minimum: 971 usec, Maximum: 1248 usec, Average: 1138 usec,
        Peak to peak: 277 usec, Stddev: 120 usec, Sum: 3414 usec
    Results over all tests:
      Probes sent: 228, Probes received: 21, Loss percentage: 90
      Measurement: Round trip time
        Samples: 21, Minimum: 971 usec, Maximum: 7228 usec, Average: 1806
usec,
        Peak to peak: 6257 usec, Stddev: 1763 usec, Sum: 37927 usec
```

2) Create a default route on R2. This route should be available only if the peering connection with Transit 1 is up. This route should be advertised into OSPF with a metric of 10.

```
[edit]
root@srx-2# show | compare
 [edit routing-options]
+   generate {
+       route 0.0.0.0/0 policy Transit-1;
+   }
[edit protocols OSPF]
-   export ospf3-to-ospf;
+   export [ ospf3-to-ospf OSPF-export ];
[edit policy-options]
+   policy-statement Transit-1 {
+       term 1 {
+           from {
+               protocol BGP;
+               neighbor 6.6.1.1;
+           }
+           then {
+               next-hop 6.6.1.1;
+               accept;
+           }
+       }
+   }
+   policy-statement OSPF-export {
+       term 1 {
+           from {
+               protocol aggregate;
+               route-filter 0.0.0.0/0 exact;
+           }
+           then {
+               metric 10;
+               accept;
+           }
+       }
+       term 2 {
+        then reject;
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+        }
+    }
```

Verification

```
root@srx-2> show route protocol aggregate 0/0 extensive

inet.0: 157 destinations, 398 routes (156 active, 0 holddown, 62 hidden)
0.0.0.0/0 (3 entries, 1 announced)
TSI:
KRT in-kernel 0.0.0.0/0 -> {indirect(262143)}
unicast reverse-path: 760
  [ge-0/0/4.101 ge-0/0/5.102]
OSPF area : 0.0.0.0, LSA ID : 0.0.0.0, LSA type : Extern
        *Aggregate Preference: 130
                Next hop type: Indirect
                Address: 0x934da80
                Next-hop reference count: 211
                Next hop type: Router, Next hop index: 262142
                Next hop: 6.6.6.1 via ge-0/0/4.101, selected
                Next hop: 6.7.7.1 via ge-0/0/5.102
                Protocol next hop: 6.6.1.1
                Indirect next hop: 95e09f8 262143
                State: <Active Int Ext>
                Local AS:   653
                Age: 1:17       Metric2: 0
                Task: Aggregate
                Announcement bits (4): 0-KRT 2-OSPF 10-KRT 12-Resolve tree 4
                AS path: I
                                Flags: Generate Resolve Depth: 1        Active
                Contributing Routes (58):
                        6.0.0.0/8 proto BGP
                        6.6.0.0/24 proto BGP
                        6.6.1.0/24 proto BGP
                        6.6.2.0/24 proto BGP
                        6.6.3.0/24 proto BGP
                        6.6.4.0/24 proto BGP
                        6.21.0.0/16 proto BGP
                        6.22.0.0/16 proto BGP
                        6.23.0.0/16 proto BGP
                        6.24.0.0/16 proto BGP
                        6.214.0.0/16 proto BGP
                        1.0.0.0/8 proto BGP
                        1.1.0.0/24 proto BGP
                        1.1.1.0/24 proto BGP
                        1.1.2.0/24 proto BGP
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                1.1.3.0/24 proto BGP
                1.1.4.0/24 proto BGP
                1.1.5.0/24 proto BGP
                1.2.0.0/16 proto BGP
                1.2.5.0/24 proto BGP
                1.3.0.0/16 proto BGP
                1.3.5.0/24 proto BGP
                1.4.0.0/16 proto BGP
                1.4.5.0/24 proto BGP
                2.0.0.0/8 proto BGP
                2.1.0.0/16 proto BGP
                2.2.1.0/24 proto BGP
                2.2.3.0/24 proto BGP
                2.2.4.0/24 proto BGP
                2.2.5.0/24 proto BGP
                2.3.0.0/16 proto BGP
                2.4.0.0/16 proto BGP
                3.0.0.0/8 proto BGP
                3.1.0.0/16 proto BGP
                3.2.0.0/16 proto BGP
                3.3.1.0/24 proto BGP
                3.3.2.0/24 proto BGP
                3.3.4.0/24 proto BGP
                3.4.0.0/16 proto BGP
                3.5.0.0/16 proto BGP
                3.6.0.0/16 proto BGP
                3.7.0.0/16 proto BGP
                3.200.0.0/16 proto BGP
                3.201.0.0/16 proto BGP
                3.202.0.0/16 proto BGP
                4.0.0.0/8 proto BGP
                4.1.0.0/16 proto BGP
                4.2.0.0/16 proto BGP
                4.3.0.0/16 proto BGP
                4.4.0.0/24 proto BGP
                4.4.1.0/24 proto BGP
                4.4.2.0/24 proto BGP
                4.4.3.0/24 proto BGP
                4.4.5.0/24 proto BGP
                4.52.0.0/16 proto BGP
                4.53.0.0/16 proto BGP
                4.54.0.0/16 proto BGP
                4.55.0.0/16 proto BGP
        Indirect next hops: 1
                Protocol next hop: 6.6.1.1
                Indirect next hop: 95e09f8 262143
                Indirect path forwarding next hops: 2
                        Next hop type: Router
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                                    Next hop: 6.6.6.1 via ge-0/0/4.101
                                    Next hop: 6.7.7.1 via ge-0/0/5.102
                                6.6.1.1/32 Originating RIB: inet.0
                                  Node path count: 1
                                  Forwarding nexthops: 2
                                    Nexthop: 6.6.6.1 via ge-0/0/4.101
                                    Nexthop: 6.7.7.1 via ge-0/0/5.102

root@srx-1-re-0> show route

inet.0: 150 destinations, 344 routes (150 active, 0 holddown, 5 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 00:02:50, metric 10, tag 0
                    > to 21.1.10.6 via ge-0/0/1.70
```

3) All RIP IPv4 learned routes should be forwarded towards R1 by default. This should not affect the route that the rest of the traffic takes. You can use a single static route for this task

On R3

```
[edit]
root@srx-3# show | compare
[edit interfaces ge-0/0/5 unit 46 family inet]
+       filter {
+           input rip-routes;
+       }
[edit]
+  routing-options {
+      rib-groups {
+          fw-rip {
+              import-rib [ R3_B.inet.0 fw-rip.inet.0 ];
+          }
+      }
+  }
+  firewall {
+      family inet {
+          filter rip-routes {
+              term 1 {
+                  from {
+                      source-address {
+                          21.1.96.0/21;
+                      }
+                  }
+                  then {
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+                        routing-instance fw-rip;
+                    }
+                }
+                term 2 {
+                    then accept;
+                }
+            }
+        }
+    }
[edit routing-instances R3_B routing-options]
+     interface-routes {
+         rib-group inet fw-rip;
+     }
[edit routing-instances R3_B protocols ospf3 realm ipv4-unicast]
+      rib-group fw-rip;
[edit routing-instances]
+   fw-rip {
+       apply-groups-except BGP-defaults;
+       instance-type forwarding;
+       routing-options {
+           static {
+               route 0.0.0.0/0 next-hop 21.1.10.37;
+           }
+       }
+   }
```

On R4

```
[edit]
root@srx-4# show | compare
[edit interfaces ge-0/0/5 unit 45 family inet]
+       filter {
+           input rip-routes;
+       }
[edit]
+  routing-options {
+      rib-groups {
+          fw-rip {
+              import-rib [ R4_B.inet.0 fw-rip.inet.0 ];
+          }
+      }
+  }
+  firewall {
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+      family inet {
+          filter rip-routes {
+              term 1 {
+                  from {
+                      source-address {
+                          21.1.96.0/21;
+                      }
+                  }
+                  then {
+                      routing-instance fw-rip;
+                  }
+              }
+              term 2 {
+                  then accept;
+              }
+          }
+      }
+  }
[edit routing-instances R4_B routing-options]
+     interface-routes {
+         rib-group inet fw-rip;
+     }
[edit routing-instances R4_B protocols ospf3 realm ipv4-unicast]
+      rib-group fw-rip;
[edit routing-instances]
+   fw-rip {
+       apply-groups-except BGP-defaults;
+       instance-type forwarding;
+       routing-options {
+           static {
+               route 0.0.0.0/0 next-hop 21.1.10.33;
+           }
+       }
+   }
```

Verification

Before Policy

```
root@VR_Router_Lab3> traceroute routing-instance VR_2 8.8.8.8
traceroute to 8.8.8.8 (8.8.8.8), 30 hops max, 40 byte packets
 1  21.1.99.1 (21.1.99.1)  8.806 ms  9.348 ms  9.852 ms
 2  21.1.10.49 (21.1.10.49)  14.290 ms  14.238 ms  19.839 ms
 3  21.1.10.17 (21.1.10.17)  14.178 ms  14.394 ms  19.867 ms
 4  8.8.8.8 (8.8.8.8)  29.579 ms  22.760 ms  29.435 ms
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

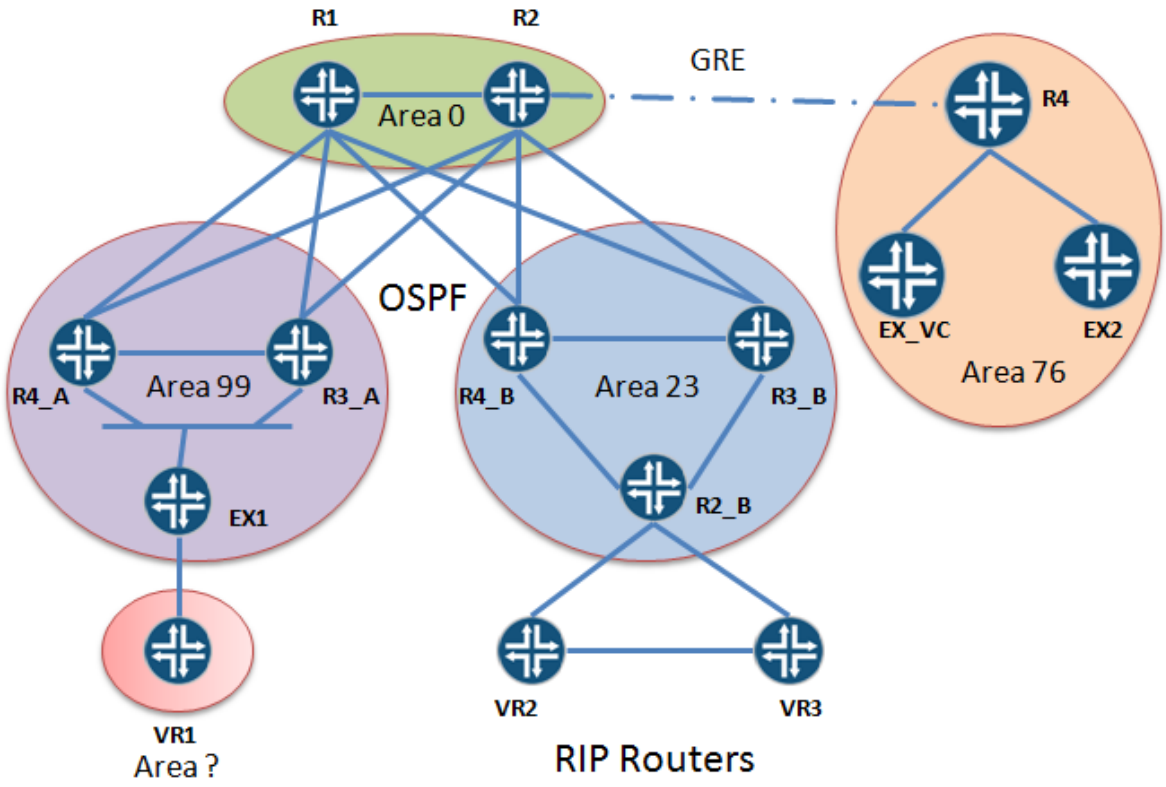
After policy

```
root@VR_Router_Lab3> traceroute routing-instance VR_2 8.8.8.8
traceroute to 8.8.8.8 (8.8.8.8), 30 hops max, 40 byte packets
 1  21.1.99.1 (21.1.99.1)  8.976 ms  9.225 ms  9.416 ms
 2  21.1.10.49 (21.1.10.49)  9.270 ms  14.115 ms  15.167 ms
 3  21.1.10.33 (21.1.10.33)  13.833 ms  14.558 ms  14.003 ms
 4  21.1.10.5 (21.1.10.5)  14.771 ms  12.985 ms  14.467 ms
 5  8.8.8.8 (8.8.8.8)  24.976 ms  24.143 ms  19.477 ms
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 4: IGP Routing



Device Name	Interface A-End	Device Name B-End	Interface B-End	IP Address	IPv6 Address	Area	Protocol
SRX1	Ge-0/0/1.69	SRX2	Ge-0/0/1.69	21.1.10.0/30	2002:95:b::a653:0/126	0	ospfv3
SRX1	Ge-0/0/1.70	SRX2	Ge-0/0/1.70	21.1.10.4/30	-	0	ospfv2
SRX1	Ge-0/0/5.30	R4_A	Ge-0/0/5.30	21.1.10.8/30	2002:95:b::a653:8/126	0	ospfv2,v3
SRX1	Ge-0/0/5.31	R3_A	Ge-0/0/5.31	21.1.10.12/30	2002:95:b::a653:c/126	0	ospfv2,v3
SRX1	Ge-0/0/4.40	R4_B	Ge-0/0/4.40	21.1.10.16/30	2002:95:b::a653:10/126	0	ospfv3
SRX1	Ge-0/0/4.41	R3_B	Ge-0/0/4.41	21.1.10.20/30	2002:95:b::a653:14/126	0	ospfv3
SRX2	Ge-0/0/5.32	R4_A	Ge-0/0/5.32	21.1.10.24/30	2002:95:b::a653:18/126	0	ospfv2,v3
SRX2	Ge-0/0/5.33	R3_A	Ge-0/0/5.33	21.1.10.28/30	2002:95:b::a653:1c/126	0	ospfv2,v3
SRX2	Ge-0/0/4.42	R4_B	Ge-0/0/4.42	21.1.10.32/30	2002:95:b::a653:20/126	0	ospfv3
SRX2	Ge-0/0/4.43	R3_B	Ge-0/0/4.43	21.1.10.36/30	2002:95:b::a653:24/126	0	ospfv3
R4_A	Ge-0/0/4.34	R3_A	Ge-0/0/4.34	21.1.10.40/30	2002:95:b::a653:28/126	99	ospfv2,v3
R4_A	Ge-0/0/12	EX1	Vlan 99	21.51.42.0/24	2002:95:b::42:0/120	99	ospfv2,v3
R3_A	Vlan.99	-	-	21.51.42.0/24	2002:95:b::42:0/120	99	ospfv2,v3
EX1	Ge-0/0/15.68	VR1	-	21.55.145.0/30	-	?	ospfv2
R4_B	Ge-0/0/5.44	R3_B	Ge-0/0/5.44	21.1.10.44/30	2002:95:b::a653:2c/126	23	ospfv3

JNCIE-ENT Workbook: Chapter 3: OSPF Challenge

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

R4_B	Ge-0/0/5.45	R2_B	Ge-0/0/5.45	21.1.10.48/30	2002:95:b::a653:30/126	23	ospfv3
R3_B	Ge-0/0/5.46	R2_B	Ge-0/0/5.46	21.1.10.52/30	2002:95:b::a653:34/126	23	ospfv3
R2_B	Ge-0/0/4.62	VR2	-	21.1.99.0/30	2002:95:b::f1:0/126	-	Rip,Ripng
R2_B	Ge-0/0/5.64	VR3	-	21.1.99.4/30	2002:95:b::f1:4/126	-	Rip,Ripng
SRX 2	Ge-0/0/4.66	VR4	-	21.1.76.0/30	-	-	BGP
SRX 4	Ge-0/0/4.67	VR4	-	21.1.76.4/30	-	76	ospfv2
SRX 4	Ge-0/0/4.69	EX_VC	Ge-0/0/14.69	21.1.76.8/30	-	76	ospfv2
SRX 4	Ge-0/0/4.70	EX2	Ge-0/0/14.70	21.1.76.12/30		76	ospfv2

Table 4.1

1) After a number of changes on both SRX1 and SRX2, the OSPF neighbor relationship between them is down. Fix this issue.

OSPFv2 Verification

On R1

```
root@srx-1-re-0> show OSPF neighbor interface ge-0/0/1.70

root@srx-1-re-0> show OSPF interface ge-0/0/1.70 detail
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/1.70         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0
0
  Type: P2P, Address: 21.1.10.5, Mask: 255.255.255.252, MTU: 1522, Cost: 1
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

On R2

```
root@srx-2> show OSPF neighbor

root@srx-2> show OSPF interface ge-0/0/1.70 detail
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/1.70         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0
0
  Type: P2P, Address: 21.1.10.6, Mask: 255.255.255.252, MTU: 1522, Cost: 1
  Adj count: 0, Passive
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Passive, Cost: 1
root@srx-2# show protocols OSPF area 0.0.0.0
interface ge-0/0/1.70;
interface ge-0/0/5.33;
interface ge-0/0/5.32;
interface lo0.0 {
    passive;
}

root@srx-2# show protocols OSPF area 0.0.0.0 interface ge-0/0/1.70 | display
inheritance
##
## 'p2p' was inherited from group 'OSPF-defaults'
##
interface-type p2p;
##
## 'passive' was inherited from group 'OSPF-defaults'
##
passive;
##
## 'no-neighbor-down-notification' was inherited from group 'OSPF-defaults'
##
no-neighbor-down-notification;

[edit]
root@srx-2# show   compare
[edit groups OSPF-defaults protocols OSPF area <*> interface <*>]
-       passive;
```

Verification

```
root@srx-2> show OSPF neighbor interface ge-0/0/1.70
Address          Interface              State     ID               Pri  Dead
21.1.10.5        ge-0/0/1.70            Full      21.100.0.100     128    37
root@srx-2> show OSPF interface ge-0/0/1.70 detail
Interface          State   Area            DR ID           BDR ID
Nbrs
ge-0/0/1.70        PtToPt  0.0.0.0         0.0.0.0         0.0.0.0
1
  Type: P2P, Address: 21.1.10.6, Mask: 255.255.255.252, MTU: 1522, Cost: 1
  Adj count: 1
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
    Acknowledgement list:  1 entries
```

OSPFv3 Verification

On R1

```
root@srx-1-re-0> show ospf3 neighbor interface ge-0/0/1.69

root@srx-1-re-0> show ospf3 interface ge-0/0/1.69 detail
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/1.69         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0
0
  Address fe80::20c:2900:45a7:4e90, Prefix-length 64
  OSPF3-Intf-index 2, Type P2P, MTU 1522, Cost 1
  Adj count: 0, Router LSA ID: -
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None

root@srx-1-re-0> show ospf3 neighbor realm ipv4-unicast interface ge-0/0/1.69
ID              Interface             State     Pri   Dead
21.100.0.2      ge-0/0/1.69           Full      128     37
  Neighbor-address fe80::20c:2900:450d:3061
```

On R2

```
root@srx-1-re-0> show ospf3 neighbor interface ge-0/0/1.69

root@srx-1-re-0> show ospf3 interface ge-0/0/1.69 detail
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/1.69         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0
0
  Address fe80::20c:2900:45a7:4e90, Prefix-length 64
  OSPF3-Intf-index 2, Type P2P, MTU 1522, Cost 1
  Adj count: 0, Router LSA ID: -
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None

root@srx-1-re-0> ...f3 neighbor realm ipv4-unicast interface ge-0/0/1.69
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
ID              Interface             State     Pri   Dead
21.100.0.2      ge-0/0/1.69           Full      128     35
  Neighbor-address fe80::20c:2900:450d:3061


[edit]
root@srx-2# show | compare
[edit protocols ospf3]
+    traceoptions {
+        file ospf3-debug;
+        flag all;
+    }


root@srx-2> show log ospf3-debug | match "mismatch|ge-0/0/1.69"
Feb 13 00:44:08.987940 OSPF rcvd Hello fe80::20c:2900:45a7:4e90 -> ff02::5
(ge-0/0/1.69 IFL 70 area 0.0.0.0)
Feb 13 00:44:08.987989 OSPF packet ignored: hello interval mismatch 10 from
fe80::20c:2900:45a7:4e90 on intf ge-0/0/1.69 area 0.0.0.0

[edit]
root@srx-2# ... area 0.0.0.0 interface ge-0/0/1.69 | display inheritance
##
## 'p2p' was inherited from group 'ospf3-defaults'
##
interface-type p2p;
##
## '1' was inherited from group 'ospf3-defaults'
##
hello-interval 1;


[edit]
root@srx-2# show | compare
[edit groups ospf3-defaults protocols ospf3 area <*> interface <*>]
-      hello-interval 1;
+      hello-interval 10;
```

Verification

```
[edit]
root@srx-2# run show ospf3 neighbor interface ge-0/0/1.69
ID              Interface             State     Pri   Dead
21.100.0.100    ge-0/0/1.69           Full      128     34
  Neighbor-address fe80::20c:2900:45a7:4e90
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

2) The OSPF neighbor relationship between R2_B and R4_B is not establishing. Resolve this issue without removing any configuration.

On SRX 2

```
root@srx-2> show ospf3 neighbor interface ge-0/0/5.45 instance R2_B

root@srx-2> show ospf3 interface ge-0/0/5.45 detail instance R2_B
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/5.45         PtToPt  0.0.0.23        0.0.0.0         0.0.0.0
0
  Address fe80::20c:2900:2d86:3394, Prefix-length 64
  OSPF3-Intf-index 2, Type P2P, MTU 1522, Cost 1
  Adj count: 0, Router LSA ID: -
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None

root@srx-2# show | compare
[edit routing-instances R2_B protocols ospf3]
+     traceoptions {
+         file ospf3-debug;
+         flag all;
+     }

[edit]
root@srx-2# run show log ospf3-debug | match mismatch
Feb  1 20:40:32.734508 OSPF packet ignored: configuration mismatch from
fe80::20c:2900:2dbf:8eb6 on intf ge-0/0/5.45 area 0.0.0.23
Feb  1 20:40:41.466861 OSPF packet ignored: configuration mismatch from
fe80::20c:2900:2dbf:8eb6 on intf ge-0/0/5.45 area 0.0.0.23
```

On R4_B

```
root@srx-4> ...ospf3 neighbor instance R4_B interface ge-0/0/5.45 detail
ID              Interface             State     Pri   Dead
21.100.0.11     ge-0/0/5.45           Init      128     39
  Neighbor-address fe80::20c:2900:2d86:3394
  Area 0.0.0.23, opt 0x13, OSPF3-Intf-Index 2
  DR-ID 0.0.0.0, BDR-ID 0.0.0.0


root@srx-4> show ospf3 interface ge-0/0/5.45 instance R4_B detail
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/5.45         DR      0.0.0.23        21.100.0.6      0.0.0.0
1
  Address fe80::20c:2900:2dbf:8eb6, Prefix-length 64
  OSPF3-Intf-index 5, Type LAN, MTU 1522, Cost 1, Priority 128
  Adj count: 0, Router LSA ID: -
  DR addr fe80::20c:2900:2dbf:8eb6
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None

[edit]
root@srx-4# show routing-instances R4_B protocols ospf3 area 23 interface ge-
0/0/5.45 | display inheritance

[edit]
root@srx-4#

[edit]
root@srx-4# show | compare
[edit groups ospf3-defaults]
-   protocols {
-       ospf3 {
-           realm ipv4-unicast {
-               area <*> {
-                   interface <*> {
-                       interface-type p2p;
-                       hello-interval 10;
-                   }
-               }
-           }
-           area <*> {
-               interface <*> {
-                   interface-type p2p;
-                   hello-interval 10;
-               }
-           }
-       }
-   }
+   routing-instances {
+       <*> {
+           protocols {
+               ospf3 {
+                   realm ipv4-unicast {
+                       area <*> {
+                           interface <*> {
+                               interface-type p2p;
+                               hello-interval 10;
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+                           }
+                       }
+                   }
+                   area <*> {
+                       interface <*> {
+                           interface-type p2p;
+                           hello-interval 10;
+                       }
+                   }
+               }
+           }
+       }
+   }
```

Verification

```
root@srx-4> show ospf3 neighbor instance R4_B interface ge-0/0/5.45
ID               Interface              State     Pri   Dead
21.100.0.13      ge-0/0/5.45            Full      128     32
  Neighbor-address fe80::20c:2900:2d86:3394

root@srx-4> show ospf3 interface ge-0/0/5.45 instance R4_B detail
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/5.45         PtToPt  0.0.0.23        0.0.0.0         0.0.0.0
1
  Address fe80::20c:2900:2dbf:8eb6, Prefix-length 64
  OSPF3-Intf-index 5, Type P2P, MTU 1522, Cost 1
  Adj count: 1, Router LSA ID: 0
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None
```

3) One of the engineers has noticed that the session between SRX1 and R4_A OSPF is down. Investigate and fix the issue.

```
root@srx-4> show OSPF neighbor instance R4_A interface ge-0/0/5.30 detail

root@srx-4> show OSPF interface ge-0/0/5.30 detail instance R4_A
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/5.30         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0
0
  Type: P2P, Address: 21.1.10.10, Mask: 255.255.255.252, MTU: 1522, Cost: 1
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1

[edit]
root@srx-4# show | compare
[edit routing-instances R4_A protocols OSPF]
+      traceoptions {
+          file OSPF-debug;
+          flag all;
+      }

[edit]
root@srx-4# run show log OSPF-debug | match "packet ignored|mismatch"
Feb 13 01:13:19.703647 OSPF packet ignored: our router ID received from
21.1.10.9 on intf ge-0/0/5.30 area 0.0.0.0

[edit]
root@srx-4# show | compare
[edit routing-instances R4_A]
+    routing-options {
+        router-id 21.100.0.4;
+    }
```

Verification

```
[edit]
root@srx-4# run clear OSPF neighbor instance R4_A

[edit]
root@srx-4# run show OSPF neighbor instance R4_A interface ge-0/0/5.30
Address          Interface              State     ID               Pri  Dead
21.1.10.9        ge-0/0/5.30            Full      21.100.0.100     128    32
```

4) Area 23 should be set up as an NSSA. No Type 3 LSA's should be allowed inside this NSSA.

On R2_B

```
[edit]
root@srx-2# show | compare
[edit routing-instances R2_B protocols ospf3 realm ipv4-unicast area 0.0.0.23]
+        nssa;
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit routing-instances R2_B protocols ospf3 area 0.0.0.23]
+       nssa;
```

On R4_B

```
 [edit]
root@srx-4# show | compare
[edit routing-instances R4_B protocols ospf3 realm ipv4-unicast area 0.0.0.23]
+        nssa no-summaries;
[edit routing-instances R4_B protocols ospf3 area 0.0.0.23]
+       nssa no-summaries;
```

On R3_B

```
root@srx-3# show | compare
[edit routing-instances R3_B protocols ospf3 realm ipv4-unicast area 0.0.0.23]
+        nssa no-summaries;
[edit routing-instances R3_B protocols ospf3 area 0.0.0.23]
+       nssa no-summaries;
```

Verification

```
root@srx-2> show ospf3 interface detail instance R2_B
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/5.45         PtToPt  0.0.0.23        0.0.0.0         0.0.0.0
1
  Address fe80::20c:2900:2d86:3394, Prefix-length 64
  OSPF3-Intf-index 2, Type P2P, MTU 1522, Cost 1
  Adj count: 1, Router LSA ID: 0
  Hello 10, Dead 40, ReXmit 5, Stub NSSA
  Protection type: None
ge-0/0/5.46         PtToPt  0.0.0.23        0.0.0.0         0.0.0.0
1
  Address fe80::20c:2900:2e86:3394, Prefix-length 64
  OSPF3-Intf-index 3, Type P2P, MTU 1522, Cost 1
  Adj count: 1, Router LSA ID: 0
  Hello 10, Dead 40, ReXmit 5, Stub NSSA
  Protection type: None

root@srx-2> show ospf3 database inter-area-prefix instance R2_B

root@srx-2>
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx-4> show ospf3 database inter-area-prefix instance R4_B

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
InterArPfx  0.0.0.1          21.100.0.4       0x80000001   597  0x2048  44
InterArPfx  0.0.0.2          21.100.0.4       0x80000002   563  0xff66  44
InterArPfx  0.0.0.3          21.100.0.4       0x80000002   564  0x5fe9  44
InterArPfx  0.0.0.1          21.100.0.5       0x8000000c   191  0x1e3d  44
InterArPfx  0.0.0.2          21.100.0.5       0x8000000a   191  0xbb81  44
InterArPfx  0.0.0.3          21.100.0.5       0x8000000a   191  0x82b1  44
InterArPfx  0.0.0.4          21.100.0.5       0x8000000a   181  0xa6ac  44
InterArPfx  0.0.0.5          21.100.0.5       0x8000000a   181  0xc29   44
InterArPfx  0.0.0.6          21.100.0.5       0x80000003   188  0x1e3f  44
InterArPfx *0.0.0.1          21.100.0.6       0x8000000c   186  0x3227  44
InterArPfx *0.0.0.2          21.100.0.6       0x8000000d   186  0xaf89  44
InterArPfx *0.0.0.3          21.100.0.6       0x8000000c   186  0x1025  44
InterArPfx *0.0.0.4          21.100.0.6       0x80000003   186  0x124d  44
InterArPfx *0.0.0.5          21.100.0.6       0x80000002   180  0xa6b2  44
InterArPfx *0.0.0.6          21.100.0.6       0x80000004   180  0x70c4  44
```

5) Advertise a default route in to Area 23 for both the IPv4 as well as the IPv6 table. All IPv4 traffic should go via R4_B and all IPv6 traffic should go via R3_B. There should be a failover when a device is unavailable.

On R4_B

```
[edit routing-instances R4_B protocols ospf3 realm ipv4-unicast area 0.0.0.23
nssa]
+         default-lsa {
+             default-metric 100;
+             metric-type 1;
+             type-7;
+         }
[edit routing-instances R4_B protocols ospf3 area 0.0.0.23 nssa]
+        default-lsa {
+            default-metric 1000;
+            metric-type 1;
+            type-7;
+        }
```

On R3_B

```
root@srx-3# show | compare
[edit routing-instances R3_B protocols ospf3 realm ipv4-unicast area 0.0.0.23
nssa]
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+         default-lsa {
+             default-metric 1000;
+             metric-type 1;
+             type-7;
+         }
[edit routing-instances R3_B protocols ospf3 area 0.0.0.23 nssa]
+        default-lsa {
+            default-metric 100;
+            metric-type 1;
+            type-7;
+        }
```

Verification

```
root@srx-2> show ospf3 database realm ipv4-unicast nssa detail instance R2_B

    OSPF3 database, Area 0.0.0.23
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
NSSA        0.0.0.1          21.100.0.5       0x80000001  1038  0xc096  28
  Prefix 0.0.0.0/0
  Prefix-options 0x0, Metric 1000, Type 1,
NSSA        0.0.0.1          21.100.0.6       0x80000001   941  0xe4e9  28
  Prefix 0.0.0.0/0
  Prefix-options 0x0, Metric 100, Type 1,

root@srx-2> show ospf3 database  nssa detail instance R2_B

    OSPF3 database, Area 0.0.0.23
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
NSSA        0.0.0.1          21.100.0.5       0x80000001  1042  0xeae4  28
  Prefix ::/0
  Prefix-options 0x0, Metric 100, Type 1,
NSSA        0.0.0.1          21.100.0.6       0x80000001   945  0xba9b  28
  Prefix ::/0
  Prefix-options 0x0, Metric 1000, Type 1,

root@srx-2> show route 0/0 detail table R2_B.inet.0

inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
0.0.0.0/0 (1 entry, 1 announced)
        *OSPF3  Preference: 150
                Next hop type: Router, Next hop index: 625
                Address: 0x9350a78
                Next-hop reference count: 4
                Next hop: 21.1.10.49 via ge-0/0/5.45, selected
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                State: <Active Int Ext>
                Age: 1:18       Metric: 200     Tag: 0
                Task: OSPF3
                Announcement bits (2): 0-KRT 6-RIPv2
                AS path: I

root@srx-2> show route ::/0 detail table R2_B.inet.0

inet6.0: 31 destinations, 36 routes (31 active, 0 holddown, 0 hidden)
::/0 (1 entry, 1 announced)
        *OSPF3  Preference: 150
                Next hop type: Router, Next hop index: 623
                Address: 0x93509e8
                Next-hop reference count: 4
                Next hop: fe80::20c:2900:2e2b:dff9 via ge-0/0/5.46, selected
                State: <Active Int Ext>
                Age: 1:55       Metric: 200     Tag: 0
                Task: OSPF3
                Announcement bits (2): 0-KRT 2-RIPng
                AS path: I
```

6) EX1 OSPF sessions are not working. One of the engineers has asked if you can have a look at this. Solve the issue.

On Ex1

```
root@ex-1> show OSPF neighbor
Address          Interface              State     ID               Pri  Dead
21.51.42.252     vlan.99                ExStart   21.100.0.3       128    38
root@ex-1> show OSPF interface vlan.99
Interface           State   Area            DR ID           BDR ID
Nbrs
vlan.99             PtToPt  0.0.0.99        0.0.0.0         0.0.0.0
1

[edit]
root@ex-1# show protocols OSPF
traceoptions {
    file OSPF-debug;
    flag all;
}

root@ex-1> show log OSPF-debug
Mar  3 19:37:26.381199 OSPF packet ignored: can't find neighbor for
21.51.42.253 (router-id 21.100.0.4) on intf vlan.99 area 0.0.0.99
```

On EX1, R3_A and R4_A

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
{master:0}[edit]
root@ex-1# show | compare
[edit protocols ospf area 0.0.0.99 interface vlan.99]
+     apply-groups-except ospf-defaults;

[edit]
root@srx-4# show | compare
[edit routing-instances R4_Aprotocols OSPF area 0.0.0.99 interface ge-
0/0/12.0]
+     apply-groups-except ospf-defaults;
```

On ex1

```
{master:0}
root@ex-1> show OSPF neighbor
Address          Interface              State     ID               Pri  Dead
21.51.42.253     vlan.99                ExStart   21.100.0.100     128    33
21.51.42.252     vlan.99                ExStart   21.100.0.3       128    37

root@ex-1> show log OSPF-debug | find ospf_process_dbd
Mar  3 16:46:32.411375 ospf_process_dbd: processing dbd from 21.51.42.252
Mar  3 16:46:32.411397 OSPF restart siganling: Received DBD with LLS data from
nbr IP=21.51.42.252 id=21.100.0.3.
Mar  3 16:46:32.411421 OSPF packet ignored: MTU mismatch from 21.51.42.252 on
intf vlan.99 area 0.0.0.99
Mar  3 16:46:32.413365 task_job_delete: delete background job OSPF process
rcvd dbd packets for task OSPF


root@ex-1> show interfaces vlan | match MTU
  Type: VLAN, Link-level type: VLAN, MTU: 1518, Speed: 1000mbps

root@srx-4> show interfaces ge-0/0/12 | match MTU
  Link-level type: Ethernet, MTU: 1540, Link-mode: Full-duplex, Speed:
1000mbps,
    Protocol inet, MTU: 1526
    Protocol inet6, MTU: 1526
```

On EX1 and R3_A

```
{master:0}[edit]
root@ex-1# show | compare rollback 1
[edit interfaces vlan]
+   mtu 1544;

{master:0}
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@ex-1> show OSPF neighbor
Address          Interface              State     ID               Pri  Dead
21.51.42.253     vlan.99                Full      21.100.0.4     128    31
21.51.42.252     vlan.99                Full      21.100.0.3       128    39
```

7) You need to make sure that 10G links have a metric of 10 and 1G links have a metric of 100 by default. Use a single command to achieve this.

On R1, R2

```
root@srx-1-re-0# show | compare
[edit protocols OSPF]
+    reference-bandwidth 100g;
[edit protocols ospf3 realm ipv4-unicast]
+     reference-bandwidth 100g;
[edit protocols ospf3]
+    reference-bandwidth 100g;
```

On Ex1, R3_A and R4_A

```
root@ex-1# show | compare
[edit protocols OSPF]
+    reference-bandwidth 100g;
 [edit protocols ospf3]
+    reference-bandwidth 100g;

root@srx-4# show | compare
[edit routing-instances R4_A protocols ospf3 realm ipv4-unicast]
+     reference-bandwidth 100g;
[edit routing-instances R4_A protocols ospf3]
+    reference-bandwidth 100g;
```

On R2_B, R3_B and R4_B

```
root@srx-4# show | compare
[edit routing-instances R4_B protocols ospf3 realm ipv4-unicast]
+     reference-bandwidth 100g;
[edit routing-instances R4_B protocols ospf3]
+    reference-bandwidth 100g;
```

8) SRX4 is connected to a provider. The OSPF session between SRX4 and the provider VR_4 is down. Fix the issue.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On R4

```
[edit]
root@srx-4# run show OSPF neighbor
Address          Interface              State     ID               Pri  Dead
21.1.76.10       ge-0/0/14.68            Full      21.100.0.12      128    33

[edit]
root@srx-4# run show OSPF interface
Interface           State   Area            DR ID           BDR ID
Nbrs
ge-0/0/14.67         PtToPt  0.0.0.76        0.0.0.0         0.0.0.0
0
ge-0/0/14.68         PtToPt  0.0.0.76        0.0.0.0         0.0.0.0
1

[edit]
root@srx-4# run ping 21.1.76.5 rapid
PING 21.1.76.5 (21.1.76.5): 56 data bytes
!!!!!
--- 21.1.76.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.885/18.305/44.645/13.435 ms

[edit]
root@srx-4# show | compare
[edit protocols OSPF]
+    traceoptions {
+        file OSPF-debug;
+        flag all;
+    }

[edit]
root@srx-4# run show log OSPF-debug | match ignored
Feb 28 19:20:08.097080 OSPF packet ignored: authentication type mismatch (1)
from 21.1.76.5
Feb 28 19:20:16.191712 OSPF packet ignored: authentication type mismatch (1)
from 21.1.76.5

root@srx-4> monitor traffic interface ge-0/0/4.67 detail matching "proto OSPF"
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on ge-0/0/4.67, capture size 1514 bytes

Reverse lookup for 224.0.0.5 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use <no-resolve> to avoid reverse lookups on IP addresses.
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
19:25:55.840777 Out IP (tos 0xc0, ttl   1, id 2518, offset 0, flags [none],
proto: OSPF (89), length: 76) 21.1.76.6 > 224.0.0.5: OSPFv2, Hello, length 56
[len 44]
        Router-ID 21.100.0.11, Area 0.0.0.76, Authentication Type: none (0)
        Options [External, LLS]
          Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.252, Priority 128
          LLS: checksum: 0xfff6, length: 3
            Extended Options (1), length: 4
              Options: 0x00000001 [LSDB resync]
19:26:00.415489  In IP (tos 0xc0, ttl   1, id 3466, offset 0, flags [none],
proto: OSPF (89), length: 76) 21.1.76.5 > 224.0.0.5: OSPFv2, Hello, length 56
[len 44]
        Router-ID 21.1.76.2, Area 0.0.0.76, Authentication Type: simple (1)
        Simple text password: jncievpn
        Options [External, LLS]
          Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.252, Priority 128
          LLS: checksum: 0xfff6, length: 3
            Extended Options (1), length: 4
              Options: 0x00000001 [LSDB resync]

root@ex-2# show | compare
[edit protocols OSPF]
-    traceoptions {
-        file OSPF-debug;
-        flag all;
-    }
[edit protocols OSPF area 0.0.0.76 interface ge-0/0/4.67]
+      authentication {
+          simple-password "$9$3PPD/A0vMX-dsEcK8xN2gfTznCu"; ## SECRET-DATA
+      }


root@srx-4> show OSPF neighbor
Address          Interface              State     ID               Pri  Dead
21.1.76.5        ge-0/0/4.67            ExStart   21.1.76.2        128    34
21.1.76.10       ge-0/0/4.68            Full      21.100.0.12      128    37

root@srx-4> monitor traffic interface ge-0/0/4.67 detail matching "proto OSPF"
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on ge-0/0/4.67, capture size 1514 bytes

Reverse lookup for 224.0.0.5 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use <no-resolve> to avoid reverse lookups on IP addresses.
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
19:30:22.132805  In IP (tos 0xc0, ttl   1, id 4487, offset 0, flags [none],
proto: OSPF (89), length: 64) 21.1.76.5 > 224.0.0.5: OSPFv2, Database
Description, length 44 [len 32]
        Router-ID 21.1.76.2, Area 0.0.0.76, Authentication Type: simple (1)
        Simple text password: jncievpn
        Options [External, LLS, Opaque], DD Flags [Init, More, Master], MTU:
1500, Sequence: 0x150813fa
          LLS: checksum: 0xfff6, length: 3
            Extended Options (1), length: 4
              Options: 0x00000001 [LSDB resync]
19:30:22.133696 Out IP (tos 0xc0, ttl   1, id 3409, offset 0, flags [none],
proto: OSPF (89), length: 64) 21.1.76.6 > 224.0.0.5: OSPFv2, Database
Description, length 44 [len 32]
        Router-ID 21.100.0.11, Area 0.0.0.76, Authentication Type: simple (1)
        Simple text password: jncievpn
        Options [External, LLS, Opaque], DD Flags [Init, More, Master], MTU:
1522, Sequence: 0x15604cf6
          LLS: checksum: 0xfff6, length: 3
            Extended Options (1), length: 4
              Options: 0x00000001 [LSDB resync]

[edit]
root@srx-4# show | compare
[edit interfaces ge-0/0/4 unit 67 family inet]
+      mtu 1500;

root@srx-4> show OSPF neighbor
Address          Interface              State     ID               Pri  Dead
21.1.76.5        ge-0/0/4.67            Full      21.1.76.2        128    33
21.1.76.10       ge-0/0/4.68            Full      21.100.0.12      128    34
```

Verification

```
root@srx-4> show OSPF interface ge-0/0/4.67 extensive
Interface          State   Area            DR ID           BDR ID
Nbrs
ge-0/0/4.67        PtToPt  0.0.0.76        0.0.0.0         0.0.0.0
1
  Type: P2P, Address: 21.1.76.6, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: Password
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

9) SRX4 should connect to SRX2 using the provider router VR_4. SRX2 is using BGP to try and connect to the provider that is using AS27. SRX4 is running OSPF. You need to connect SRX4 to SRX2 though the use of a GRE tunnel. Make sure that the GRE tunnel is sourced from the loopback interface on both sides of the tunnel. Additional information is provided in the table below (Some configuration has already been done).

Device	IP address	Area
SRX 2	192.168.1.1/30	76
SRX 4	192.168.1.2/30	76

On EX2

```
root@srx-4> show route 21.100.0.2

inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

21.100.0.2/32      *[OSPF/150] 00:09:21, metric 0, tag 0
                    > to 21.1.76.5 via ge-0/0/4.67


root@srx-4> ping 21.100.0.2 source 21.100.0.11 rapid
PING 21.100.0.2 (21.100.0.2): 56 data bytes
!!!!!
--- 21.100.0.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.091/17.727/20.045/2.544 ms
[edit]
root@srx-4# show | compare
[edit interfaces gr-0/0/0]
+       unit 1 {
+           description "GRE to SRX2";
+           tunnel {
+               source 21.100.0.24;
+               destination 21.100.0.2;
+           }
+           family inet {
+               address 192.168.1.2/30;
+           }
+       }
 [edit protocols OSPF area 0.0.0.76]
      interface lo0.0 { ... }
+     interface gr-0/0/0.0 {
+         interface-type p2p;
+     }
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

On SRX2

```
[edit]
root@srx-2# show | compare
[edit interfaces]
+   gr-0/0/0 {
+       unit 0 {
+           description "GRE to SRX4";
+           tunnel {
+               source 21.100.0.2;
+               destination 21.100.0.24;
+           }
+           family inet {
+               address 192.168.1.1/30;
+           }
+       }
+   }
```

```
root@srx-2> ping 192.168.1.2 rapid
PING 192.168.1.2 (192.168.1.2): 56 data bytes
!!!!!
--- 192.168.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 13.334/14.556/16.140/1.277 ms
```

```
root@srx-2> show OSPF neighbor interface gr-0/0/0.0 detail
Address          Interface              State     ID               Pri  Dead
192.168.1.2      gr-0/0/0.0             Init      21.100.0.11      128    33
  Area 0.0.0.76, opt 0x12, DR 0.0.0.0, BDR 0.0.0.0
```

```
root@srx-4> show route 21.100.0.2

inet.0: 14 destinations, 16 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

21.100.0.2/32      *[OSPF/150] 00:00:10, metric 0, tag 0
                    > to 21.1.76.5 via ge-0/0/4.67

root@srx-4> show route 21.100.0.2
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
inet.0: 43 destinations, 44 routes (43 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

21.100.0.2/32      *[OSPF/10] 00:00:00, metric 1
                    > via gr-0/0/0.0

[edit]
root@srx-2# show | compare
[edit protocols BGP group as27]
+    import as27-import;
[edit protocols OSPF]
     area 0.0.0.0 { ... }
+    area 0.0.0.76 {
+        network-summary-export OSPF-gre-summary;
+    }
[edit policy-options]
+   policy-statement as27-import {
+       term 1 {
+           from {
+               route-filter 21.100.0.24/32 exact;
+           }
+           then {
+               preference 9;
+           }
+       }
+   }

+      policy-statement ospf-gre-summary {
+          term 1 {
+              from {
+                  route-filter 21.100.0.2/32 exact;
+              }
+              then reject;
+          }
+          term 2 {
+              then accept;
+          }
+      }
```

Verification

```
root@srx-2> show OSPF neighbor interface gr-0/0/0.0
Address          Interface              State     ID               Pri  Dead
192.168.1.2      gr-0/0/0.0             Full      21.100.0.11      128    32

root@srx-2> show route 21.100.0.24

inet.0: 156 destinations, 395 routes (155 active, 0 holddown, 62 hidden)
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+ = Active Route, - = Last Active, * = Both

21.100.0.24/32     *[BGP/9] 00:27:08, MED 1, localpref 100
                      AS path: 27 I
                    > to 21.1.76.2 via ge-0/0/4.66
                    [OSPF/10] 00:00:14, metric 1
                    > via gr-0/0/0.0

root@srx-2> show route 21.100.0.12

inet.0: 156 destinations, 395 routes (155 active, 0 holddown, 62 hidden)
+ = Active Route, - = Last Active, * = Both

21.100.0.12/32     *[OSPF/10] 00:00:47, metric 2
                    > via gr-0/0/0.0
                    [BGP/170] 00:27:41, MED 2, localpref 100
                      AS path: 27 I
                    > to 21.1.76.2 via ge-0/0/4.66
```

10) VR1 has been connected to EX1. The area is unknown. Establish neighbor relationship between EX1 and VR1.

Method 1 Traceoptions

```
[edit]
root@ex-1# show | compare
[edit protocols OSPF]
+    traceoptions {
+        file OSPF-debug;
+        flag all;
+    }
[edit protocols OSPF area 0.0.0.99]
      interface lo0.0 { ... }
+     interface ge-0/0/15.68;

root@ex-1> show log OSPF-debug
Feb 17 01:55:35.307924 OSPF packet ignored: area mismatch (1.0.5.5) from
21.55.145.2 on intf ge-0/0/5.68 area 0.0.0.99
Feb 17 01:55:35.307930 OSPF rcvd Hello 21.55.145.2 -> 224.0.0.5 (ge-0/0/5.68
IFL 71 area 0.0.0.99)
Feb 17 01:55:35.307936   Version 2, length 44, ID 21.55.140.1, area 1.0.5.5
Feb 17 01:55:35.307942   checksum 0x9dc5, authtype 0
Feb 17 01:55:35.307948   mask 255.255.255.252, hello_ivl 5, opts 0x12, prio
128
Feb 17 01:55:35.307954   dead_ivl 20, DR 21.55.145.2, BDR 0.0.0.0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit protocols OSPF]
     area 0.0.0.99 { ... }
+    area 1.0.5.5 {
+        interface ge-0/0/15.68 {
+            hello-interval 5;
+        }
+    }

root@ex-1> show OSPF neighbor interface ge-0/0/15.68
Address          Interface              State     ID               Pri  Dead
21.55.145.2      ge-0/0/15.68            ExStart   21.55.140.1      128    17

root@ex-1> show log OSPF-debug
Feb 17 02:13:26.465600 OSPF rcvd DbD 21.55.145.2 -> 224.0.0.5 (ge-0/0/5.68 IFL
71 area 1.0.5.5)
Feb 17 02:13:26.465606   Version 2, length 32, ID 21.55.140.1, area 1.0.5.5
Feb 17 02:13:26.465612   checksum 0x0, authtype 0
Feb 17 02:13:26.465618   options 0x52, i 1, m 1, ms 1, r 0, seq 0x15337655,
mtu 1500
Feb 17 02:13:26.974766 OSPF retransmitting DBD to 21.55.145.2
Feb 17 02:13:26.974792 OSPF sent DbD 21.55.145.1 -> 224.0.0.5 (ge-0/0/5.68 IFL
71 area 1.0.5.5)
Feb 17 02:13:26.974886   Version 2, length 32, ID 21.100.0.10, area 1.0.5.5
Feb 17 02:13:26.974892   options 0x52, i 1, m 1, ms 1, r 0, seq 0x1567642c,
mtu 1522

[edit interfaces ge-0/0/15 unit 68 family inet]
+      mtu 1500;
```

Method 2 Packet Capture

```
[edit protocols OSPF area 0.0.0.99]
      interface lo0.0 { ... }
+     interface ge-0/0/15.68;

monitor traffic interface ge-0/0/15.68 matching "proto OSPF"
01:57:47.087867  In
        Juniper PCAP Flags [Ext, no-L2, In], PCAP Extension(s) total length 22
          Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
          Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
          Device Interface Index Extension TLV #1, length 2, value: 138
          Logical Interface Index Extension TLV #4, length 4, value: 71
          Logical Unit Number Extension TLV #5, length 4, value: 68
        -----original packet-----
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
        PFE proto 2 (ipv4): (tos 0xc0, ttl   1, id 28197, offset 0, flags
[none], proto: OSPF (89), length: 76) 21.55.145.2 > 224.0.0.5: OSPFv2, Hello,
length 56 [len 44]
        Router-ID 21.55.140.1, Area 1.0.5.5, Authentication Type: none (0)
        Options [External, LLS]
          Hello Timer 5s, Dead Timer 20s, Mask 255.255.255.252, Priority 128
          Designated Router 21.55.145.2
          LLS: checksum: 0xfff6, length: 3
            Extended Options (1), length: 4
              Options: 0x00000001 [LSDB resync]

[edit protocols OSPF]
     area 0.0.0.99 { ... }
+    area 1.0.5.5 {
+        interface ge-0/0/15.68 {
+            hello-interval 5;
+        }
+    }

root@ex-1> show OSPF neighbor interface ge-0/0/15.68
Address          Interface              State     ID               Pri  Dead
21.55.145.2      ge-0/0/15.68            ExStart   21.55.140.1      128    17

02:09:58.905420  In
        Juniper PCAP Flags [Ext, no-L2, In], PCAP Extension(s) total length 22
          Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
          Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
          Device Interface Index Extension TLV #1, length 2, value: 138
          Logical Interface Index Extension TLV #4, length 4, value: 71
          Logical Unit Number Extension TLV #5, length 4, value: 68
        -----original packet-----
        PFE proto 2 (ipv4): (tos 0xc0, ttl   1, id 29864, offset 0, flags
[none], proto: OSPF (89), length: 64) 21.55.145.2 > 224.0.0.5: OSPFv2,
Database Description, length 44 [len 32]
        Router-ID 21.55.140.1, Area 1.0.5.5, Authentication Type: none (0)
        Options [External, LLS, Opaque], DD Flags [Init, More, Master], MTU:
1500, Sequence: 0x15337655
          LLS: checksum: 0xfff6, length: 3
            Extended Options (1), length: 4
              Options: 0x00000001 [LSDB resync]
02:10:00.539353 Out
        Juniper PCAP Flags [Ext], PCAP Extension(s) total length 22
          Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
          Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
          Device Interface Index Extension TLV #1, length 2, value: 138
          Logical Interface Index Extension TLV #4, length 4, value: 71
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
          Logical Unit Number Extension TLV #5, length 4, value: 68
        -----original packet-----
        00:0c:29:3b:c5:40 > 01:00:5e:00:00:05, ethertype 802.1Q (0x8100),
length 82: VLAN 68, p 6, ethertype IPv4, (tos 0xc0, ttl   1, id 4611, offset
0, flags [none], proto: OSPF (89), length: 64) 21.55.145.1 > 224.0.0.5:
OSPFv2, Database Description, length 44 [len 32]
        Router-ID 21.100.0.10, Area 1.0.5.5, Authentication Type: none (0)
        Options [External, LLS, Opaque], DD Flags [Init, More, Master], MTU:
1522, Sequence: 0x1567642c
          LLS: checksum: 0xfff6, length: 3
            Extended Options (1), length: 4
              Options: 0x00000001 [LSDB resync]
```

Verification

```
root@ex-1> show OSPF neighbor interface ge-0/0/5.68 detail
Address          Interface              State     ID               Pri  Dead
21.55.145.2      ge-0/0/5.68            Full      21.55.140.1      128    19
  Area 1.0.5.5, opt 0x52, DR 0.0.0.0, BDR 0.0.0.0
  Up 00:02:44, adjacent 00:02:44

root@ex-1> show route terse | match 21.55.145.2
* 21.55.140.0/24     O 150          0            >21.55.145.2
* 21.55.141.0/24     O 150          0            >21.55.145.2
* 21.55.142.0/24     O 150          0            >21.55.145.2
* 21.55.143.0/24     O 150          0            >21.55.145.2
```

11) You need to ensure that routes from VR1 are correctly propagated towards area 0. Use the following information to do this. You are allowed to use a single static route to complete this task.

Device	IP address	Area
VR1	21.10.10.2/30	0
R4_A	21.10.10.1/30	0

```
[edit]
root@srx-4# show | compare
[edit routing-instances R4_A routing-options]
+     static {
+         route 21.55.140.0/22 next-hop 21.51.42.254;
+     }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx-4> monitor traffic interface ge-0/0/.0 no-resolve matching "proto 47"
extensive
Address resolution is OFF.
Listening on ge-0/0/6.0, capture size 1514 bytes

03:20:03.854433  In
        Juniper PCAP Flags [Ext, no-L2, In], PCAP Extension(s) total length 16
          Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
          Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
          Device Interface Index Extension TLV #1, length 2, value: 140
          Logical Interface Index Extension TLV #4, length 4, value: 82
        -----original packet-----
        PFE proto 2 (ipv4): (tos 0x0, ttl  63, id 0, offset 0, flags [DF],
proto: GRE (47), length: 100) 21.55.142.99 > 21.100.0.4: gre 21.55.142.99 >
21.100.0.4: [] (tos 0xc0, ttl   1, id 39534, offset 0, flags [none], proto:
OSPF (89), length: 76) 21.10.10.2 > 224.0.0.5: OSPFv2, Hello, length 56 [len
44]
        Router-ID 21.55.140.1, Backbone Area, Authentication Type: none (0)
        Options [External, LLS]
          Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.252, Priority 128
          LLS: checksum: 0xfff6, length: 3
            Extended Options (1), length: 4
              Options: 0x00000001 [LSDB resync]

[edit]
root@srx-4# show | compare
[edit interfaces]
+   gr-0/0/0 {
+       unit 0 {
+           tunnel {
+               source 21.100.0.4;
+               destination 21.55.142.99;
+               routing-instance {
+                   destination R4_A;
+               }
+           }
+           family inet {
+               address 21.10.10.1/30;
+           }
+       }
+   }
[edit routing-instances R4_A]
+    interface gr-0/0/0.0;
[edit routing-instances R4_A routing-options static]
+      route 21.55.142.99/32 next-hop 21.51.42.254;
-      route 21.55.140.0/22 next-hop 21.51.42.254;
[edit routing-instances R4_A protocols OSPF area 0.0.0.0]
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+       interface gr-0/0/0.0;
        interface ge-0/0/5.30 { ... }
```

Verification

```
root@srx-4> show OSPF neighbor instance R4_A interface gr-0/0/0.0 detail
Address          Interface              State     ID               Pri  Dead
21.10.10.2       gr-0/0/0.0             Full      21.55.140.1      128    38
  Area 0.0.0.0, opt 0x52, DR 0.0.0.0, BDR 0.0.0.0
  Up 00:00:18, adjacent 00:00:18
    Link state retransmission list:  1 entries

root@srx-4> show route table R4_A terse | match gr-0/0/0
* 21.10.10.0/30      D   0                       >gr-0/0/0.0
                     O  10        125            >gr-0/0/0.0
* 21.55.140.0/24     O 150          0            >gr-0/0/0.0
* 21.55.141.0/24     O 150          0            >gr-0/0/0.0
* 21.55.142.0/24     O 150          0            >gr-0/0/0.0
* 21.55.143.0/24     O 150          0            >gr-0/0/0.0
* 21.55.145.0/30     O  10        126            >gr-0/0/0.0
```

12) Set up routing towards VR2 and VR3 using RIP and RIPng. Do not accept any routes with a prefix length longer than a /29 for IPv4 or a /125 for IPv6.

On R2_B

```
[edit]
root@srx-2# show | compare
[edit routing-instances R2_B protocols]
+   rip {
+       group remote {
+           import rip-import;
+           neighbor ge-0/0/4.62;
+           neighbor ge-0/0/5.64;
+       }
+   }
+   ripng {
+       group remote-v6 {
+           import ripng-import;
+           neighbor ge-0/0/4.62;
+           neighbor ge-0/0/5.64;
+       }
+   }
[edit]
+  policy-options {
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+      policy-statement rip-import {
+          term 1 {
+              from {
+                  protocol rip;
+                  route-filter 0.0.0.0/0 upto /29;
+              }
+              then accept;
+          }
+          term 2 {
+              then reject;
+          }
+      }
+      policy-statement ripng-import {
+          term 1 {
+              from {
+                  protocol ripng;
+                  route-filter ::/0 upto /125;
+              }
+              then accept;
+          }
+          term 2 {
+              then reject;
+          }
+      }
+  }
```

Verification

```
root@ex-2> show rip neighbor instance R2_B
                   Local  Source          Destination     Send   Receive   In
Neighbor           State  Address         Address         Mode   Mode     Met
--------           -----  -------         -----------     ----   -------  ---
ge-0/0/4.62           Up 21.1.99.1       224.0.0.9       mcast  both       1
ge-0/0/5.64           Up 21.1.99.5       224.0.0.9       mcast  both       1

root@ex-2> show ripng neighbor instance R2_B
                     Source                            Dest               In
Neighbor     State  Address                           Address  Send Recv Met
--------     -----  -------                           -------  ---- ---- ---
ge-0/0/4.62     Up fe80::20c:2900:3e86:338a         ff02::9  yes yes    1
ge-0/0/5.64     Up fe80::20c:2900:4086:3394         ff02::9  yes yes    1

root@ex-2> show route protocol rip table R2_B.inet.0

R2_B.inet.0: 31 destinations, 31 routes (31 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

21.1.96.0/24       *[RIP/100] 00:04:34, metric 2, tag 0
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                    > to 21.1.99.6 via ge-0/0/5.64
21.1.97.0/24       *[RIP/100] 00:04:34, metric 2, tag 0
                    > to 21.1.99.6 via ge-0/0/5.64
21.1.98.0/24       *[RIP/100] 00:04:34, metric 2, tag 0
                    > to 21.1.99.6 via ge-0/0/5.64
21.1.99.0/24       *[RIP/100] 00:04:34, metric 2, tag 0
                    > to 21.1.99.6 via ge-0/0/5.64
21.1.100.0/24      *[RIP/100] 00:08:29, metric 2, tag 0
                    > to 21.1.99.2 via ge-0/0/4.62
21.1.101.0/24      *[RIP/100] 00:08:29, metric 2, tag 0
                    > to 21.1.99.2 via ge-0/0/4.62
21.1.102.0/24      *[RIP/100] 00:08:29, metric 2, tag 0
                    > to 21.1.99.2 via ge-0/0/4.62
21.1.103.0/24      *[RIP/100] 00:08:29, metric 2, tag 0
                    > to 21.1.99.2 via ge-0/0/4.62
224.0.0.9/32       *[RIP/100] 00:00:14, metric 1
                      MultiRecv

root@ex-2> show route protocol ripng table R2_B.inet6.0

R2_B.inet6.0: 32 destinations, 38 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2002:95:b::f0:0/112*[RIPng/100] 00:04:41, metric 2, tag 0
                    > to fe80::20c:2900:4073:a775 via ge-0/0/5.64
2002:95:b::f2:0/112*[RIPng/100] 00:08:30, metric 2, tag 0
                    > to fe80::20c:2900:3e73:a775 via ge-0/0/4.62
2002:95:b::f3:0/112*[RIPng/100] 00:08:30, metric 2, tag 0
                    > to fe80::20c:2900:3e73:a775 via ge-0/0/4.62
2002:95:b::f4:0/112*[RIPng/100] 00:08:30, metric 2, tag 0
                    > to fe80::20c:2900:3e73:a775 via ge-0/0/4.62
2002:95:b::f5:0/112*[RIPng/100] 00:08:30, metric 2, tag 0
                    > to fe80::20c:2900:3e73:a775 via ge-0/0/4.62
2002:95:b::f6:0/112*[RIPng/100] 00:04:41, metric 2, tag 0
                    > to fe80::20c:2900:4073:a775 via ge-0/0/5.64
2002:95:b::f7:0/112*[RIPng/100] 00:04:41, metric 2, tag 0
                    > to fe80::20c:2900:4073:a775 via ge-0/0/5.64
ff02::9/128        *[RIPng/100] 00:00:19, metric 1
                      MultiRecv
```

13) SRX1 and SRX2 need to be able to reach the RIP routes. You should be able to ping 21.1.100.1 and 21.1.96.1 from the loopback of R1 and R2.

On R2_B

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit routing-instances R2_B protocols rip group remote]
+    export rip-export;
[edit routing-instances R2_B protocols ripng group remote-v6]
+    export ripng-export;
[edit policy-options]
+   policy-statement rip-export {
+       term 1 {
+           from {
+               protocol ospf3;
+               route-filter 0.0.0.0/0 exact;
+           }
+           then accept;
+       }
+       term 2 {
+           then reject;
+       }
+   }
+   policy-statement ripng-export {
+       term 1 {
+           from {
+               protocol ospf3;
+               route-filter ::/0 exact;
+           }
+           then accept;
+       }
+       term 2 {
+           then reject;
+       }
+   }

[edit routing-instances R2_B protocols ospf3 realm ipv4-unicast]
+    export ospf3-v4-export;
[edit routing-instances R2_B protocols ospf3]
+   export ospf3-v6-export;
[edit policy-options]
+   policy-statement ospf3-v4-export {
+       term 1 {
+           from {
+               protocol rip;
+               route-filter 21.1.96.0/21 orlonger;
+           }
+           then {
+               external {
+                   type 1;
+               }
+               accept;
+           }
+       }
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+       term 2 {
+           then reject;
+       }
+   }
+   policy-statement ospf3-v6-export {
+       term 1 {
+           from {
+               protocol ripng;
+               route-filter 2002:95:b::f0:0/109 orlonger;
+           }
+           then {
+               external {
+                   type 1;
+               }
+               accept;
+           }
+       }
+       term 2 {
+           then reject;
+       }
+   }
```

Verification

On SRX1

```
root@srx-1-re-0> ping 21.1.96.1 source 21.100.0.1 rapid
PING 21.1.96.1 (21.1.96.1): 56 data bytes
!!!!!
--- 21.1.96.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 18.755/19.856/21.219/0.789 ms

root@srx-1-re-0> ping 21.1.100.1 source 21.100.0.1 rapid
PING 21.1.100.1 (21.1.100.1): 56 data bytes
!!!!!
--- 21.1.100.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 15.878/18.644/20.019/1.635 ms
```

14) Make sure that 21.1.96.0/21 and 2002:95:b::f0:0/109 are the only RIP routes advertised towards other areas.

On R3_B

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit routing-instances R3_B protocols ospf3 realm ipv4-unicast area 0.0.0.23
nssa]
+         area-range 21.1.96.0/21;
```

On R4_B

```
[edit routing-instances R4_B protocols ospf3 realm ipv4-unicast area 0.0.0.23
nssa]
+         area-range 21.1.96.0/21;
```

Verification

```
root@srx-1-re-0> show route 21.1.96.0

inet.0: 141 destinations, 214 routes (141 active, 0 holddown, 7 hidden)
+ = Active Route, - = Last Active, * = Both

21.1.96.0/21       *[OSPF3/150] 00:32:10, metric 152, tag 0
                    > to 21.1.10.18 via ge-0/0/4.40
                    [OSPF/150] 00:32:05, metric 153, tag 20
                    > to 21.1.10.6 via ge-0/0/1.70

root@srx-1-re-0> show route 21.1.99.0

inet.0: 141 destinations, 214 routes (141 active, 0 holddown, 7 hidden)
+ = Active Route, - = Last Active, * = Both

21.1.96.0/21       *[OSPF3/150] 00:32:17, metric 152, tag 0
                    > to 21.1.10.18 via ge-0/0/4.40
                    [OSPF/150] 00:32:12, metric 153, tag 20
                    > to 21.1.10.6 via ge-0/0/1.70

root@srx-1-re-0> ping 21.1.96.1 rapid
PING 21.1.96.1 (21.1.96.1): 56 data bytes
!!!!!
--- 21.1.96.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.189/17.365/21.032/2.350 ms
```

15) Ensure that the OSPFv2 networks are able to reach IPv4 networks known to OSPFv3.

On R1

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit]
root@srx-1-re-0# show | compare
[edit protocols OSPF]
+       export ospf3-to-ospf;
 [edit protocols ospf3 realm ipv4-unicast]
+    export ospf-to-ospf3;
[edit]
+  policy-options {
+      policy-statement ospf-to-ospf3 {
+          term 1 {
+              from tag 20;
+              then reject;
+          }
+          term 2 {
+              from protocol OSPF;
+              then {
+                  tag 10;
+                  accept;
+              }
+          }
+      }
+      policy-statement ospf3-to-ospf {
+          term 1 {
+              from tag 10;
+              then reject;
+          }
+          term 2 {
+              from {
+                  family inet;
+                  protocol ospf3;
+              }
+              then {
+                  tag 20;
+                  accept;
+              }
+          }
+      }
+  }

[edit]
root@srx-1-re-0# show protocols ospf3 realm ipv4-unicast area 0 interface ge-
0/0/4.40
metric 50;
```

On R2

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit]
root@srx-2# show | compare
[edit protocols OSPF]
+       export ospf3-to-ospf;
 [edit protocols ospf3 realm ipv4-unicast]
+    export ospf-to-ospf3;
[edit]
+  policy-options {
+      policy-statement ospf-to-ospf3 {
+          term 1 {
+              from tag 20;
+              then reject;
+          }
+          term 2 {
+              from protocol OSPF;
+              then {
+                  tag 10;
+                  accept;
+              }
+          }
+      }
+      policy-statement ospf3-to-ospf {
+          term 1 {
+              from tag 10;
+              then reject;
+          }
+          term 2 {
+              from {
+                  family inet;
+                  protocol ospf3;
+              }
+              then {
+                  tag 20;
+                  accept;
+              }
+          }
+      }
+  }

[edit]
root@srx-1# show protocols ospf3 realm ipv4-unicast area 0 interface ge-
0/0/4.42
metric 50;
```

On R4_B

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit]
root@srx-4# show | compare
[edit routing-instances R4_B protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.42]
+        metric 50;
[edit routing-instances R4_B protocols ospf3 realm ipv4-unicast area 0.0.0.0
interface ge-0/0/4.40]
+        metric 50;
```

Verification

```
root@ex-1> show route 21.100.0.13

inet.0: 37 destinations, 38 routes (37 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

21.100.0.13/32     *[OSPF/150] 00:41:47, metric 150, tag 20
                      to 21.51.42.252 via ge-0/0/6.0
                    > to 21.51.42.253 via ge-0/0/6.0

root@ex-1> ping 21.100.0.13 rapid
PING 21.100.0.13 (21.100.0.13): 56 data bytes
!!!!!
--- 21.100.0.11 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 13.709/20.194/27.210/5.246 ms

root@ex-1> show route 21.1.96.0

inet.0: 37 destinations, 38 routes (37 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

21.1.96.0/21       *[OSPF/150] 00:42:13, metric 352, tag 20
                    > to 21.51.42.252 via ge-0/0/6.0
                      to 21.51.42.253 via ge-0/0/6.0

root@ex-1> ping 21.1.96.1 rapid
PING 21.1.96.1 (21.1.96.1): 56 data bytes
!!!!!
--- 21.1.96.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 25.113/28.928/33.039/2.842 ms

root@ex-1> traceroute 21.1.96.1 no-resolve
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

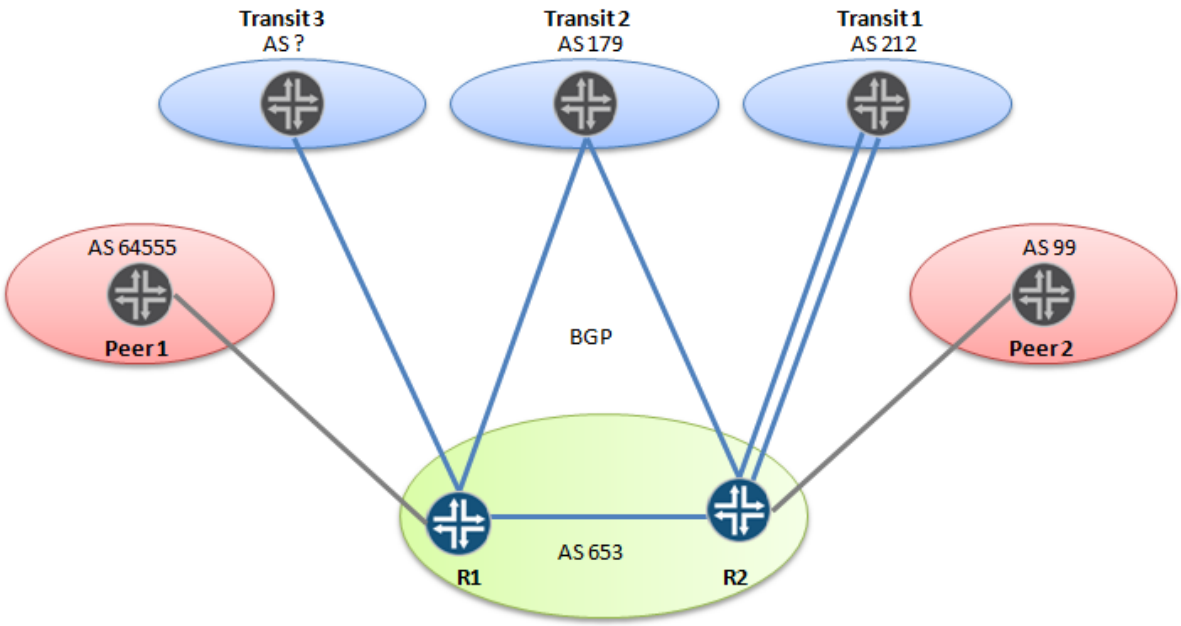
```
traceroute to 21.1.96.1 (21.1.96.1), 30 hops max, 40 byte packets
 1  21.51.42.252  12.529 ms  18.293 ms  9.608 ms
 2  21.1.10.29  15.580 ms  13.669 ms  18.889 ms
 3  21.1.10.34  19.983 ms  20.165 ms  16.038 ms
 4  21.1.10.50  28.560 ms  34.662 ms  23.820 ms
 5  21.1.96.1  30.882 ms  39.194 ms  29.838 ms

root@ex-1>
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 5: BGP Routing



Device	Interface A-End	VLAN	IP Address	Device B-End	Auth key
SRX1	Ge-0/0/4.59	59	8.254.157.2 /30	Transit3	n/a
SRX1	Ge-0/0/5.99	99	7.16.1.2/30	Transit2	jnciebgp
			2002:a179:ef:5413::2/126		n/a
SRX2	Ge-0/0/5.70	70	21.24.1.1/30	Peer2	n/a
SRX2	Ge-0/0/4.12	12	7.16.15.2/30	Transit2	jnciebgp
			00		n/a
SRX2	Ge-0/0/4.101	101	6.6.6.2/30	Transit1-A	n/a
SRX2	Ge-0/0/5.102	102	6.7.7.2/30	Transit1-B	n/a
SRX1	Ge-0/0/4.3	3	21.21.21.1/30	Peer1	n/a

Table 5.1

NOTE: *Some of the information is not provided and you will need to troubleshoot to find the required information.*

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

1) A junior engineer has set up IBGP between R1 and R2. The session is not working. Investigate the issues and fix them.

root@srx-1-re-0> show **route 21.100.0.2**

inet.0: 141 destinations, 214 routes (141 active, 0 holddown, 7 hidden)
+ = Active Route, - = Last Active, * = Both

21.100.0.2/32 *[OSPF/10] 00:48:05, metric 1
> to 21.1.10.6 via ge-0/0/1.70
[OSPF3/10] 00:47:21, metric 100
> to 21.1.10.2 via ge-0/0/1.69
to 21.1.10.18 via ge-0/0/4.40

root@srx-1-re-0> **ping 21.100.0.2 rapid**
PING 21.100.0.2 (21.100.0.2): 56 data bytes
!!!!!
--- 21.100.0.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.979/4.327/6.225/1.636 ms

root@srx-1-re-0> show configuration **protocols BGP group ibgp**
type internal;
neighbor 21.100.0.2

On R1

```
 [edit protocols BGP group ibgp]
+            local-address 24.100.0.1;
```

On R2

```
[edit protocols BGP group ibgp]
+            local-address 24.100.0.2;
```

Verification

```
root@srx-1-re-0> show BGP summary | match 653
21.100.0.2              653        115        142       0       0       46:12
37/38/38/0           0/0/0/0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx-1-re-0> show BGP neighbor 21.100.0.2
Peer: 21.100.0.2+179 AS 653    Local: 21.100.0.1+55109 AS 653
  Type: Internal    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ BGP-export ]
  Options: <Preference LocalAddress LogUpDown Refresh>
  Options: <MtuDiscovery>
  Local Address: 21.100.0.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 21.100.0.2      Local ID: 21.100.0.100      Active Holdtime: 90
  Keepalive Interval: 30         Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 653)
  Peer does not support Addpath
  Table inet.0 Bit: 10001
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              37
    Received prefixes:            38
    Accepted prefixes:            38
    Suppressed due to damping:    0
    Advertised prefixes:          92
  Last traffic (seconds): Received 0    Sent 17   Checked 77
  Input messages:  Total 116    Updates 13      Refreshes 0     Octets 3040
  Output messages: Total 142    Updates 37      Refreshes 0     Octets 4708
  Output Queue[0]: 0

root@srx-1-re-0> show system connections | match 179
tcp4       0      0  21.100.0.1.55109
21.100.0.2.179                         ESTABLISHED
```

2) You have two peering sessions with Transit 1. Update the configuration in such a way that you only have a single BGP session with Transit 1. This session should use both of the available links. You are allowed to use a single static route to complete this task.

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx-2> show BGP summary | match 212
6.6.6.1                 212        281        299       0       0     2:05:15
38/92/92/0           0/0/0/0
6.7.7.1                 212        279        298       0       0     2:05:11
0/92/92/0            0/0/0/0

[edit routing-options]
+   static {
+       route 6.6.1.1/32 next-hop [ 6.6.6.1 6.7.7.1 ];
+   }
+   forwarding-table {
+       export lb-per-packet;
+   }
[edit protocols BGP]
+    group AS_179 {
+        type external;
+        neighbor 6.6.1.1 {
+            multihop {
+                ttl 10;
+            }
+            local-address 21.100.0.2;
+            peer-as 212;
+        }
+    }
     group ibgp { ... }
[edit protocols BGP]
-    group Transit1 {
-        type external;
-        neighbor 6.6.6.1 {
-            peer-as 212;
-        }
-        neighbor 6.7.7.1 {
-            peer-as 212;
-        }
-    }
[edit policy-options]
+   policy-statement lb-per-packet {
+       then {
+           load-balance per-packet;
+       }
+   }
```

Verification

```
root@srx-2> show BGP summary
Groups: 2 Peers: 2 Down peers: 0
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.0               184         96          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
6.6.1.1                 212         12         16       0       0          22
38/92/92/0           0/0/0/0

root@srx-2> show route forwarding-table family inet matching 6.6.1.1
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
6.6.1.1/32         user     1                    ulst 262150     4
                              6.6.6.1           ucst   651     4 ge-0/0/4.101
                              6.7.7.1           ucst   650     3 ge-0/0/5.102
```

3) The connection towards peer 1 is not working. Peer 1 has been configured to peer with 65522. Set up the peering connection by altering the configuration on R1.

```
root@srx-1-re-0> show BGP summary | match 64
21.21.21.2            64555          1          8       0       0     3:02:49
Active

[edit]
root@srx-1-re-0# show | compare
[edit protocols BGP group AS_64555 neighbor 21.21.21.2]
+      local-as 65522;
```

Verification

```
root@srx-1-re-0> show BGP summary | match 64
21.21.21.2            64555          4         27       0       0           9
2/3/3/0              0/0/0/0
```

4) Transit 3 is a new transit. A junior engineer has given you the IP address and interface information (Table 5.1) but not the ASN. An SLA requirement for Peer 1 is a working session with Transit 3. Get the peering with Transit 3 working ASAP.

```
[edit]
root@srx-1-re-0# run show BGP summary | match 8.254.157.1
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
8.254.157.1             100          0          0       0       0     3:06:41
Connect

[edit]
root@srx-1-re-0# run ping  8.254.157.1 rapid
PING 8.254.157.1 (8.254.157.1): 56 data bytes
.....
--- 8.254.157.1 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss

[edit]
root@srx-1-re-0# run show interfaces ge-0/0/4.59
  Logical interface ge-0/0/4.59 (Index 76) (SNMP ifIndex 533)
    Description: Transit3
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.58 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 2093
    Security: Zone: Null
    Protocol inet, MTU: 1522
      Flags: Sendbcast-pkt-to-re, uRPF, uRPF-loose
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 8.254.157.0/30, Local: 8.254.157.2, Broadcast:
8.254.157.3
    Protocol inet6, MTU: 1522
      Flags: None
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::20c:2900:3aa7:4eae

[edit]
root@srx-1-re-0# show | compare
[edit interfaces ge-0/0/4]
+    unit 59 {
+        VLAN-id 59;
+    }
```

```
root@srx-1-re-0> ping 8.254.157.1 rapid
PING 8.254.157.1 (8.254.157.1): 56 data bytes
!!!!!
--- 8.254.157.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.713/4.927/6.932/1.260 ms
```

Method 1 Traceoptions

```
[edit]
root@srx-1-re-0# show | compare
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit protocols BGP group AS_Transit3]
+     traceoptions {
+         file BGP-debug;
+         flag all;
+     }


root@srx-1-re-0> how log BGP-debug | match "NOTIFICATION sent to 8.254.157.1"
Mar  2 03:10:07.891416 bgp_process_open:2789: NOTIFICATION sent to 8.254.157.1
(External AS 100): code 2 (Open Message Error) subcode 2 (bad peer AS number),
Reason: peer 8.254.157.1 (External AS 100) claims 735, 100 configured
```

Method 2 Packet Capture

```
 root@srx-1-re-0> monitor traffic interface ge-0/0/4.59 matching "port BGP"
extensive
03:14:23.999420 Out
        Juniper PCAP Flags [Ext], PCAP Extension(s) total length 22
          Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
          Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
          Device Interface Index Extension TLV #1, length 2, value: 137
          Logical Interface Index Extension TLV #4, length 4, value: 85
          Logical Unit Number Extension TLV #5, length 4, value: 59
        -----original packet-----
        00:0c:29:a7:4e:ae > 00:0c:29:73:a7:75, ethertype 802.1Q (0x8100),
length 70: VLAN 59, p 6, ethertype IPv4, (tos 0xc0, ttl   1, id 17903, offset
0, flags [DF], proto: TCP (6), length: 52) 8.254.157.2.56558 >
8.254.157.1.bgp: . 1:1(0) ack 1 win 17376 <nop,nop,timestamp 5953536 2384707>
03:14:23.999698 Out
        Juniper PCAP Flags [Ext], PCAP Extension(s) total length 22
          Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
          Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
          Device Interface Index Extension TLV #1, length 2, value: 137
          Logical Interface Index Extension TLV #4, length 4, value: 85
          Logical Unit Number Extension TLV #5, length 4, value: 59
        -----original packet-----
        00:0c:29:a7:4e:ae > 00:0c:29:73:a7:75, ethertype 802.1Q (0x8100),
length 129: VLAN 59, p 6, ethertype IPv4, (tos 0xc0, ttl   1, id 17904, offset
0, flags [DF], proto: TCP (6), length: 111) 8.254.157.2.56558 >
8.254.157.1.bgp: P 1:60(59) ack 1 win 17376 <nop,nop,timestamp 5953536
2384707>: BGP, length: 59
        Open Message (1), length: 59
          Version 4, my AS 653, Holdtime 90s, ID 21.100.0.100
          Optional parameters, length: 30
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
              Option Capabilities Advertisement (2), length: 6
                Multiprotocol Extensions (1), length: 4
                  AFI IPv4 (1), SAFI Unicast (1)
                  0x0000: 0001 0001
              Option Capabilities Advertisement (2), length: 2
                Route Refresh (Cisco) (128), length: 0
              Option Capabilities Advertisement (2), length: 2
                Route Refresh (2), length: 0
              Option Capabilities Advertisement (2), length: 4
                Graceful Restart (64), length: 2
                  Restart Flags: [none], Restart Time 120s
                  0x0000: 0078
              Option Capabilities Advertisement (2), length: 6
                32-Bit AS Number (65), length: 4
                   4 Byte AS 653
                  0x0000: 0000 028d
03:14:24.009291  In
        Juniper PCAP Flags [Ext, no-L2, In], PCAP Extension(s) total length 22
          Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
          Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
          Device Interface Index Extension TLV #1, length 2, value: 137
          Logical Interface Index Extension TLV #4, length 4, value: 85
          Logical Unit Number Extension TLV #5, length 4, value: 59
        -----original packet-----
        PFE proto 2 (ipv4): (tos 0xc0, ttl   1, id 30448, offset 0, flags
[none], proto: TCP (6), length: 127) 8.254.157.1.bgp > 8.254.157.2.56558: P
1:76(75) ack 60 win 16384 <nop,nop,timestamp 2384709 5953536>: BGP, length: 75
        Open Message (1), length: 75
          Version 4, my AS 735, Holdtime 90s, ID 8.8.0.1
          Optional parameters, length: 46
            Option Capabilities Advertisement (2), length: 6
              Multiprotocol Extensions (1), length: 4
                AFI IPv4 (1), SAFI Unicast (1)
                0x0000: 0001 0001
            Option Capabilities Advertisement (2), length: 6
              Multiprotocol Extensions (1), length: 4
                AFI IPv6 (2), SAFI Unicast (1)
                0x0000: 0002 0001
            Option Capabilities Advertisement (2), length: 2
              Route Refresh (Cisco) (128), length: 0
            Option Capabilities Advertisement (2), length: 2
              Route Refresh (2), length: 0
            Option Capabilities Advertisement (2), length: 12
              Graceful Restart (64), length: 10
                Restart Flags: [none], Restart Time 120s
                  AFI IPv4 (1), SAFI Unicast (1), Forwarding state preserved:
yes
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                  AFI IPv6 (2), SAFI Unicast (1), Forwarding state preserved:
yes
                0x0000: 0078 0001 0180 0002 0180
             Option Capabilities Advertisement (2), length: 6
               32-Bit AS Number (65), length: 4
                 4 Byte AS 735
                0x0000: 0000 02df
root@srx-1-re-0# show | compare
[edit protocols BGP group AS_Transit3 neighbor 8.254.157.1]
-     peer-as 100;
+     peer-as 735;
```

Verification

```
root@srx-1-re-0> show BGP summary | match 735
8.254.157.1             735         18         31       0       0          43
26/64/64/0           0/0/0/0

root@srx-1-re-0> show BGP neighbor 8.254.157.1
Peer: 8.254.157.1+62767 AS 735 Local: 8.254.157.2+179 AS 653
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ BGP-export ]
  Options: <Preference LogUpDown PeerAS Refresh>
  Options: <MtuDiscovery>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 8.8.0.1         Local ID: 21.100.0.100      Active Holdtime: 90
  Keepalive Interval: 30         Peer index: 0
  BFD: disabled, down
  Local Interface: ge-0/0/4.59
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast inet6-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI peer can save forwarding state: inet-unicast inet6-unicast
  NLRI that peer saved forwarding for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 735)
  Peer does not support Addpath
  Table inet.0 Bit: 10003
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              26
    Received prefixes:            64
    Accepted prefixes:            64
    Suppressed due to damping:    0
    Advertised prefixes:          105
  Last traffic (seconds): Received 27   Sent 3    Checked 55
  Input messages:  Total 18     Updates 15      Refreshes 0     Octets 1303
  Output messages: Total 32     Updates 27      Refreshes 0     Octets 2090
  Output Queue[0]: 0
  Trace options:  all
  Trace file: /var/log/BGP-debug size 131072 files 10
```

5) You have been told that Transit 3 has BFD sessions configured for the BGP peering session. Configure BFD for your side with an interval of 1000ms and a multiplier of 3.

```
[edit]
root@srx-1-re-0# show | compare
[edit protocols BGP group AS_Transit3 neighbor 8.254.157.1]
+      bfd-liveness-detection {
+          minimum-interval 1000;
+          multiplier 3;
+      }
```

Verification

```
root@srx-1-re-0> show bfd session
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval
Multiplier
8.254.157.1              Up        ge-0/0/4.59    3.000     1.000        3

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

6) Exchange IPv6 routes using the IPv4 session you have with Transit 3. You can use a mapped IPv4 address.

```
[edit]
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx-1-re-0# show | compare
[edit interfaces ge-0/0/4 unit 59]
+      family inet6 {
+          address ::ffff:8.254.157.2/126;
+      }
[edit protocols BGP group AS_Transit3]
+     family inet {
+         unicast;
+     }
+     family inet6 {
+         unicast;
+     }
```

Verification

```
root@srx-1-re-0> show BGP summary
Groups: 4 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.0               166         95          0          0          0
0
inet6.0                8          7          0          0          0
0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
8.254.157.1             735         21         38       0       0        1:37
Establ
  inet.0: 26/64/64/0
  inet6.0: 7/8/8/0


root@srx-1-re-0> show route protocol BGP table inet6.0

inet6.0: 54 destinations, 64 routes (54 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::ffff:8.254.157.0/126
                    [BGP/170] 00:02:17, localpref 100, from 8.254.157.1
                      AS path: 735 I
                    > to ::ffff:8.254.157.1 via ge-0/0/4.59
2002:ac34:5ea5:1::/64
                   *[BGP/170] 00:02:17, localpref 100, from 8.254.157.1
                      AS path: 735 I
                    > to ::ffff:8.254.157.1 via ge-0/0/4.59
2002:ac34:5ea5:2::/64
                   *[BGP/170] 00:02:17, localpref 100, from 8.254.157.1
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                      AS path: 735 I
                    > to ::ffff:8.254.157.1 via ge-0/0/4.59
2002:ac34:5ea5:3::/64
                    *[BGP/170] 00:02:17, localpref 100, from 8.254.157.1
                      AS path: 735 I
                    > to ::ffff:8.254.157.1 via ge-0/0/4.59
2002:ac34:5ea5:4::/64
                    *[BGP/170] 00:02:17, localpref 100, from 8.254.157.1
                      AS path: 735 I
                    > to ::ffff:8.254.157.1 via ge-0/0/4.59
2002:ac34:5ea5:e::/64
                    *[BGP/170] 00:02:17, localpref 100, from 8.254.157.1
                      AS path: 735 I
                    > to ::ffff:8.254.157.1 via ge-0/0/4.59
2002:ac34:5ea5:f::/64
                    *[BGP/170] 00:02:17, localpref 100, from 8.254.157.1
                      AS path: 735 I
                    > to ::ffff:8.254.157.1 via ge-0/0/4.59
2002:ef54:ac11:56b1::/126
                    *[BGP/170] 00:02:17, localpref 100, from 8.254.157.1
                      AS path: 735 I
                    > to ::ffff:8.254.157.1 via ge-0/0/4.59
```

7) Advertise your IP range, 21.0.0.0/8, to all your EBGP neighbors. Tag the exported route with community 653:101. You can use an aggregate route to complete this task. You are not allowed to use any prefix-lists or route-filter commands.

On R1

```
[edit]
root@srx-1-re-0# show | compare
[edit routing-options]
+   aggregate {
+       route 21.0.0.0/8 community 653:101;
+   }
[edit protocols BGP group AS_179]
-    export BGP-export-defaults;
+    export as653-export;
[edit protocols BGP group AS_Transit3]
+    export as653-export;
 [edit policy-options]
+   policy-statement as653-export {
+       term 1 {
+           from community agg;
+           then accept;
+       }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+       term default-reject {
+           then reject;
+       }
+   }
[edit policy-options]
+   community agg members 653:101;
```

On R2

```
[edit]
root@srx-2# show | compare
[edit routing-options]
+   aggregate {
+       route 21.0.0.0/8 community 653:101;
+   }
[edit protocols BGP group AS_212]
+    export as653-export;
[edit policy-options]
+   policy-statement as653-export {
+       term 1 {
+           from community local;
+           then accept;
+       }
+       term default-reject {
+           then reject;
+       }
+   }
[edit policy-options]
+   community local members 653:101;
```

Verification

```
root@srx-1-re-0> show route advertising-protocol BGP 8.254.157.1 detail

inet.0: 144 destinations, 358 routes (144 active, 0 holddown, 2 hidden)
* 21.0.0.0/8 (2 entries, 1 announced)
 BGP group AS_Transit3 type External
     Nexthop: Self
     AS path: [653] I (LocalAgg)
     Communities: 653:101

root@srx-1-re-0> ping 8.8.8.8 source 21.100.0.1 rapid
PING 8.8.8.8 (8.8.8.8): 56 data bytes
!!!!!
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
--- 8.8.8.8 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.455/4.562/5.732/0.949 ms

root@srx-1-re-0> ping 1.1.1.1 source 21.100.0.1 rapid
PING 1.1.1.1 (1.1.1.1): 56 data bytes
!!!!!
--- 1.1.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.878/5.357/7.311/1.266 ms
```

8) Advertise Peer 1's routes to Transit 3 only. Transit 3 will accept routes with your community 653:* on it. Use a single new policy to achieve this. Once it is working, you should be able to ping "1.1.1.1" from Peer's 1 loopback address.

```
[edit]
root@srx-1-re-0# show | compare
[edit protocols BGP group AS_Transit3]
-    export as653-export;
+    export [ as_735-export as653-export ];
+     remove-private;
[edit protocols BGP group AS_64555]
+    import as_64555-import;
[edit policy-options]
+   policy-statement as_64555-import {
+       term 1 {
+           from {
+               route-filter 11.0.0.0/8 exact;
+           }
+           then {
+               community set as64555;
+           }
+       }
+       term 2 {
+           then accept;
+       }
+   }
+   policy-statement as_735-export {
+       term 1 {
+           from community as64555;
+           then accept;
+       }
+       term last {
+           then next policy;
+       }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+    }
[edit policy-options]
+    community as64555 members 653:102;
```

Verification

```
root@srx-1-re-0> show route advertising-protocol BGP 8.254.157.1 detail

inet.0: 144 destinations, 266 routes (144 active, 0 holddown, 1 hidden)
* 11.0.0.0/8 (1 entry, 1 announced)
 BGP group AS_Transit3 type External
     Nexthop: Self
     AS path: 65522 64555 I
     Communities: 653:102

* 21.0.0.0/8 (2 entries, 1 announced)
 BGP group AS_Transit3 type External
     Nexthop: Self
     AS path: [653] I (LocalAgg)
     Communities: 653:101

root@srx-1-re-0> show route advertising-protocol BGP 7.16.1.1

inet.0: 144 destinations, 266 routes (144 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 21.0.0.0/8              Self                                    I


root@VR_Router_Lab3> traceroute routing-instance peer_1 1.1.1.1 source
11.0.0.1
traceroute to 1.1.1.1 (1.1.1.1) from 11.0.0.1, 30 hops max, 40 byte packets
 1  21.21.21.1 (21.21.21.1)  8.506 ms  9.300 ms  9.746 ms
 2  7.16.1.1 (7.16.1.1)  19.860 ms  17.476 ms  19.793 ms
 3  1.1.1.1 (1.1.1.1)  18.807 ms  13.288 ms  19.704 ms
```

9) Transit 1 should be used to reach all routes, except for the routes that originate from the connected transit providers. Ensure SRX 1 and SRX 2 use the correct exit point.

On SRX 1

```
[edit]
root@srx-1-re-0# show | compare
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit protocols BGP group AS_179]
+    import as179-import;
[edit protocols BGP group AS_Transit3]
+    import as735-import;
[edit protocols BGP group ibgp]
+    export ibg-export;
[edit policy-options]
+   policy-statement as179-import {
+       term 1 {
+           from as-path as179;
+           then {
+               local-preference 1000;
+               accept;
+           }
+       }
+       term 2 {
+           then accept;
+       }
+   }
+   policy-statement as735-import {
+       term 1 {
+           from as-path as735;
+           then {
+               local-preference 1000;
+               accept;
+           }
+       }
+       term 2 {
+           then accept;
+       }
+   }
+   policy-statement ibg-export {
+       term 1 {
+           then {
+               next-hop self;
+           }
+       }
+   }
[edit policy-options]
+   as-path as179 "^179$";
+   as-path as735 "^735$";
```

On SRX 2

```
root@srx-2# show | compare
[edit protocols BGP group as179]
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+    import as179-import;
[edit protocols BGP group AS_212]
+    import as212-import;
[edit protocols BGP group ibgp]
+    export ibg-export;
[edit policy-options]
+   policy-statement as179-import {
+       term 1 {
+           from as-path as179;
+           then {
+               local-preference 1000;
+               accept;
+           }
+       }
+       term 2 {
+           then accept;
+       }
+   }
+   policy-statement as212-import {
+       term 1 {
+           from as-path as212;
+           then {
+               local-preference 1000;
+               accept;
+           }
+       }
+       term 2 {
+           then {
+               local-preference 500;
+               accept;
+           }
+       }
+   }
+   policy-statement ibg-export {
+       term 1 {
+           then {
+               next-hop self;
+           }
+       }
+   }
[edit policy-options]
+   as-path as179 "^179$";
+   as-path as212 "^212$";
```

Verification

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
root@srx-2> show route 1.1.1.1

inet.0: 142 destinations, 169 routes (142 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.0/24          *[BGP/170] 03:09:28, localpref 500, from 6.6.1.1
                      AS path: 212 83 179 1144 I
                    > to 6.6.6.1 via ge-0/0/4.101
                      to 6.7.7.1 via ge-0/0/5.102

root@srx-2> show route 8.8.8.8

inet.0: 145 destinations, 200 routes (145 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

8.0.0.0/8           *[BGP/170] 00:00:00, localpref 1000, from 21.100.0.1
                      AS path: 735 I
                    > to 21.1.10.5 via ge-0/0/1.70
                    [BGP/170] 03:09:29, localpref 500, from 6.6.1.1
                      AS path: 212 83 179 1144 775 I
                      to 6.6.6.1 via ge-0/0/4.101
                    > to 6.7.7.1 via ge-0/0/5.102

root@srx-2> show route 7.7.7.7

inet.0: 145 destinations, 200 routes (145 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

7.0.0.0/8           *[BGP/170] 00:00:17, MED 1000, localpref 1000, from
21.100.0.1
                      AS path: 179 I
                    > to 21.1.10.5 via ge-0/0/1.70
                    [BGP/170] 03:09:46, localpref 500, from 6.6.1.1
                      AS path: 212 83 179 I
                    > to 6.6.6.1 via ge-0/0/4.101
                      to 6.7.7.1 via ge-0/0/5.102
```

10) Set up the connection with Peer 2 using the information from table 5.1. Peer 2 thinks we are using AS654. Ensure that AS654 is not seen in the AS-path when advertising routes out towards EBGP neighbors. Advertise only the 15.0.0.0/8 prefix to Transit 1 when it has our local community attached. All routes should be advertised to Peer 1. Make sure we can learn a maximum of 10 prefixes from Peer 2.

```
[edit]
root@srx-2# show | compare
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit protocols BGP]
     group as179 { ... }
+    group as99 {
+        type external;
+        import as99-import;
+        neighbor 21.24.1.2 {
+            family inet {
+                unicast {
+                    prefix-limit {
+                        maximum 10;
+                    }
+                }
+            }
+            peer-as 99;
+            local-as 654 private;
+        }
+    }
 [edit policy-options]
+   policy-statement as99-import {
+       term 1 {
+           from {
+               route-filter 15.0.0.0/8 exact;
+           }
+           then {
+               community set local;
+           }
+       }
+       term 2 {
+           then accept;
+       }
+   }
```

Verification

```
[edit]
root@srx-2# run show BGP summary | match 99
21.24.1.2               99         41         43       0       0       12:37
1/67/66/0            0/0/0/0


[edit]
root@srx-2# run show route 15/8 detail

inet.0: 148 destinations, 269 routes (148 active, 0 holddown, 1 hidden)
15.0.0.0/8 (2 entries, 1 announced)
        *BGP    Preference: 170/-501
                Next hop type: Indirect
                Address: 0x934d1c8
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                    Next-hop reference count: 220
                    Source: 6.6.1.1
                    Next hop type: Router, Next hop index: 262142
                    Next hop: 6.6.6.1 via ge-0/0/4.101
                    Next hop: 6.7.7.1 via ge-0/0/5.102, selected
                    Protocol next hop: 6.6.1.1
                    Indirect next hop: 949d3f0 262143
                    State: <Active Ext>
                    Local AS:   653 Peer AS:   212
                    Age: 4:44:59    Metric2: 0
                    Task: BGP_212.6.6.1.1+63036
                    Announcement bits (4): 0-KRT 10-KRT 11-BGP_RT_Background 12-
Resolve tree 4
                    AS path: 212 824 99 ?
                    Accepted
                    Localpref: 500
                    Router ID: 6.6.0.1
            BGP     Preference: 170/-101
                    Next hop type: Router, Next hop index: 690
                    Address: 0x934fa00
                    Next-hop reference count: 69
                    Source: 21.24.1.2
                    Next hop: 21.24.1.2 via ge-0/0/5.70, selected
                    State: <Ext>
                    Inactive reason: Local Preference
                    Local AS:   653 Peer AS:    99
                    Age: 12:59
                    Task: BGP_99_654.21.24.1.2+179
                    AS path: 99 99 99 99 99 I
                    Communities: 653:101
                    Accepted
                    Localpref: 100
                    Router ID: 15.1.1.1

[edit]
root@srx-2# run show route receive-protocol BGP 21.24.1.2 active-path

inet.0: 148 destinations, 269 routes (148 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 4.40.4.0/24             21.24.1.2                               99 99 99 99
99 I
```

Verifiy AS path to peer 1

```
[edit]
root@srx-1-re-0# run show route advertising-protocol BGP 21.21.21.2
4.40.4.0/24
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
inet.0: 145 destinations, 332 routes (145 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 4.40.4.0/24             Self                                    653 99 99 99
99 99 I
```

11) Ensure that all traffic destined to Peer 2's network is using the local peering session with Peer 2. Peer 2 instructed us to make sure our prefix (21.0.0.0/8) and Peer 1's prefix (11.0.0.0/8) is tagged with community 99:250. All other routes should be advertised without this community. Only accept prefixes that originated from within Peer 2's AS. Complete this task without the use of an AS-Path filter.

```
root@srx-2# show | compare
[edit protocols BGP group as99]
+    export as99-export;
[edit policy-options]
+   policy-statement as99-export {
+       term 1 {
+           from {
+               route-filter 21.0.0.0/8 exact;
+           }
+           then {
+               community set as99;
+           }
+       }
+       term 2 {
+           from {
+               route-filter 11.0.0.0/8 exact;
+           }
+           then {
+               community set as99;
+           }
+       }
+       term 3 {
+           from {
+               protocol BGP;
+               route-filter 0.0.0.0/0 upto /24;
+            }
+           then accept;
+       }
+   }
[edit policy-options policy-statement as99-import term 1 then]
+      local-preference 1000;
[edit policy-options policy-statement as99-import term 2]
+     from {
+         route-filter 15.0.0.0/8 orlonger;
+     }
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
[edit policy-options policy-statement as99-import term 2 then]
+      local-preference 1000;
[edit policy-options policy-statement as99-import]
     term 2 { ... }
+    term 3 {
+        then reject;
+    }
[edit policy-options]
+   community as99 members 99:250;
```

Verification

```
root@srx-2> show route receive-protocol BGP 21.24.1.2

inet.0: 148 destinations, 301 routes (147 active, 0 holddown, 93 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 15.0.0.0/8              21.24.1.2                               99 99 99 99
99 I
* 15.100.1.0/24           21.24.1.2                               99 99 99 99
99 I
* 15.100.10.0/24          21.24.1.2                               99 99 99 99
99 I
* 15.100.20.0/24          21.24.1.2                               99 99 99 99
99 I
* 15.100.30.0/24          21.24.1.2                               99 99 99 99
99 I
* 15.100.40.0/24          21.24.1.2                               99 99 99 99
99 I

root@srx-2> show route 15.0.0.0/8 detail | match local
                Local AS:   653 Peer AS:    99
                Localpref: 1000


root@srx-2> show route advertising-protocol BGP 21.24.1.2 detail 21/8

inet.0: 148 destinations, 301 routes (147 active, 0 holddown, 93 hidden)
* 21.0.0.0/8 (1 entry, 1 announced)
 BGP group as99 type External
     Nexthop: Self
     AS path: [654] I (LocalAgg)
     Communities: 99:250

root@srx-2> show route advertising-protocol BGP 21.24.1.2 detail 11/8

inet.0: 148 destinations, 301 routes (147 active, 0 holddown, 93 hidden)
* 11.0.0.0/8 (1 entry, 1 announced)
 BGP group as99 type External
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
     Nexthop: Self
     AS path: 653 65522 64555 I
     Communities: 99:250

root@srx-2> show route advertising-protocol BGP 21.24.1.2 detail 1/8

inet.0: 148 destinations, 301 routes (147 active, 0 holddown, 93 hidden)
* 1.0.0.0/8 (2 entries, 1 announced)
 BGP group as99 type External
     Nexthop: Self
     AS path: 653 212 83 179 1144 I
```

12) Set up IPv6 peering with Transit 2. Use the Link-local address for the peering session.

On R1

```
root@srx-1-re-0> ping ff02::1 interface ge-0/0/4.99
PING6(56=40+8+8 bytes) fe80::ef13:7332:1911:aed1 --> ff02::1
16 bytes from fe80::20c:2900:633b:cf30, icmp_seq=0 hlim=64 time=9.498 ms
16 bytes from fe80::20c:2900:633b:cf30, icmp_seq=1 hlim=64 time=9.995 ms
16 bytes from fe80::20c:2900:633b:cf30, icmp_seq=2 hlim=64 time=8.280 ms
^C
--- ff02::1 ping6 statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 8.280/9.258/9.995/0.720 ms

[edit]
root@srx-1-re-0# show | compare
[edit protocols BGP]
     group ibgp { ... }
+    group as179-ipv6 {
+        type external;
+        neighbor fe80::20c:2900:633b:cf30 {
+            local-interface ge-0/0/4.99;
+            peer-as 179;
+        }
+    }
```

On R2

```
root@srx-2> ping ff02::1 interface ge-0/0/4.12
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
PING6(56=40+8+8 bytes) fe80::ef13:7332:1911:aed5 --> ff02::1
16 bytes from fe80::20c:2900:633b:af30, icmp_seq=0 hlim=64 time=12.397 ms
16 bytes from fe80::20c:2900:633b:af30, icmp_seq=1 hlim=64 time=11.330 ms
^C
--- ff02::1 ping6 statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 6.171/9.966/12.397/2.719 ms

root@srx-2# show | compare
[edit protocols BGP]
     group as99 { ... }
+    group as179-ipv6 {
+        type external;
+        neighbor fe80::20c:2900:633b:af30 {
+            local-interface ge-0/0/4.12;
+            peer-as 179;
+        }
+    }
```

Verification

```
root@srx-1-re-0> show BGP summary | match fe80
fe80::20c:2900:633b:cf30         179          6          5       0       0
1:05 Establ

root@srx-1-re-0> show route receive-protocol BGP fe80::20c:2900:633b:cf30

inet6.0: 64 destinations, 74 routes (64 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2002:a179:ef:5411::/64
*                         fe80::20c:2900:633b:cf30               179 I
  2002:a179:ef:5412::/64
*                         fe80::20c:2900:633b:cf30               179 I
  2002:a179:ef:5413::/64
*                         fe80::20c:2900:633b:cf30               179 I
  2002:ef54:ac11:56b2::/64
*                         fe80::20c:2900:633b:cf30               179 1144 I
  2002:ef54:ac11:56b3::/64
*                         fe80::20c:2900:633b:cf30               179 1144 I
  2002:ef54:ac11:99a1::/64
*                         fe80::20c:2900:633b:cf30               179 1144 I
  2002:ef54:ac11:99a2::/64
*                         fe80::20c:2900:633b:cf30               179 1144 I
  2002:ef54:ac11:99a3::/64
*                         fe80::20c:2900:633b:cf30               179 1144 I
  2002:ef54:ac11:fe11::/64
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
*                          fe80::20c:2900:633b:cf30                179 1144 I
  2002:ef54:ac11:fe12::/64
*                          fe80::20c:2900:633b:cf30                179 1144 I

root@srx-2> show BGP summary | match fe80
fe80::20c:2900:633b:af30      179       3       3     0     0       15 Establ
```

13) Make sure that both inbound as well as outbound IPv6 traffic will traverse the link between R2 and Transit 2. Advertise your IPv6 prefix, 2002:95:b::/64, to all IPv6 peers. Do so using an aggregate.

On SRX 1

```
root@srx-1-re-0# show | compare
[edit routing-options]
+   rib inet6.0 {
+       aggregate {
+           route 2002:95:b::/64;
+       }
+   }
[edit protocols BGP group as179-ipv6]
+    import as179-ipv6-import;
+    export as179-ipv6-export;
[edit protocols BGP]
     group as179-ipv6 { ... }
+    group ibgp-ipv6 {
+        type internal;
+        neighbor 2002:95:b::100:2 {
+            local-address 2002:95:b::100:1;
+        }
+    }
[edit policy-options]
+   policy-statement as179-ipv6-export {
+       term 1 {
+           from {
+               protocol aggregate;
+               route-filter 2002:95:b::/64 exact;
+           }
+           then {
+               metric 1000;
+               accept;
+           }
+       }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+    }
+    policy-statement as179-ipv6-import {
+        term 1 {
+            then {
+                local-preference 100;
+            }
+        }
+    }
[edit policy-options policy-statement as_735-export]
     term 1 { ... }
+    term 2 {
+        from {
+            protocol aggregate;
+            route-filter 2002:95:b::/64 exact;
+        }
+        then accept;
+    }
     term last { ... }
```

On SRX 2

```
[edit]

root@srx-2# show | compare
[edit routing-options]
+    rib inet6.0 {
+        aggregate {
+            route 2002:95:b::/64;
+        }
+    }
[edit protocols BGP group as179-ipv6]
+    import as179-ipv6-import;
+    export as179-ipv6-export;
[edit protocols BGP]
     group as179-ipv6 { ... }
+    group ibgp-ipv6 {
+        type internal;
+        neighbor 2002:95:b::100:1 {
+            local-address 2002:95:b::100:2;
+        }
+    }
[edit policy-options]
+    policy-statement as179-ipv6-export {
+        term 1 {
+            from {
+                protocol aggregate;
+                route-filter 2002:95:b::/64 exact;
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+           }
+           then {
+               metric 100;
+               accept;
+           }
+       }
+   }
+   policy-statement as179-ipv6-import {
+       term 1 {
+           then {
+               local-preference 1000;
+           }
+       }
+   }
```

Verification

```
root@srx-2> show route 2002:ef54:ac11:56b2::1

inet6.0: 59 destinations, 70 routes (59 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2002:ef54:ac11:56b2::/64
                   *[BGP/170] 01:56:15, localpref 1000
                      AS path: 179 1144 I
                    > to fe80::20c:2900:633b:af30 via ge-0/0/4.12

root@srx-1-re-0> show route 2002:ef54:ac11:56b2::1

inet6.0: 65 destinations, 85 routes (65 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2002:ef54:ac11:56b2::/64
                   *[BGP/170] 00:01:15, localpref 1000, from 2002:95:b::100:2
                      AS path: 179 1144 I
                    > to fe80::20c:2900:450d:3061 via ge-0/0/1.69
                    [BGP/170] 02:07:10, localpref 100
                      AS path: 179 1144 I
                    > to fe80::20c:2900:633b:cf30 via ge-0/0/4.99

root@srx-1-re-0> traceroute 2002:ef54:ac11:56b2::1
traceroute6 to 2002:ef54:ac11:56b2::1 (2002:ef54:ac11:56b2::1) from
2002:95:b::a653:1, 64 hops max, 12 byte packets
 1  2002:95:b::a653:2 (2002:95:b::a653:2)  1027.344 ms  5.717 ms  5.394 ms
 2  2002:a179:ef:5413::5 (2002:a179:ef:5413::5)  23.281 ms  11.710 ms  15.329
ms
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
 3  2002:ef54:ac11:56b2::1 (2002:ef54:ac11:56b2::1)  13.253 ms  10.900 ms
10.001 ms
```

Transit_2.inet6.0: 32 destinations, 35 routes (32 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2002:95:b::/64 *[BGP/170] 00:07:21, **MED 100**, localpref 100
AS path: **653** I
> to **fe80::ef13:7332:1911:aed5** via **ge-0/0/2.12**
[BGP/170] 00:07:18, **MED 1000**, **localpref 100**
AS path: 653 I
> to **fe80::ef13:7332:1911:aed1** via ge-0/0/1.99

14) SRX 1 should be the preferred route for inbound IPv4 traffic passing the Transit 2 network. You are not allowed to use a policy to complete this task.

```
[edit]
root@srx-1-re-0# show | compare
[edit protocols BGP group AS_179 neighbor 7.16.1.1]
+      metric-out 10;
```

Verification

```
root@VR_Router_Lab3> show route table Transit_2.inet.0 21

Transit_2.inet.0: 125 destinations, 126 routes (125 active, 0 holddown, 0
hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

21.0.0.0/8          *[BGP/170] 00:06:44, MED 10, localpref 100
                       AS path: 653 I
                     > to 7.16.1.2 via ge-0/0/1.99
```

15) You want to load balance traffic across your transit providers. Advertise 21.50.22.0/24 and 21.50.42.0/24 to both Transit 3 and Transit 2. Transit 3 should be preferred for 21.50.22.0/24 and Transit 2 should be preferred for 21.50.42.0/24. Complete this task without manipulating the AS-path. You are allowed to create 2 static routes on SRX 1.

```
[edit routing-options static]
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+    route 21.50.22.0/24 discard;
+    route 21.50.42.0/24 discard;
[edit protocols BGP group AS_179]
-    export as653-export;
+    export [ load-balance-transit2 as653-export ];
[edit protocols BGP group AS_Transit3]
-    export [ as_735-export as653-export ];
+    export [ as_735-export load-balance-transit3 as653-export ];
[edit policy-options]
+   policy-statement load-balance-transit2 {
+       term 1 {
+           from {
+               protocol static;
+               route-filter 21.50.22.0/24 exact;
+           }
+           then {
+               origin incomplete;
+               accept;
+           }
+       }
+       term 2 {
+           from {
+               protocol static;
+               route-filter 21.50.42.0/24 exact;
+           }
+           then {
+               origin igp;
+               accept;
+           }
+       }
+       term 3 {
+           then next policy;
+       }
+   }
+   policy-statement load-balance-transit3 {
+       term 1 {
+           from {
+               protocol static;
+               route-filter 21.50.22.0/24 exact;
+           }
+           then {
+               origin igp;
+               accept;
+           }
+       }
+       term 2 {
+           from {
+               protocol static;
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+               route-filter 21.50.42.0/24 exact;
+           }
+           then {
+               origin incomplete;
+               accept;
+           }
+       }
+       term 3 {
+           then next policy;
+       }
+   }
```

Verification

```
root@srx-1-re-0> show route advertising-protocol BGP 7.16.1.1

inet.0: 146 destinations, 334 routes (146 active, 0 holddown, 3 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 15.0.0.0/8              Self                 10                 99 99 99 99
99 I
* 21.0.0.0/8              Self                 10                 I
* 21.50.22.0/24           Self                 10                 ?
* 21.50.42.0/24           Self                 10                 I

root@srx-1-re-0> show route advertising-protocol BGP 8.254.157.1 table inet.0

inet.0: 146 destinations, 334 routes (146 active, 0 holddown, 3 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 11.0.0.0/8              Self                                    I
* 15.0.0.0/8              Self                                    99 99 99 99
99 I
* 21.0.0.0/8              Self                                    I
* 21.50.22.0/24           Self                                    I
* 21.50.42.0/24           Self                                    ?
```

16) For security reasons, you need to make sure that routes with a private ASN in the AS-path are not accepted from any peer. You also need to reject routes with a prefix length greater than /24. Once executed successfully, the 2.2.2.0/24 network should be hidden.

```
[edit]
root@srx-1-re-0# run show route 2.2.2.0/24

inet.0: 146 destinations, 334 routes (146 active, 0 holddown, 3 hidden)
+ = Active Route, - = Last Active, * = Both
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

2.2.2.0/24 *[BGP/170] 09:17:31, MED 1000, localpref 100
AS path: 179 64513 83 I
> to 7.16.1.1 via ge-0/0/4.99
[BGP/170] 09:17:22, localpref 100
AS path: 735 775 1144 179 **64513** 83 I
> to 8.254.157.1 via ge-0/0/4.59

On SRX 1

```
[edit]
root@srx-1-re-0# show | compare

[edit protocols BGP group AS_179]
-    import as179-import;
+    import [ block-private-as as179-import ];
[edit protocols BGP group AS_Transit3]
-    import as735-import;
+    import [ block-private-as as735-import ];
[edit policy-options]
+   policy-statement block-private-as {
+       term 1 {
+           from as-path private-ASN;
+           then reject;
+       }
+       term 2 {
+           then next policy;
+       }
+   }
[edit policy-options]
    as-path as735 { ... }
+   as-path private-ASN ".* 64512-65534 .*";
```

On SRX 2

```
[edit]
root@srx-2# show | compare
[edit protocols BGP group AS_212]
-    import as212-import;
+    import [ block-private-as as212-import ];
[edit protocols BGP group as179]
-    import as179-import;
+    import [ block-private-as as179-import ];
[edit policy-options]
+   policy-statement block-private-as {
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+       term 1 {
+           from as-path private-ASN;
+           then reject;
+       }
+       term 2 {
+           then next policy;
+       }
+   }
[edit policy-options]
+   as-path private-ASN ".* 64512-65534 .*";
    as-path as179 { ... }
```

Verification

```
[edit]
root@srx-1-re-0# run show route 2.2.2.0/24

inet.0: 146 destinations, 334 routes (145 active, 0 holddown, 5 hidden)
```

17) You have been told that the Transit 2 IPv4 connection towards R2 is not working. Solve the problem.

```
root@srx-2> show BGP summary | match Active

Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
7.16.15.1               179          0          0       0       0     6:25:11
Active

root@srx-2> ping 7.16.15.1 rapid
PING 7.16.15.1 (7.16.15.1): 56 data bytes
!!!!!
--- 7.16.15.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.513/4.968/5.968/1.039 ms

[edit]
root@srx-2# show | compare
[edit protocols BGP group as179]
+     traceoptions {
+         file BGP-debug;
+         flag all;
+     }

root@srx-2> show log BGP-debug
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Mar  3 01:42:22.721372 task_addr_local: task BGP_179.7.16.15.1 address
7.16.15.2
Mar  3 01:42:22.721436 task_connect: task BGP_179.7.16.15.1+179 addr
7.16.15.1+179task_timer_reset: reset BGP_179.7.16.15.1+179_Connect
Mar  3 01:42:22.721456 task_timer_set_oneshot_latest: timer
BGP_179.7.16.15.1+179_Connect interval set to 32
Mar  3 01:42:54.723457 bgp_connect_timeout: BGP_179.7.16.15.1+179_Connect
Mar  3 01:42:54.723582 task_timer_reset: reset BGP_179.7.16.15.1+179_Connect
Mar  3 01:42:55.685472 bgp_down_peer_timeout: BGP_179.7.16.15.1+179_Down Peer
Timer
Mar  3 01:42:55.685643 task_timer_delete: BGP_179.7.16.15.1+179_Down Peer
Timer <Processing>



root@srx-2> show system statistics tcp | match auth
         0 send packets dropped by TCP due to auth errors
         10 rcv packets dropped by TCP due to auth errors

[edit]
root@srx-2# show | compare
[edit protocols BGP group as179 neighbor 7.16.15.1]
-     authentication-key "$9$xoO-b2.mT3nCZUqfQzCAIEcSK8X7-"; ## SECRET-DATA
+     authentication-key "$9$wV2oGzF/tpBk.Qn9CB1revMxd"; ## SECRET-DATA
```

Verification

```
root@srx-2> show BGP summary | match 179
7.16.15.1               179         16         14       0       0           6
3/97/97/0            0/0/0/0
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 6: Multicast Routing

1) EX1, EX2 and the EX_VC should be enabled for SSM. They should also stop sending multicast traffic as soon as the last receiver has finished receiving traffic.

```
{master:0}[edit]
root@ex_vc# show | compare
[edit protocols]
+    igmp {
+        interface all {
+            version 3;
+            immediate-leave;
+        }
+    }
```

Verification

```
root@ex_vc> show igmp interface
Interface: ge-0/0/14.69
    Querier: 21.1.76.9
    State:         Up Timeout:     214 Version:  3 Groups:      0
    Immediate leave: On
    Promiscuous mode: Off
    Passive: Off
Interface: ge-0/0/7.0
    Querier: 21.30.0.2
    State:         Up Timeout:    None Version:  3 Groups:      4
    Immediate leave: On
    Promiscuous mode: Off
    Passive: Off
Interface: ge-0/0/22.0
    Querier: 21.30.10.8
    State:         Up Timeout:    None Version:  3 Groups:      3
    Immediate leave: On
    Promiscuous mode: Off
    Passive: Off
Interface: vlan.200
    Querier: 21.17.200.1
    State:         Up Timeout:     223 Version:  3 Groups:      0
    Immediate leave: On
    Promiscuous mode: Off
    Passive: Off
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

2) A PIM neighbor relationship between R1 and R2 is not forming. Fix the issue.

```
root@srx-1-re-0> show pim neighbors
Instance: PIM.master

unknown command.
root@srx-1-re-0> show configuration protocols pim | display inheritance
interface all {
    ##
    ## 'inet' was inherited from group 'inet'
    ##
    family inet {
        ##
        ## 'disable' was inherited from group 'inet'
        ##
        disable;
    }
    ##
    ## 'inet6' was inherited from group 'inet'
    ##
    family inet6 {
        ##
        ## 'disable' was inherited from group 'inet'
        ##
        disable;
    }
}
```

On R1

```
[edit]
root@srx-1-re-0# show | compare
[edit groups inet protocols pim interface <*>]
-       family inet {
-           disable;
-       }
```

On R2

```
[edit]
root@srx-2# show | compare
[edit groups inet protocols pim interface <*>]
-       family inet {
-           disable;
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
-       }
[edit protocols pim]
     interface ge-0/0/1.70 { ... }
+    interface all;
```

Verification

```
[edit]
root@srx-1-re-0# run show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Interface           IP V Mode        Option      Uptime Neighbor addr
ge-0/0/1.69          4 2             HPLGT      00:00:51 21.1.10.2
ge-0/0/1.70          4 2             HPLGT      00:00:51 21.1.10.6
```

3) Ensure that all devices are configured to run PIM sparse mode on every interface.

On R3 and R4

```
[edit]
root@srx-3# show | compare
[edit routing-instances R3_A protocols]
+     pim {
+         interface all {
+             mode sparse;
+         }
+     }
[edit routing-instances R3_B protocols]
+     pim {
+         interface all {
+             mode sparse;
+         }
+     }
```

On EX switches and R4

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
{master:0}[edit]
root@ex-1# show | compare
[edit protocols]
+   pim {
+       interface all {
+           mode sparse;
+       }
+   }
```

Verification

```
[edit]
root@srx-1-re-0# run show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Interface           IP V Mode        Option      Uptime Neighbor addr
ge-0/0/1.69          4 2             HPLGT      00:05:54 21.1.10.2
ge-0/0/1.70          4 2             HPLGT      00:05:54 21.1.10.6
ge-0/0/4.40          4 2             HPLGT      00:00:58 21.1.10.18
ge-0/0/4.41          4 2             HPLGT      00:02:15 21.1.10.22
ge-0/0/5.30          4 2             HPLGT      00:00:58 21.1.10.10
ge-0/0/5.31          4 2             HPLGT      00:02:15 21.1.10.14
```

4) Configure anycast RP on R1 and R2. Use address 21.100.0.100 as the anycast RP address. Use not other protocol than the PIM protocol.

```
[edit]
root@srx-1-re-0# show | compare
[edit protocols pim]
+    rp {
+        local {
+            family inet {
+                address 21.100.0.100;
+                anycast-pim {
+                    rp-set {
+                        address 21.100.0.2;
+                    }
+                    local-address 21.100.0.1;
+                }
+            }
+        }
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+    }
[edit]
root@srx-2# show | compare
[edit interfaces lo0 unit 0 family inet]
        address 21.100.0.2/32 { ... }
+       address 21.100.0.100/32;
[edit protocols pim]
+    rp {
+        local {
+            family inet {
+                address 21.100.0.100;
+                anycast-pim {
+                    rp-set {
+                        address 21.100.0.1;
+                    }
+                    local-address 21.100.0.2;
+                }
+            }
+        }
+    }
```

Verification

```
[edit]
root@srx-1-re-0# run show pim rps extensive
Instance: PIM.master
Address family INET

RP: 21.100.0.100
Learned via: static configuration
Time Active: 01:12:34
Holdtime: 0
Device Index: 131
Subunit: 32769
Interface: ppd0.32769
Static RP Override: Off
Group Ranges:
        224.0.0.0/4
Anycast PIM rpset:
        21.100.0.2
Anycast PIM local address used: 21.100.0.1

Address family INET6
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

5) Ensure that the RP address is propagated across the entire network without the use of any dense groups. SRX_4A and SRX_3B should both propagate the RP information. SRX4_A should be the primary.

```
[edit]
root@srx-4# show | compare
[edit interfaces lo0 unit 1 family inet]
!        inactive: address 21.100.0.100/32 { ... }
[edit routing-instances R4_A protocols pim]
+      rp {
+          bootstrap {
+              family inet {
+                  priority 200;
+              }
+          }
+      }

[edit]
root@srx-3# show | compare
[edit routing-instances R3_B protocols pim]
+      rp {
+          bootstrap {
+              family inet {
+                  priority 150;
+              }
+          }
+      }
```

Verification

```
[edit]
root@srx-4# run show pim bootstrap instance R4_A
Instance: PIM.R4_A

BSR                     Pri Local address           Pri State      Timeout
21.100.0.4              200 21.100.0.4              200 Elected         33
None                      0 2002:95:b::100:4          0 InEligible       0

[edit]
root@srx-4# run show pim rps instance R4_A
Instance: PIM.R4_A

Address family INET
RP address      Type        Mode   Holdtime Timeout Groups Group prefixes
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
21.100.0.100    bootstrap   sparse      150     117      0 224.0.0.0/4

[edit]
root@srx-3# run show pim rps instance R3_B
Instance: PIM.R3_B

Address family INET
RP address      Type        Mode   Holdtime Timeout Groups Group prefixes
21.100.0.100    bootstrap   sparse      150     141      0 224.0.0.0/4
```

6) Peer 1 has been configured to run MSDP with SRX1. Set up the MSDP peering to Peer1 on SRX1.

```
[edit]
root@srx-1-re-0# show | compare
[edit protocols]
+   msdp {
+       group peer1 {
+           local-address 21.21.21.1;
+           peer 21.21.21.2;
+       }
+   }


[edit]
root@srx-1-re-0# run show msdp
Peer address    Local address   State       Last up/down Peer-Group   SA Count
21.21.21.2      21.21.21.1      Inactive        00:00:08 peer1        0/0

[edit]
root@srx-1-re-0# run show system connections | match 639
tcp4       0      0  21.21.21.1.639
11.0.1.1.59794                             TIME_WAIT
tcp4       0      0  21.21.21.1.639
11.0.1.1.49336                             TIME_WAIT

[edit]
root@srx-1-re-0# show | compare
[edit protocols msdp group peer1]
+     peer 11.0.1.1;
-     peer 21.21.21.2;
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification

```
[edit]
root@srx-1-re-0# run show msdp detail
Peer: 11.0.1.1
Peer Group: peer1
Local address: 21.21.21.1
State: Established
Peer Connect Retries: 0
State timer expires: 24
Peer Times out: 51
SA accepted: 0
SA received: 0

[edit]
root@srx-1-re-0# run show system connections | match 639
tcp4       0      0  21.21.21.1.639
11.0.1.1.61373                                 ESTABLISHED
```

7) Configure the EX_VC to use the inet.2 table for RPF lookups.

```
{master:0}[edit]
root@ex_vc# show | compare
[edit routing-options]
+   interface-routes {
+       rib-group inet rpf-table;
+   }
+   rib-groups {
+       rpf-rib {
+           import-rib inet.2;
+       }
+       rpf-table {
+           import-rib [ inet.0 inet.2 ];
+       }
+   }
[edit protocols ospf]
+   rib-group rpf-table;
[edit protocols pim]
+    rib-group inet rpf-rib;
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Verification

```
{master:0}
root@ex_vc> show multicast rpf summary
Multicast RPF table: inet.2 , 41 entries

Multicast RPF table: inet6.0 , 2 entries

{master:0}
root@ex_vc> show multicast rpf inet
Multicast RPF table: inet.2 , 41 entries

10.10.0.0/24
    Protocol: Direct
    Interface: me0.0

10.10.0.7/32
    Protocol: Local

21.1.10.4/30
    Protocol: OSPF
    Interface: ge-0/0/14.69
    Neighbor: 21.1.76.9

21.1.10.8/30
    Protocol: OSPF
    Interface: ge-0/0/14.69
    Neighbor: 21.1.76.9

{master:0}
root@ex_vc> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.2.127.254
    Source: *
    RP: 21.100.0.100
    Flags: sparse,rptree,wildcard
    Upstream interface: ge-0/0/14.69

Group: 239.17.200.1
    Source: *
    RP: 21.100.0.100
    Flags: sparse,rptree,wildcard
    Upstream interface: ge-0/0/14.69
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Group: 239.17.200.2
    Source: *
    RP: 21.100.0.100
    Flags: sparse,rptree,wildcard
    Upstream interface: ge-0/0/14.69
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Task 7: Class of Service

1) Configure forwarding-classes on all of the SRX devices according to the information provided in the table below:

forwarding-classes	Queue	priority	rate
Data	0	Low	50%
Assured-Data	1	medium-high	25%
VoIP	2	high	10%
Network-Control	3	strict-high	10%
Trading	4	medium-high	5%

Table 7.1

```
[edit]
root@srx-1-re-0# show | compare
[edit]
+  class-of-service {
+      forwarding-classes {
+          queue 0 Data;
+          queue 1 Assured-Data;
+          queue 2 VoIP;
+          queue 3 Network-Control;
+          queue 4 Trading;
+      }
+      scheduler-maps {
+          cos {
+              forwarding-class Data scheduler Data;
+              forwarding-class Assured-Data scheduler Assured-Data;
+              forwarding-class VoIP scheduler VoIP;
+              forwarding-class Network-Control scheduler Network-Control;
+              forwarding-class Trading scheduler Trading;
+          }
+      }
+      schedulers {
+          Data {
+              transmit-rate percent 50;
+              priority low;
+          }
+          Assured-Data {
+              transmit-rate percent 25;
+              priority medium-high;
+          }
+          VoIP {
+              transmit-rate percent 10;
+              priority high;
+          }
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+          Network-Control {
+              transmit-rate percent 10;
+              priority strict-high;
+          }
+          Trading {
+              transmit-rate percent 5;
+              priority medium-high;
+          }
+      }
+  }
```

Verification

```
[edit]
root@srx-1-re-0# run show interfaces ge-0/0/1 extensive | find queue
  CoS queues      : 8 supported, 8 maximum usable queues
  Hold-times      : Up 0 ms, Down 0 ms
  Current address: 00:0c:29:fd:79:be, Hardware address: 00:0c:29:fd:79:be
  Last flapped    : 2016-03-10 12:41:23 UTC (04:58:37 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                   0 bps
   Output bytes  :              3020028                1064 bps
   Input  packets:                    0                   0 pps
   Output packets:                22542                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 248, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped
packets
    0 Data                           0                   0
0
    1 Assured-Data                   0                   0
0
    2 VoIP                           0                   0
0
    3 Network-Cont               22542               22542
0
    4 Trading                        0                   0
0
  Queue number:         Mapped forwarding classes
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    0                   Data
    1                   Assured-Data
    2                   VoIP
    3                   Network-Control
    4                   Trading
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                         85275          2118734
    Total packets                          597            22541
    Unicast packets                        597            22541
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
    Input packet count                     597
    Input packet rejects                   349
    Input DA rejects                       349
    Input SA rejects                         0
    Output packet count                                     0
    Output packet pad count                                 0
    Output packet error count                               0
    CAM destination filters: 8, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: None, Remote fault: OK
    Local resolution:
        Flow control: None, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 0
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                                   %           bps     %           usec
  0 Data                          50     500000000     r              0      low
none
  1 Assured-Data                  25     250000000     r              0     high
none
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
 2 VoIP                  10      100000000     r              0      medium-
high none
 3 Network-Control       10      100000000     r              0      strict-
high none
 4 Trading                5       50000000     r              0      medium-
high none
```

2) Configure forwarding-classes all of the EX devices according to the information provided in the table below:

forwarding-classes	Queue	priority	rate
Data	0	Low	50%
Assured-Data	1	Low	25%
VoIP	4	strict-high	-
Network-Control	3	strict-high	-
Trading	2	low	5%

Table 7.2

```
root@ex1# show | compare
[edit]
+  class-of-service {
+      forwarding-classes {
+          queue 0 Data;
+          queue 1 Assured-Data;
+          queue 2 Trading;
+          queue 3 Network-Control;
+          queue 4 VoIP;
+      }
+      scheduler-maps {
+          cos {
+              forwarding-class Data scheduler Data;
+              forwarding-class Assured-Data scheduler Assured-Data;
+              forwarding-class VoIP scheduler VoIP;
+              forwarding-class Network-Control scheduler Network-Control;
+              forwarding-class Trading scheduler Trading;
+          }
+      }
+      schedulers {
+          Data {
+              transmit-rate percent 50;
+              priority low;
+          }
+          Assured-Data {
+              transmit-rate percent 25;
+              priority low;
+          }
+          VoIP {
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+              transmit-rate percent 10;
+              priority strict-high;
+          }
+          Network-Control {
+              transmit-rate percent 10;
+              priority strict-high;
+          }
+          Trading {
+              transmit-rate percent 5;
+              priority low;
+          }
+      }
+  }

{master:0}[edit]
root@ex1# set class-of-service tri-color
```

3) Make sure that the queues on every SRX device have the correct RED configuration applied to them. The following table provides you with the necessary parameters:

Drop profile	priority	fill level	Drop probability
Low-data	medium-high	10	25
		50	50
		90	90
High-data	Low	10	15
		50	25
		90	35

Table 7.3

The drop profile Low-data should be applied to the forwarding classes Data and Trading. The drop profile High-data should be applied to the forwarding class VoIP. Make sure that the queue configuration is applied to all interfaces on the SRX devices.

```
[edit class-of-service]
root@srx-1-re-0# show | compare
[edit class-of-service]
+ drop-profiles {
+     low-data {
+         interpolate {
+             fill-level [ 10 50 90 ];
+             drop-probability [ 25 50 90 ];
+         }
+     }
+     high-data {
+         interpolate {
+             fill-level [ 10 50 90 ];
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
+             drop-probability [ 15 25 35 ];
+         }
+     }
+ }
+ interfaces {
+     ge-* {
+         scheduler-map cos;
+     }
+ }
[edit class-of-service schedulers Data]
+   drop-profile-map loss-priority medium-high protocol any drop-profile low-
data;
[edit class-of-service schedulers VoIP]
+   drop-profile-map loss-priority low protocol any drop-profile high-data;
[edit class-of-service schedulers Trading]
+   drop-profile-map loss-priority medium-high protocol any drop-profile low-
data;
```

Verification

```
[edit]
root@srx-1-re-0# run show class-of-service interface
Physical interface: ge-0/0/0, Index: 133
Queues supported: 8, Queues in use: 5
  Scheduler map: cos, Index: 33170
  Congestion-notification: Disabled

Physical interface: ge-0/0/1, Index: 134
Queues supported: 8, Queues in use: 5
  Scheduler map: cos, Index: 33170
  Congestion-notification: Disabled

  Logical interface: ge-0/0/1.32767, Index: 77

  Logical interface: ge-0/0/1.69, Index: 71
    Object                  Name                    Type
Index
    Classifier              dscp-ipv6-compatibility dscp-ipv6
9
    Classifier              ipprec-compatibility    IP
13

  Logical interface: ge-0/0/1.70, Index: 72
    Object                  Name                    Type
Index
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    Classifier              dscp-ipv6-compatibility dscp-ipv6
9
    Classifier              ipprec-compatibility   IP
13

root@ srx-1-re-0> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: high-data, Type: interpolated, Index: 62291
  Fill level    Drop probability
           0                   0
           1                   1
           2                   3
           4                   6
           5                   7
           6                   9
           8                  12
          10                  15
          12                  15
          14                  16
          15                  16
          16                  16
          18                  17
          20                  17
          22                  18
          24                  18
          25                  18
          26                  19
          28                  19
          30                  20
          32                  20
          34                  21
          35                  21
          36                  21
          38                  22
          40                  22
          42                  23
          44                  23
          45                  23
          46                  24
          48                  24
          49                  24
          51                  25
          52                  25
          54                  26
          55                  26
          56                  26
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
                58                  27
                60                  27
                62                  28
                64                  28
                65                  28
                66                  29
                68                  29
                70                  30
                72                  30
                74                  31
                75                  31
                76                  31
                78                  32
                80                  32
                82                  33
                84                  33
                85                  33
                86                  34
                88                  34
                90                  35
                92                  48
                94                  61
                95                  67
                96                  74
                98                  87
                99                  93
               100                 100
Drop profile: low-data, Type: interpolated, Index: 42361
  Fill level    Drop probability
                0                   0
                1                   2
                2                   5
                4                  10
                5                  12
                6                  15
                8                  20
               10                  25
               12                  26
               14                  27
               15                  28
               16                  28
               18                  30
               20                  31
               22                  32
               24                  33
               25                  34
               26                  35
               28                  36
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

30	37
32	38
34	40
35	40
36	41
38	42
40	43
42	45
44	46
45	46
46	47
48	48
49	49
51	51
52	52
54	54
55	55
56	56
58	58
60	60
62	62
64	64
65	65
66	66
68	68
70	**70**
72	72
74	74
75	75
76	76
78	78
80	80
82	82
84	84
85	85
86	86
88	88
90	90
92	92
94	94
95	95
96	96
98	98
99	99
100	100

```
root@srx-2> show class-of-service scheduler-map cos
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Scheduler map: cos, Index: 33170

  Scheduler: Data, Forwarding class: Data, Index: 35522
    Transmit rate: 50 percent, Rate Limit: none, Buffer size: remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium low      any             1    <default-drop-profile>
      Medium high     any         42361    low-data
      High            any             1    <default-drop-profile>

  Scheduler: Assured-Data, Forwarding class: Assured-Data, Index: 14297
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: VoIP, Forwarding class: VoIP, Index: 46426
    Transmit rate: 10 percent, Rate Limit: none, Buffer size: remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         62291    high-data
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: Network-Control, Forwarding class: Network-Control, Index: 40528
    Transmit rate: 10 percent, Rate Limit: none, Buffer size: remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: Trading, Forwarding class: Trading, Index: 19345
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: remainder,
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Drop profiles:
  Loss priority   Protocol    Index    Name
  Low             any             1    <default-drop-profile>
  Medium low      any             1    <default-drop-profile>
  Medium high     any         42361    low-data
  High            any             1    <default-drop-profile>
```

4) Configure a traffic classifier according to the following table:

Classes	COS	Loss-priority
Data	001	high
Assured-Data	101	medium- high
VoIP	011	low
Network-Control	111	low
Trading	010	medium-high

Table 7.4

```
[edit class-of-service]
+  classifiers {
+      inet-precedence cos {
+          forwarding-class Data {
+              loss-priority high code-points 001;
+          }
+          forwarding-class Assured-Data {
+              loss-priority medium-high code-points 101;
+          }
+          forwarding-class VoIP {
+              loss-priority low code-points 011;
+          }
+          forwarding-class Network-Control {
+              loss-priority low code-points 111;
+          }
+          forwarding-class Trading {
+              loss-priority medium-high code-points 010;
+          }
+      }
+  }
```

5) Configure rewrite rules according to the following table:

Classes	COS	Loss-priority
Data	001	high
Assured-Data	101	medium- high
VoIP	011	low

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

Network-Control	111	low
Trading	010	medium-high

Table 7.5

```
[edit class-of-service]
+  rewrite-rules {
+      inet-precedence cos-rewrite {
+          forwarding-class Data {
+              loss-priority high code-point 001;
+          }
+          forwarding-class Assured-Data {
+              loss-priority medium-high code-point 101;
+          }
+          forwarding-class VoIP {
+              loss-priority low code-point 011;
+          }
+          forwarding-class Network-Control {
+              loss-priority low code-point 111;
+          }
+          forwarding-class Trading {
+              loss-priority medium-high code-point 010;
+          }
+      }
+  }
```

6) Apply the classifier and rewrite rules to all of the interfaces.

```
[edit class-of-service interfaces ge-*]
+    unit * {
+        classifiers {
+            inet-precedence cos;
+        }
+        rewrite-rules {
+            inet-precedence cos-rewrite;
+        }
+    }
```

Verification

```
root@srx-2> show class-of-service interface ge-0/0/1
Physical interface: ge-0/0/1, Index: 134
Queues supported: 8, Queues in use: 5
  Scheduler map: cos, Index: 33170
  Congestion-notification: Disabled
```

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
  Logical interface: ge-0/0/1.32767, Index: 74

  Logical interface: ge-0/0/1.69, Index: 72
    Object                  Name                    Type
Index
    Rewrite                 cos-rewrite             IP
45290
    Classifier              dscp-ipv6-compatibility dscp-ipv6
9
    Classifier              cos                     IP
33575
```

7) SRX2 is not classifying traffic correctly. SSH traffic is ending up in the Data queue but it should be classified as COS 101. The firewall filter and apply-group on the interface level have been deactivated and will need to be reactivated on both srx1 and srx2. Fix this issue.

```
[edit]
root@srx-2# show | compare
[edit interfaces]1
!  active: apply-groups apply-firewall-classify;
[edit firewall family inet]
!      active: filter traffic-class { ... }

{master:0}
root@ex-1> traceroute 21.100.0.13
traceroute to 21.100.0.13 (21.100.0.13), 30 hops max, 40 byte packets
 1  21.51.42.252 (21.51.42.252)  8.881 ms  7.608 ms  3.366 ms
 2  21.1.10.29 (21.1.10.29)  7.554 ms  27.582 ms  7.213 ms
 3  21.1.10.34 (21.1.10.34)  1.904 ms  4.739 ms  1.848 ms
 4  21.100.0.13 (21.100.0.13)  2.259 ms  7.761 ms  9.128 ms

{master:0}
root@ex-1> ssh 21.100.0.13
ssh: connect to host 21.100.0.13 port 22: Connection refused

root@srx-2> show interfaces ge-0/0/4 extensive | find Que
  CoS queues      : 8 supported, 8 maximum usable queues
  Hold-times      : Up 0 ms, Down 0 ms
  Current address: 78:19:f7:18:3d:84, Hardware address: 78:19:f7:18:3d:84
  Last flapped    : 2016-05-07 15:20:24 UTC (03:44:04 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              1198734                 200 bps
   Output bytes  :              2584224                1096 bps
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
   Input  packets:                12552                    0 pps
   Output packets:                17604                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 182393, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped
packets
    0 Data                        7394               7394                    0
    1 Assured-Data                   0                  0                    0
    2 VoIP                           0                  0                    0
    3 Network-Cont               10210              10210                    0

[edit]
root@srx-2# show | compare
[edit groups accept-telnet firewall family inet filter <*> term <*> from]
-         protocol tcp;
-         port 22;
+         protocol tcp;
+         port 23;
```

Verification

```
{master:0}
root@ex-1> ssh 21.100.0.13
ssh: connect to host 21.100.0.13 port 22: Connection refused

root@srx-2> show interfaces ge-0/0/4 extensive | find Que
  CoS queues      : 8 supported, 8 maximum usable queues
  Hold-times      : Up 0 ms, Down 0 ms
  Current address: 78:19:f7:18:3d:84, Hardware address: 78:19:f7:18:3d:84
  Last flapped    : 2016-05-07 15:20:24 UTC (03:44:17 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              1200677                 4368 bps
   Output bytes  :              2587127                 6160 bps
   Input  packets:                12566                    1 pps
   Output packets:                17622                    2 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 182600, L2 mismatch timeouts: 0,
```


http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets

```
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped
packets
    0 Data                       7397               7397                    0
    1 Assured-Data                  1                  1                    0
    2 VoIP                          0                  0                    0
    3 Network-Cont              10224              10224                    0
    4 Trading                       0                  0                    0
```

END

http://www.inetzero.com - Copyright 2017 iNET ZERO, The Netherlands. All rights reserved

https://t.me/learningnets