



Product Manual

TNSR v20.02

© Copyright 2023 Rubicon Communications LLC

Mar 15, 2023

CONTENTS

1	Introduction	2
1.1	TNSR Secure Networking	2
1.2	TNSR Trial	2
1.3	TNSR Architecture	2
1.4	Technology Stack	4
1.5	Basic Assumptions	4
2	Supported Platforms	5
3	Installation	6
3.1	Hardware Installation	10
4	Default Behavior	12
4.1	Default Accounts and Passwords	12
4.2	Default TNSR Permissions	13
4.3	Default Allowed Traffic	13
5	Zero-to-Ping	15
5.1	First Login	15
5.2	Interface Configuration	16
5.3	TNSR Interfaces	17
5.4	NAT	18
5.5	DHCP Server	18
5.6	DNS Server	19
5.7	Ping	19
5.8	Save the TNSR Configuration	20
5.9	Next Steps	21
6	Command Line Basics	22
6.1	Working in the TNSR CLI	22
6.2	Finding Help	23
6.3	Starting TNSR	24
6.4	Entering the TNSR CLI	25
6.5	Configuration Database	26
6.6	Configuration Mode	28
6.7	Configuration Backups	31
6.8	Viewing Status Information	32
6.9	Service Control	33
6.10	Diagnostic Utilities	33
6.11	Basic System Information	35

7	Basic Configuration	37
7.1	Setup Interfaces	37
7.2	Disable Host OS NICs for TNSR	38
7.3	Setup NICs in Dataplane	39
7.4	Setup QAT Compatible Hardware	42
7.5	Remove TNSR NIC for Host Use	46
8	Updates and Packages	50
8.1	Generate a Key Pair	50
8.2	Generate a Certificate Signing Request	51
8.3	Submit the Certificate Signing Request	52
8.4	Retrieve the signed certificate	54
8.5	Install the certificate	54
8.6	Package Management	55
8.7	Package Information Commands	55
8.8	Package Installation	56
8.9	Updating TNSR	56
9	Interfaces	60
9.1	Locate Interfaces	60
9.2	Configure Interfaces	61
9.3	Types of Interfaces	64
10	Static Routing	86
10.1	Route Tables	86
10.2	Neighbors	86
10.3	Viewing Routes	87
10.4	Managing Routes	88
10.5	Default Route	90
11	Dynamic Routing	91
11.1	Dynamic Routing Manager	91
11.2	Border Gateway Protocol	100
11.3	Open Shortest Path First v2 (OSPF)	116
11.4	Open Shortest Path First v3 (OSPF6)	129
11.5	Routing Information Protocol (RIP)	139
11.6	Dynamic Routing Protocol Lists	145
12	Virtual Router Redundancy Protocol	147
12.1	VRRP Compatibility	148
12.2	VRRP Example	148
12.3	VRRP Configuration	152
12.4	VRRP Status	154
13	Access Lists	155
13.1	Standard ACLs	155
13.2	MACIP ACLs	157
13.3	Viewing ACL and MACIP Information	158
13.4	ACL and NAT Interaction	159
13.5	Host ACLs	159
14	IPsec	162
14.1	Required Information	162
14.2	IPsec Example	163
14.3	IPsec Configuration	165

14.4	IPsec Status Information	174
14.5	IPsec Cryptographic Acceleration	176
15	Network Address Translation	178
15.1	Dataplane NAT Modes	178
15.2	NAT Options	179
15.3	NAT Pool Addresses	180
15.4	Outbound NAT	181
15.5	Static NAT	182
15.6	Dual-Stack Lite	183
15.7	Deterministic NAT	185
15.8	NAT Status	185
15.9	NAT Examples	187
16	MAP (Mapping of Address and Port)	193
16.1	MAP Configuration	193
16.2	MAP Parameters	196
16.3	MAP Example	197
16.4	MAP Types	198
17	Dynamic Host Configuration Protocol	200
17.1	DHCP Configuration	200
17.2	DHCP Logging	205
17.3	DHCP Service Control and Status	207
17.4	DHCP Service Example	208
18	DNS Resolver	210
18.1	DNS Resolver Configuration	210
18.2	DNS Resolver Service Control and Status	216
18.3	DNS Resolver Examples	217
19	Network Time Protocol	219
19.1	NTP Configuration	219
19.2	NTP Service Control and Status	222
19.3	NTP Service Example	223
19.4	NTP Best Practices	224
20	Link Layer Discovery Protocol	225
20.1	Configuring the LLDP Service	225
20.2	LLDP Status	226
21	Public Key Infrastructure	227
21.1	Key Management	227
21.2	Certificate Signing Request Management	229
21.3	Certificate Management	231
21.4	Certificate Authority Management	232
22	Bidirectional Forwarding Detection	236
22.1	BFD Sessions	236
22.2	BFD Session Authentication	238
22.3	BFD Example	239
23	User Management	243
23.1	User Configuration	243
23.2	Authentication Methods	243

24	NETCONF Access Control Model (NACM)	245
24.1	NACM Example	245
24.2	View NACM Configuration	246
24.3	Enable or Disable NACM	247
24.4	NACM Default Policy Actions	248
24.5	NACM Username Mapping	248
24.6	NACM Groups	248
24.7	NACM Rule Lists	249
24.8	NACM Rules	249
24.9	NACM Rule Processing Order	251
24.10	Regaining Access if Locked Out by NACM	251
24.11	NACM Defaults	252
25	HTTP Server	254
25.1	HTTP Server Configuration	254
25.2	HTTPS Encryption	254
25.3	Authentication	255
25.4	RESTCONF Server	256
26	Monitoring	257
26.1	Monitoring Interfaces	257
26.2	Simple Network Management Protocol	260
27	TNSR Configuration Example Recipes	265
27.1	RESTCONF Service Setup with Certificate-Based Authentication and NACM	265
27.2	TNSR IPsec Hub for pfSense	274
27.3	Edge Router Speaking eBGP with Static Redistribution for IPv4 And IPv6	294
27.4	Service Provider Route Reflectors and Client for iBGP IPv4	304
27.5	LAN + WAN with NAT (Basic SOHO Router Including DHCP and DNS Resolver)	322
27.6	Using Access Control Lists (ACLs)	325
27.7	Inter-VLAN Routing	328
27.8	GRE ERSPAN Example Use Case	331
27.9	OSPF Router with Multiple Areas and Summarization	333
27.10	TNSR Remote Office With Existing IPsec Hub	337
28	Advanced Configuration	354
28.1	Dataplane Configuration	354
28.2	Host Memory Management Configuration	359
28.3	IP Reassembly	360
29	Troubleshooting	363
29.1	Ping and traceroute do not function without host OS default route	363
29.2	Unrecognized routes in a routing table	363
29.3	Services do not receive traffic on an interface with NAT enabled	363
29.4	NAT session limits / “Create NAT session failed” error	363
29.5	ACL rules do not match NAT traffic as expected	363
29.6	Some Traffic to the host OS management interface is dropped	364
29.7	Locked out by NACM Rules	364
29.8	How to gain access to the root account	364
29.9	Console Messages Obscure Prompts	364
29.10	OSPF Neighbors Stuck in ExStart State	364
29.11	Large packets fail to pass over IPsec	364
29.12	IPsec packets are dropped or fail to pass with QAT enabled	365
29.13	Console DMA / PTE Read access Error Messages	365
29.14	Diagnosing Service Issues	365

29.15 Debugging TNSR	365
29.16 Diagnostic Information for Support	367
30 Commands	368
30.1 Mode List	368
30.2 Master Mode Commands	370
30.3 Config Mode Commands	370
30.4 Show Commands in Both Master and Config Modes	373
30.5 Access Control List Modes	375
30.6 MACIP ACL Mode	376
30.7 GRE Mode	377
30.8 HTTP mode	377
30.9 Interface Mode	378
30.10 Loopback Mode	379
30.11 Bridge Mode	379
30.12 NAT Commands in Configure Mode	380
30.13 DS-Lite Commands in Configure Mode	380
30.14 Tap Mode	380
30.15 BFD Key Mode	381
30.16 BFD Mode	381
30.17 Host Interface Mode	382
30.18 IPsec Tunnel Mode	383
30.19 IKE mode	383
30.20 IKE Peer Authentication Mode	384
30.21 IKE Peer Authentication Round Mode	384
30.22 IKE Child SA Mode	384
30.23 IKE Child SA Proposal Mode	385
30.24 IKE Peer Identity Mode	385
30.25 IKE Proposal Mode	386
30.26 Map Mode	386
30.27 Map Parameters Mode	387
30.28 memif Mode	387
30.29 Dynamic Routing Access List Mode	388
30.30 Dynamic Routing Prefix List Mode	388
30.31 Dynamic Routing Route Map Rule Mode	389
30.32 Dynamic Routing BGP Mode	390
30.33 Dynamic Routing BGP Server Mode	391
30.34 Dynamic Routing BGP Neighbor Mode	392
30.35 Dynamic Routing BGP Address Family Mode	392
30.36 Dynamic Routing BGP Address Family Neighbor Mode	394
30.37 Dynamic Routing BGP Community List Mode	395
30.38 Dynamic Routing BGP AS Path Mode	395
30.39 Dynamic Routing OSPF Mode	396
30.40 Dynamic Routing OSPF Server Mode	396
30.41 Dynamic Routing OSPF Interface Mode	397
30.42 Dynamic Routing OSPF Area Mode	398
30.43 Dynamic Routing OSPF6 Mode	398
30.44 Dynamic Routing OSPF6 Server Mode	399
30.45 Dynamic Routing OSPF6 Interface Mode	399
30.46 Dynamic Routing OSPF6 Area Mode	400
30.47 Dynamic Routing RIP Mode	400
30.48 Dynamic Routing RIP Server Mode	401
30.49 Dynamic Routing RIP Interface Mode	401
30.50 Dynamic Routing RIP Key Chain Mode	402

30.51	Dynamic Routing Manager Mode	402
30.52	IPv4 Route Table Mode	403
30.53	IPv6 Route Table Mode	403
30.54	IPv4 or IPv6 Next Hop Mode	404
30.55	SPAN Mode	404
30.56	VXLAN Mode	405
30.57	User Authentication Configuration Mode	405
30.58	NTP Configuration Mode	406
30.59	NTP Restrict Mode	406
30.60	NTP Upstream Server Mode	407
30.61	NACM Group Mode	407
30.62	NACM Rule-list Mode	408
30.63	NACM Rule Mode	408
30.64	DHCP IPv4 Server Config Mode	409
30.65	DHCP4 Subnet4 Mode	409
30.66	DHCP4 Subnet4 Pool Mode	410
30.67	DHCP4 Subnet4 Reservation Mode	410
30.68	Kea DHCP4, Subnet4, Pool, or Reservation Option Mode	411
30.69	Kea DHCP4 Option Definition Mode	411
30.70	DHCP4 Log Mode	412
30.71	DHCP4 Log Output Mode	412
30.72	Unbound Server Mode	412
30.73	Unbound Forward-Zone Mode	413
30.74	Subif Mode	414
30.75	Bond Mode	414
30.76	Host ACL Mode	415
30.77	Host ACL Rule Mode	415
30.78	VRRP Mode	416
31	API Endpoints	417
31.1	YANG Data Models	417
31.2	RESTCONF API	417
32	Netgate TNSR Releases	418
32.1	TNSR 20.02.2 Release Notes	418
32.2	TNSR 20.02.1 Release Notes	419
32.3	TNSR 20.02 Release Notes	419
32.4	TNSR 19.12 Release Notes	429
32.5	TNSR 19.08 Release Notes	438
32.6	TNSR 19.05 Release Notes	445
32.7	TNSR 19.02.1 Release Notes	452
32.8	TNSR 19.02 Release Notes	455
32.9	TNSR 18.11 Release Notes	460
32.10	TNSR 18.08 Release Notes	464
32.11	TNSR 18.05 Release Notes	468
32.12	TNSR 0.1.0 Release Notes	470
33	Licensing	472
33.1	Apache 2.0 License	472
33.2	BSD 2-Clause License	476
33.3	BSD 3-Clause License	477
33.4	CentOS EULA License	477
33.5	GPLv2.0 License	478
33.6	LGPLv2.1 License	484

33.7	MIT License	493
33.8	Mozilla Public License 2.0 (MPL-2.0)	493
33.9	Net SNMP License	500
33.10	NTP License	505
33.11	Joint OpenSSL and SSLeay License	507

This documentation has all the details needed to fully configure your TNSR platform, from the basics of TNSR all the way to the complexities of implementing different applications. For quotes, updates, and more information about TNSR, please visit tnsr.com or [contact TNSR sales](#).

INTRODUCTION

TNSR is an open-source based packet processing platform that delivers superior secure networking solution performance, manageability, and services flexibility. TNSR can scale packet processing from 1 to 10 to 100 Gbps, even 1 Tbps and beyond on commercial-off-the-shelf (COTS) hardware - enabling routing, firewall, VPN and other secure networking applications to be delivered for a fraction of the cost of legacy brands. TNSR features a RESTCONF API - enabling multiple instances to be orchestration managed - as well as a CLI for single instance management.

1.1 TNSR Secure Networking

TNSR is a full-featured software solution designed to provide secure networking from 1 Gbps to 400 Gbps. With graduated pricing based on performance increments, TNSR is a viable option for users with moderate bandwidth needs to the demanding requirements of enterprise and service providers.

Each licensed instance comes bundled with TNSR Technical Assistance from our 24/7 world-wide team of support engineers. Find out more about the [included support](#) available with TNSR.

[Contact us](#) to begin a conversation about how TNSR can help meet your needs.

1.2 TNSR Trial


A 120-day trial version is also available. You can visit the Trials page of tnsr.com to find out full details on how the trial works.

1.3 TNSR Architecture


TNSR runs on a Linux host operating system. Initial configuration of TNSR includes installing associated services and configuring network interfaces. It is important to note that network interfaces can be managed by the host OS or by TNSR, but not by both. In other words, once a network interface is assigned to TNSR, it is no longer available - or even visible - to the host OS.

A little background. TNSR is the result of Netgate development, using many open source technologies to create a product that can be supported and easily implemented in production environments.

Without TNSR, Linux systems use drivers to plumb the connections from hardware interfaces (NICs) to the OS kernel. The Linux kernel then handles all I/O between these NICs. The kernel also handles all other I/O tasks, as well as memory and process management.



In high I/O situations, the kernel can be tasked with servicing millions of requests per second. TNSR uses two open source technologies to simplify this problem and service terabits of data in user space. Data Plane Development Kit (DPDK) bypasses the kernel, delivering network traffic directly to user space, and Vector Packet Processing (VPP) accelerates traffic processing.



In practical terms, this means that once a NIC is assigned to TNSR, **that NIC is attached to** a fast data plane, but it is no longer available to the host OS. All management - including configuration, troubleshooting and update - of TNSR is performed either at the console or via RESTCONF. In cloud or virtual environments, console access may be available, but the recommended configuration is still to dedicate a host OS interface for RESTCONF API access.

The recommended configuration of a TNSR system includes one host NIC for the host OS and all other NICs assigned to TNSR.

This is important and bears repeating:

- The host OS cannot access NICs assigned to TNSR
- In order to manage TNSR, you must be able to connect to the console

1.4 Technology Stack

TNSR is designed and built from the ground up, using open source software projects including:

- [Vector Packet Processing](#) (VPP)
- [Data Plane Developer Kit](#) (DPDK)
- [YANG](#) for data modeling
- [Clixon](#) for system management
 - Command Line Interface (CLI)
 - [RESTCONF](#) for REST API configuration served by [nginx](#)
- [FRR](#) for routing protocols
- [strongSwan](#) for IPsec key management
- [Kea](#) for DHCP services
- [net-snmp](#) for SNMP
- [ntp.org](#) daemon for NTP
- [Unbound](#) for DNS
- [CentOS](#) as the base operating system

See also:

What is Vector Packet Processing? Vector processing handles more than one packet at a time, as opposed to scalar processing which handles packets individually. The vector approach fixes problems that scalar processing has with cache efficiency, read latency, and issues related to stack depth/misses.

For technical details on how VPP accomplishes this feat, see the [VPP Wiki](#).

1.5 Basic Assumptions

This documentation assumes the reader has moderate to advanced networking knowledge and some familiarity with the CentOS Linux distribution.

SUPPORTED PLATFORMS

There are three tiers of [Supported Platforms](#) for TNSR that serve as a guide to deploying TNSR successfully. The information and specifications in each tier listed below meet different requirements for different deployment types:

- Tier One - Tested by Netgate
- Tier Two - DPDK and VPP Compatible
- Tier Three - Community Reported

Documented Platforms

INSTALLATION

Use the following instructions to install TNSR 20.02.1 from an .ISO image. Ensure that the target hardware meets the minimum specifications for a [TNSR Supported Platform](#).

1. Obtain the TNSR .iso image file from Netgate®.
2. Write the .iso image to bootable media (DVD or USB drive) for hardware installations, or copy the .iso image to a location readable by the hypervisor for virtual machine installations.
3. Connect to the system or VM console.

Note: The installer supports both VGA and serial console output, with VGA as the default.

4. Insert or attach the boot media to the target system.
5. Boot the system using the TNSR image.

Note: If the optical drive or removable media is not set as the primary boot device for the hardware, then use the system boot menu to manually select the boot device.

6. After a few seconds, the installer displays a TNSR 20.02.1 screen.
7. **Press any key**, such as `space`, to stop the 60-second timer. The menu contains, at minimum, the following two choices:
 - **Install TNSR (serial-console) <version>**: Select this option for hardware that uses serial port 0 for its console.
 - **Install TNSR <version>**: Select this option for installation via VGA console
8. Highlight the correct option for this hardware and press `Enter` to begin the installation of **TNSR**. It may take a few seconds for the installer to display output to the console.

Note: If the console does not display a visual indication of which item is selected, reboot the device and use the BIOS boot selection menu to choose UEFI as the boot method. For example, on the SG-5100, press `Esc` during POST to access this menu, and of the two entries in the menu for the USB drive, choose the line that starts with `UEFI :`.

9. Once the installer launches, it displays a menu labeled **Installation** with nine choices. All options marked with `[!]` must be configured to resolve all installation requirements.

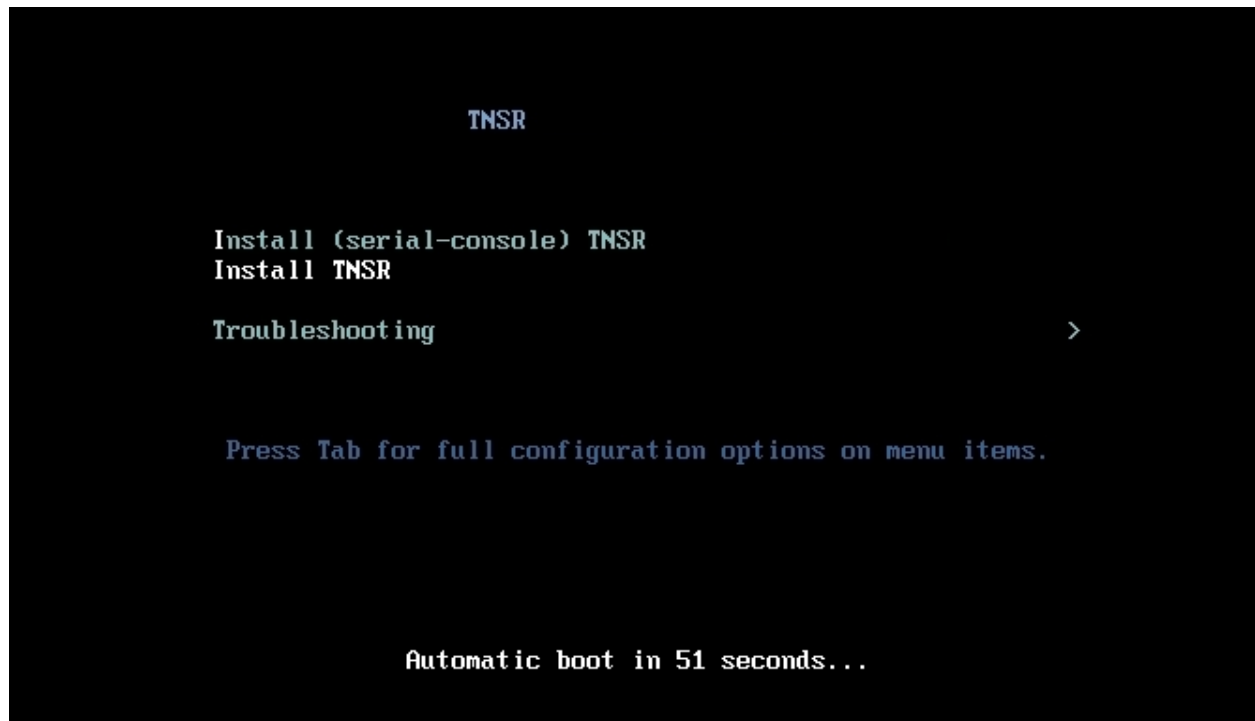


Fig. 1: TNSR 20.02.1 Installation Menu

Note: Some items marked with a ! will resolve on their own a few moments after the installer launches, such as options **3** and **4**. Wait a few moments and enter `r` to refresh the screen.

At a minimum, configure **2) Time Settings**, **5) Installation Destination**, and an administrator account with **9) User creation** to allow the installer to proceed. These are covered next.

10. Configure the time zone

- Enter `2` to start the time zone configuration process.
- Enter `1` to enter the time zone selection screen.
- Continue through the available options until the correct zone is located.
For example, Enter `3` for America, then `36` for Chicago.
- Enter the number corresponding to the region and zone, or type out the zone name.

After selecting a zone, the installer will return to the main menu.

11. Configure the installation destination.

- Enter `5` to start the installation destination configuration process.
- Select the correct target disk on the next screen.
The installer will select the disk automatically when only one is present.
- Enter `c` to continue.
- Choose how to partition the disk.
The default **Use All Space** is the best practice.

```
Starting installer, one moment...
anaconda 21.48.22.147-1 for TNSR started.
* installation log files are stored in /tmp during the installation
* shell is available on TTY2
* when reporting a bug add logs from /tmp as separate text/plain attachments
18:03:25 Not asking for UNC because we don't have a network
=====
Installation

1) [x] Language settings                2) [!] Time settings
    (English (United States))           (Timezone is not set.)
3) [x] Installation source              4) [x] Software selection
    (Local media)                       (TNSR Install)
5) [!] Installation Destination         6) [x] Kdump
    (No disks selected)                 (Kdump is enabled)
7) [ ] Network configuration           8) [x] Root password
    (Not connected)                     (Root account is disabled.)
9) [ ] User creation
    (No user will be created)
Please make your choice from above ['q' to quit ; 'b' to begin installation ;
'r' to refresh]:
```

```
[anaconda] 1:main* 2:shell 3:log 4:storage-log 5:program-log Switch tab: Alt+Tab ; Help: F1
```

Fig. 2: TNSR 20.02.1 Setup Menu

- Enter `c` to continue.
- Choose the partition scheme.
The default **LVM** is the best practice.
- Enter `c` to finish and return to the main menu.

12. Add an administrator account.

Note: Security best practices dictate that it is best not to enable interactive logon for the `root` account. As such, the `root` account will be locked out after the installation. Use this process to add at least one alternate administrator account.

- Enter 9 to start the user configuration process.
- Enter 1 to create a new user.
- Enter 3 to enter the username.
- Enter 4 to configure the account to use a password.
- Enter 5 to set and confirm the password for the user.
- Enter 6 to mark the user as an Administrator.
- Enter `c` to finish and return to the main menu.

13. Optionally configure a Host OS interface.

This will enable a network interface in the host OS for use as a management interface. This interface can then be used to access the system for troubleshooting or maintenance.

<p>Warning: Though this is technically optional, using a management interface is the best practice.</p>
--

- Enter 7 to start the interface configuration.
- Choose one of the listed network interfaces.
- Configure interface parameters on this screen as needed, such as a static IP address.

Note: The default behavior is to use DHCP to obtain the interface address. If this is the desired behavior, then leave the address options as-is.

- Enter 7 to enable **Connect automatically after reboot**.
- Enter 8 to enable **Apply configuration in installer**.
- Enter `c` to complete the interface configuration and continue back to the interface list.
- Enter `c` again to exit the network configuration.

14. Once all options with [!] have been resolved, enter `b` from the main menu to begin the installation. Messages are displayed indicating the progress of the installation. When the installer finishes its tasks, it displays message that says **Installation complete. Press return to quit**. At that point, press `Enter` and the system will reboot.

Note: The installer may spend several minutes displaying the message **Performing post-installation setup tasks**, but it will eventually continue.

15. When the system is restarting, remove the DVD or USB drive while the system reboots. CentOS 7 will start up automatically from the disk to which it was installed. If the installation media remains inserted, the system may boot into the installer again.

Note: The boot options in the system BIOS may need changed if it does not boot automatically into CentOS 7.

16. After the system finishes rebooting, log in with the user and password chosen during the installation.

Note: Once the system reboots, network interfaces not configured in the installer will be disabled in CentOS. Depending on the hardware, these interfaces may automatically be enabled in TNSR. If TNSR does not see any interfaces, they will need to be manually configured in TNSR. See [Setup NICs in Dataplane](#) for details.

Tip: One network interface should be enabled in the host OS as a **management** interface to allow access to the system for troubleshooting or maintenance. This can be configured in the installer, as mentioned above, or afterward.

Warning: Once the Host OS is capable of reaching the Internet, check for updates ([Updating TNSR](#)) before proceeding. This ensures the security and integrity of the router before TNSR interfaces are exposed to the Internet.

Before continuing on, check the next section for hardware-specific installation guidance. Certain hardware may require additional configuration before it is usable by TNSR.

3.1 Hardware Installation

This section includes information about known actions required for certain hardware components to function with TNSR. If this TNSR installation does not include any of the hardware listed here, skip ahead to [Default Behavior](#).

3.1.1 Mellanox ConnectX-5 Firmware Requirements

Mellanox ConnectX-5 network interface cards (mlx5) in the MT27800 family are currently shipping with firmware revision 16.26.1040 which is not compatible with TNSR.

The incompatible firmware can be identified by errors in the log (e.g. `sudo vppctl show errors`), such as:

```
net_mlx5: probe of PCI device [PCI ID] aborted after encountering an error: Operation_
↪not supported
```

Or:

```
Interface WAN error -12: Unknown error -12
net_mlx5: Failed to query QP using DevX
net_mlx5: Fail to query port 0 Tx queue 0 QP TIS transport domain
net_mlx5: port 0 Tx queue allocation failed: Cannot allocate memory
Device with port_id=0 already stopped
```

Firmware version 16.24.1000 is compatible with TNSR, and can be manually downgraded using the following procedure.

First, identify the first PCI ID of the card in question. This can be found by looking in the boot logs, the output of a utility such as `lspci`, or similar methods. The ID will take the form of `xx:yy.z`, for example `65:00.0`. The ID will be used in the following set of commands.

Next, download and decompress the appropriate firmware:

```
$ curl -LO http://www.mellanox.com/downloads/firmware/fw-ConnectX5-rel-16_24_1000-
↪MCX516A-CCA_Ax-UEFI-14.17.11-FlexBoot-3.5.603.bin.zip
$ unzip fw-ConnectX5-rel-16_24_1000-MCX516A-CCA_Ax-UEFI-14.17.11-FlexBoot-3.5.603.bin.
↪zip
```

Now install the Mellanox firmware tool and perform the firmware downgrade:

```
$ sudo yum install mstflint
$ sudo mstconfig q > mst.log
$ sudo mstfwmanager -d <PCI ID> -u -f -i fw-ConnectX5-rel-16_24_1000-MCX516A-CCA_Ax-
↪UEFI-14.17.11-FlexBoot-3.5.603.bin
```

Replace `<PCI ID>` in the last command with the first PCI ID of the card.

With the appropriate firmware loaded on the card, it will no longer produce errors and it will be usable by TNSR.

Warning: Reboot after performing the firmware downgrade to ensure the card is fully reinitialized with the appropriate firmware.

DEFAULT BEHAVIOR

After the installation completes and TNSR boots for the first time, TNSR has an empty default configuration. This means that TNSR has no pre-configured interfaces, addresses, routing behavior, and so on.

The host OS defaults are set during installation, and depend on the base OS, CentOS 7.7. For example, host management interfaces may have been configured by the installer.

4.1 Default Accounts and Passwords

By default, the TNSR installation includes host OS accounts for `root` with interactive login disabled, and a `tnsr` account.

For ISO installations, the best practice is to create at least one additional initial administrator account during the installation process. That user is custom created by the person performing the installation, and thus is not a common default that can be listed here.

Warning: When installing TNSR from an ISO image, the installer allows the `root` account to be unlocked and assigned a password. The best practice, however, is to leave the `root` account locked and create at least one additional administrator account using the installer. These additional accounts may use `sudo` to elevate privileges. Any users added to the `wheel` group later may also use `sudo` to execute commands as `root`.

The default behavior of the `tnsr` account varies by platform:

ISO/Bare Metal The `tnsr` user is available with a default password of `tnsr-default`.

Appliances Shipped with TNSR Pre-installed The `tnsr` user is available with a default password of `tnsr-default`.

AWS The `tnsr` account is present but restricted to key-based authentication via SSH, using a key selected when launching the TNSR instance.

Azure The `tnsr` account is present but restricted to key-based authentication via SSH, using a key selected when launching the TNSR instance.

The password for the `tnsr` account can be reset by any other account with access to the shell and `sudo`. For example, the command `shell sudo passwd tnsr` run at a TNSR prompt will set and confirm a new password for the `tnsr` user. The same action may also be performed for the `root` account (`shell sudo passwd root`). As mentioned in the previous warning, it is best to leave interactive logins for `root` disabled.

Warning: Change default passwords, even randomized default passwords or passwords pre-configured when launching a cloud-based instance, after the first login. Do not leave default passwords active!

Note: User authentication is performed by the host OS. Though users may be created inside TNSR (*User Management*), these users are propagated to the host. To control what users may access, see *NETCONF Access Control Model (NACM)*.

4.2 Default TNSR Permissions

By default, there is no TNSR configuration present. As such, there are no pre-configured access permissions for users to restrict access to TNSR. Thus, any operating system user on the TNSR host will be able to reach the TNSR CLI and make changes.

To restrict which accounts have access to TNSR, see *NETCONF Access Control Model (NACM)*.

4.3 Default Allowed Traffic

For the default behavior of allowed traffic to and from TNSR, there are two separate areas to consider:

- Traffic flowing through TNSR
- Traffic for the host OS management interface

4.3.1 TNSR

By default, there is no TNSR configuration present. As such, there are no default access lists (ACLs) and once TNSR is able to route traffic, all packets flow freely. See *Access Lists* for information on configuring access lists.

4.3.2 Host OS

The TNSR installation configures a default set of Netfilter rules for the host OS management interface. The following traffic is allowed to pass into and out of the host operating system interfaces:

- ICMP / ICMP6
- SSH (TCP/22)
- HTTP (TCP/80)
- HTTPS (TCP/443)
- BGP (TCP/179)
- OSPF (Protocol 89)
- RIP (UDP/520)
- ISAKMP (UDP/500)
- NTP (UDP/123, TCP/123)
- DNS (UDP/53, TCP/53)
- SNMP (UDP/161)
- DHCP Server (UDP/67)
- UDP Traceroute (UDP ports 33434-33524 with TTL=1)

To manage host ACLs which can override this behavior, see [Host ACLs](#).

ZERO-TO-PING

This document is a crash course in getting TNSR up and running quickly after installation. The topics included here are covered in more detail throughout the remainder of the documentation.

Each section contains a list of additional related resources with more detail in a **See Also** box. Follow these links for more information on each topic.

5.1 First Login

When TNSR boots, it will present a login prompt on the console (video and serial). Login at this prompt using either the default `tnsr` account or an administrator account created during the installation process.

Note: For installations from ISO and for hardware shipped with TNSR preinstalled, the default password for the `tnsr` user is `tnsr-default`. For cloud-based installs such as AWS and Azure, by default the `tnsr` account can only login with key-based ssh authentication. See [Default Accounts and Passwords](#) for more information.

The `tnsr` user automatically enters the TNSR CLI when used to login interactively. Manually created administrative users do not have this behavior, and using them to login interactively will result in a login shell.

Alternately, if the host OS management interface was configured in the installer, login using an SSH client connecting to that interface.

See also:

- [Installation](#)
- [Default Accounts and Passwords](#)

5.1.1 Changing the Password

The password for administrator accounts was set during the installation process, but the default `tnsr` account should have its password reset before making other changes.

Login to the `tnsr` account with the default `tnsr-default` password and change it using the `shell passwd` command from the TNSR CLI:

```
tnsr# shell passwd
Changing password for user tnsr.
Changing password for tnsr.
(current) UNIX password:
New password:
```

(continues on next page)

(continued from previous page)

```
Retype new password:
passwd: all authentication tokens updated successfully.
tnsr#
```

Alternately, login in as an administrator and change the password for the default `tnsr` account using `sudo`:

```
tnsr# shell sudo passwd tnsr
Changing password for user tnsr.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
tnsr#
```

Note: These examples use the TNSR prompt and `shell` command. The same commands may be used without the `shell` prefix from a non-TNSR shell prompt.

Warning: Use a strong password for this account as it will be able to make changes to the TNSR configuration, unless restricted by a custom *NACM configuration*.

See also:

- *Installation*
- *Default Accounts and Passwords*
- *NETCONF Access Control Model (NACM)*

5.2 Interface Configuration

There are two types of interfaces on a TNSR system: Host OS interfaces for managing the device and dataplane interfaces which are available for use by TNSR.

5.2.1 Host OS Management Interface

A host management interface may be configured manually in the installer or later in TNSR or in CentOS. See *Installation* for the full procedure to configure a host OS management interface during installation, and *Host Interfaces* for information on configuring host OS interfaces from within TNSR.

At a minimum, the host OS interface must have an IP address, subnet mask, and a default gateway configured. The default gateway is necessary so that the host OS may retrieve updates as that traffic does not flow through TNSR, but over the management interface. Additionally, other host traffic may flow through the management interface, such as the `ping` command from within the TNSR CLI.

If an interface was not configured for management in the installer, it will need to be manually changed back to host OS control and then configured for network access. See *Remove TNSR NIC for Host Use* for instructions on how to return an interface from TNSR back to host OS control so it can be used for management. This procedure will require rebooting the TNSR device.

Consult CentOS 7.7 documentation for the specifics of network configuration for other environments.

Warning: Once the Host OS is capable of reaching the Internet, check for updates (*Updating TNSR*) before proceeding. This ensures the security and integrity of the router before TNSR interfaces are exposed to the Internet.

See also:

- [Installation](#)
- [Disable Host OS NICs for TNSR](#)
- [Host Interfaces](#)
- [Remove TNSR NIC for Host Use](#)

5.2.2 Dataplane Interfaces

Interfaces not configured for host OS management control in the installer will be setup in such a way that they are available for use by the dataplane and thus TNSR.

Enter the TNSR CLI (*Entering the TNSR CLI*) and configure the network interfaces:

```
tnsr# configure
tnsr(config)# dataplane dpdk dev ?
0000:00:14.0      Ethernet controller: Intel Corporation Ethernet
                  Connection I354 (rev 03)
0000:00:14.1      Ethernet controller: Intel Corporation Ethernet
                  Connection I354 (rev 03)
0000:00:14.2      Ethernet controller: Intel Corporation Ethernet
                  Connection I354 (rev 03)
0000:00:14.3      Ethernet controller: Intel Corporation Ethernet
                  Connection I354 (rev 03)
0000:03:00.0      Ethernet controller: Intel Corporation I211 Gigabit
                  Network Connection (rev 03)
0000:04:00.0      Ethernet controller: Intel Corporation I211 Gigabit
                  Network Connection (rev 03) ( Active Interface enp4s0 )
tnsr(config)# dataplane dpdk dev 0000:00:14.1 network
tnsr(config)# dataplane dpdk dev 0000:00:14.2 network
tnsr(config)# service dataplane restart
tnsr(config)# exit
```

See also:

- [Installation](#)
- [Setup NICs in Dataplane](#)

5.3 TNSR Interfaces

Next, the interfaces inside TNSR must be configured with addresses and routing.

5.3.1 WAN DHCP Client

In this example, WAN will be set as a DHCP client and configured as the `outside` NAT interface:

```
tnsr# configure terminal
tnsr(config)# interface GigabitEthernet0/14/1
tnsr(config-interface)# description Internet
tnsr(config-interface)# dhcp client ipv4
tnsr(config-interface)# enable
tnsr(config-interface)# ip nat outside
tnsr(config-interface)# exit
tnsr(config)# exit
```

See also:

- [DHCP Client Example](#)
- [Configure Interfaces](#)

5.3.2 LAN Interface

Next, configure an address for the internal network and set it as the inside NAT interface:

```
tnsr(config)# interface GigabitEthernet0/14/2
tnsr(config-interface)# ip address 172.16.1.1/24
tnsr(config-interface)# description Local
tnsr(config-interface)# ip nat inside
tnsr(config-interface)# enable
tnsr(config-interface)# exit
```

See also:

- [Configure Interfaces](#)

5.4 NAT

Configure TNSR to use the WAN interface address for NAT, and enable NAT forwarding:

```
tnsr(config)# nat pool interface GigabitEthernet0/14/1
tnsr(config)# nat global-options nat44 forwarding true
```

See also:

- [Network Address Translation](#)
- [NAT Pool Addresses](#)
- [NAT Forwarding](#)

5.5 DHCP Server

Setup a basic DHCP server on the LAN side to hand out addresses, also instruct clients to use TNSR as their gateway and DNS server.

```
tnsr(config)# dhcp4 server
tnsr(config-kea-dhcp4)# description LAN DHCP Server
tnsr(config-kea-dhcp4)# interface listen GigabitEthernet0/14/2
tnsr(config-kea-dhcp4)# subnet 172.16.1.0/24
```

(continues on next page)

(continued from previous page)

```
tnsr(config-kea-subnet4)# pool 172.16.1.100-172.16.1.245
tnsr(config-kea-subnet4-pool)# exit
tnsr(config-kea-subnet4)# interface GigabitEthernet0/14/2
tnsr(config-kea-subnet4)# option domain-name-servers
tnsr(config-kea-subnet4-opt)# data 172.16.1.1
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-subnet4)# option routers
tnsr(config-kea-subnet4-opt)# data 172.16.1.1
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-subnet4)# exit
tnsr(config-kea-dhcp4)# exit
tnsr(config)# dhcp4 enable
```

See also:

- *Dynamic Host Configuration Protocol*

5.6 DNS Server

Configure TNSR to act as a DNS server for local clients, using upstream forwarding DNS servers of 8.8.8.8 and 8.8.4.4:

```
tnsr# configure
tnsr(config)# unbound server
tnsr(config-unbound)# interface 127.0.0.1
tnsr(config-unbound)# interface 172.16.1.1
tnsr(config-unbound)# access-control 172.16.1.0/24 allow
tnsr(config-unbound)# forward-zone .
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.8.8
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.4.4
tnsr(config-unbound-fwd-zone)# exit
tnsr(config-unbound)# exit
tnsr(config)# unbound enable
```

See also:

- *DNS Resolver*

5.7 Ping

5.7.1 From the Host

The TNSR CLI includes a ping utility which will send an ICMP echo request out.

```
tnsr# ping 203.0.113.1
PING 203.0.113.1 (203.0.113.1) 56(84) bytes of data.
64 bytes from 203.0.113.1: icmp_seq=1 ttl=64 time=0.680 ms
64 bytes from 203.0.113.1: icmp_seq=2 ttl=64 time=0.176 ms
64 bytes from 203.0.113.1: icmp_seq=3 ttl=64 time=0.505 ms
64 bytes from 203.0.113.1: icmp_seq=4 ttl=64 time=0.453 ms
64 bytes from 203.0.113.1: icmp_seq=5 ttl=64 time=0.420 ms
64 bytes from 203.0.113.1: icmp_seq=6 ttl=64 time=0.144 ms
```

(continues on next page)

(continued from previous page)

```
64 bytes from 203.0.113.1: icmp_seq=7 ttl=64 time=0.428 ms
64 bytes from 203.0.113.1: icmp_seq=8 ttl=64 time=0.494 ms
64 bytes from 203.0.113.1: icmp_seq=9 ttl=64 time=0.163 ms
64 bytes from 203.0.113.1: icmp_seq=10 ttl=64 time=0.346 ms

--- 203.0.113.1 ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 9000ms
rtt min/avg/max/mdev = 0.144/0.380/0.680/0.167 ms

tnsr#
```

By default this will follow the host OS routing table, but by specifying a source address, it will use addresses from TNSR:

```
tnsr# ping 203.0.113.1 source 203.0.113.2
PING 203.0.113.1 (203.0.113.1) from 203.0.113.2 : 56(84) bytes of data.
64 bytes from 203.0.113.1: icmp_seq=1 ttl=64 time=0.700 ms
64 bytes from 203.0.113.1: icmp_seq=2 ttl=64 time=0.353 ms
64 bytes from 203.0.113.1: icmp_seq=3 ttl=64 time=0.590 ms
64 bytes from 203.0.113.1: icmp_seq=4 ttl=64 time=0.261 ms
64 bytes from 203.0.113.1: icmp_seq=5 ttl=64 time=0.395 ms
64 bytes from 203.0.113.1: icmp_seq=6 ttl=64 time=0.598 ms
64 bytes from 203.0.113.1: icmp_seq=7 ttl=64 time=0.490 ms
64 bytes from 203.0.113.1: icmp_seq=8 ttl=64 time=0.790 ms
64 bytes from 203.0.113.1: icmp_seq=9 ttl=64 time=0.155 ms
64 bytes from 203.0.113.1: icmp_seq=10 ttl=64 time=0.430 ms

--- 203.0.113.1 ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 9001ms
rtt min/avg/max/mdev = 0.155/0.476/0.790/0.187 ms
```

See also:

- [Diagnostic Utilities](#)

5.7.2 From LAN Client

At this stage a LAN client will be able to connect to the network (port or switch) connected to the LAN interface. It can pull an IP address and other configuration via DHCP, resolve domain names via DNS, and reach hosts beyond TNSR using it as a gateway.

A ping executed on a client will flow through TNSR and replies will return.

5.8 Save the TNSR Configuration

TNSR maintains three separate *configuration databases*: startup, candidate, and running. The running copy is the active configuration. TNSR loads the startup copy at boot time.

To ensure the expected configuration is loaded when TNSR is rebooted, copy the running configuration to the startup configuration after making changes:

```
tnsr# configure
tnsr(config)# configuration copy running startup
```

Optionally, create a *backup copy* of the configuration which can be loaded later if necessary:

```
tnsr(config)# configuration save running backup.xml
```

See also:

- [Configuration Database](#)
- [Configuration Backups](#)

5.9 Next Steps

From here, click the **Next** button at the bottom of the page to continue on to the next section of the documentation, or choose a topic from the table of contents to the left.

Other suggested next steps include:

- [Configure updates](#) (non-trial version only)
- See [more practical examples](#), such as setting up the RESTCONF API
- Configure [IPsec tunnels](#)
- Configure [time synchronization](#)

COMMAND LINE BASICS

The TNSR command line interface (CLI) may seem familiar to administrators who are familiar the CLI of other routers or networking equipment. However, the specific behavior and structure of the TNSR CLI differs in several aspects.

Tip: For a full TNSR CLI command reference, visit [Commands](#).

6.1 Working in the TNSR CLI

6.1.1 Command Prompt

The TNSR CLI command prompt has a several components:

```
<hostname> tnsr<(mode)># <user input>
```

These components are:

hostname The fully qualified hostname of the router.

mode This section of the prompt changes depending on the current mode to indicate that a different subset of commands is available.

See also:

For a list of modes and prompt strings, see [Mode List](#).

user input This area is where a user enters commands and other input.

In this brief example, the router hostname is `router`, and the mode section of the prompt is shown changing when a command enters or exits a mode.

```
router tnsr# configure
router tnsr(config)# interface GigabitEthernet3/0/0
router tnsr(config-interface)# description Management
router tnsr(config-interface)# exit
router tnsr(config)# exit
router tnsr#
```

6.1.2 Command History

The TNSR CLI stores the last 300 commands across sessions. This command history is kept in `~/ .tnsr_history`.

The command history is accessed by pressing `Ctrl-P` (previous command), `Ctrl-N` (next command), or by using the up and down arrow keys.

The number of commands stored by TNSR can be controlled by the `cli history-config lines <count>` command. To restore the default value, use `no cli history-config lines`.

This behavior may also be disabled by the `cli history-config disable` or `no cli history-config enable` commands. Use `cli history-config enable` to turn it back on.

6.1.3 Autocomplete

The TNSR CLI supports case-sensitive tab expansion and prediction for input to speed up interactive work. For example, the first few letters of a command or entity may be typed, depending on context, and then pressing the tab key will complete a portion or all of the remaining input where possible. Additionally, in cases when there is an existing entry or only one possible choice, pressing tab will automatically insert the entire entry. Commands or entities may also be shortened provided the input is not ambiguous.

Tip: Press `?` to show possible completions of the current command when in the middle of a word, or press it between words to show the next available parameter (*Finding Help*).

6.1.4 Keyboard Shortcuts

The TNSR CLI supports several CLI navigation and editing key combinations, including:

Command	Keys
Previous History Command	<code>Ctrl-P</code> or up arrow
Next History Command	<code>Ctrl-N</code> or down arrow
Erase Character	Backspace or <code>Ctrl-H</code>
Erase Word	<code>Ctrl-W</code>
Cursor to Start of Line	<code>Ctrl-A</code>
Cursor to End of Line	<code>Ctrl-E</code>
Clear and Redraw Screen	<code>Ctrl-L</code>
Exit the CLI	<code>Ctrl-D</code>
Context-Sensitive Help	<code>?</code>

6.2 Finding Help

The CLI includes context-sensitive help. At any point, enter a `?` and TNSR will print a list of available commands or keywords that are valid in the current context. Enter a space before the `?` to ensure correct context.

Additionally, the `help` command can be issued in any mode. There are three variations:

help, help commands These are equivalent and print a list of available commands in the current mode.

help mode Prints information about the current mode, including whether or not exiting the mode will cause a commit (*Configuration Database*).

6.3 Starting TNSR

The services required by TNSR to run are enabled by the installer, and they will automatically start at boot time. There is no need to manually stop or start TNSR services during normal operation.

6.3.1 TNSR Service Relationships

TNSR requires the `vpp`, `clixon-backend`, and `clixon-restconf` services.

The `clixon-backend` service is configured to depend on `vpp`, thus:

- If the `vpp` service is restarted, `clixon-backend` will also restart if it is running.
- If the `vpp` service is stopped, `clixon-backend` will stop if it is running.
- If both `vpp` and `clixon-backend` are stopped, then starting `clixon-backend` will also start `vpp`.

Note: TNSR may require additional services depending on features enabled by the TNSR configuration. These will be automatically managed as needed.

6.3.2 Manual TNSR Service Operations

Stop TNSR services:

```
$ sudo systemctl stop vpp clixon-restconf
```

Start TNSR services:

```
$ sudo systemctl start clixon-backend clixon-restconf
```

Restarting TNSR services if they are already running:

```
$ sudo systemctl restart vpp clixon-restconf
```

These services are all daemons and not interactive. To configure TNSR, the administrator must initiate the TNSR CLI separately, as described in [Entering the TNSR CLI](#).

Convenience Alias

For convenience, an alias in the shell can be used to handle this task. For example, the following single line can be added to `~/.bashrc`:

```
alias restarttnsr='sudo systemctl stop vpp clixon-restconf;  
sudo systemctl start clixon-backend clixon-restconf'
```

Note: The changes to `~/.bashrc` will not take effect immediately. Either logout and login again, or source the file by running `source ~/.bashrc` or `. ~/.bashrc`.

The above actions can then be accomplished all at once by running `restarttnsr`.

6.4 Entering the TNSR CLI

The TNSR CLI can be started a few different ways. The command to start the CLI is `/usr/bin/clixon_cli`, but the exact method varies, as discussed in this section.

When started, the TNSR CLI will print the hostname followed by the prompt:

```
tnsr#
```

From that prompt, commands can be entered to view status information or perform other tasks. Throughout this documentation, the router hostname will typically be omitted unless it is required for clarification.

6.4.1 Using the `tnsr` account

TNSR includes a `tnsr` user by default, and this user will automatically load the TNSR CLI at login. To take advantage of this user, login to it directly using `ssh`, or switch to it using `sudo` or `su` from another account.

The behavior of the `tnsr` account varies by platform, and its password can be reset using any account with access to `sudo` (See [Default Accounts and Passwords](#)).

To switch from another user to the `tnsr` user, use `sudo`:

```
$ sudo su - tnsr
```

Alternately, use `su` and enter the password for the `tnsr` user:

```
$ su - tnsr
Password:
```

6.4.2 Using another account

The TNSR CLI can also be started manually from any user.

This command will start the TNSR CLI as the current user, which is ideal to use in combination with [NACM](#):

```
$ /usr/bin/clixon_cli
```

6.4.3 Using root

This command will start the TNSR CLI as root, which generally should be avoided unless absolutely necessary (for example, recovering from a flawed NACM configuration):

```
$ sudo /usr/bin/clixon_cli
```

6.4.4 Current User

From inside TNSR, check the current user as seen by TNSR with `whoami`:

```
tnsr# whoami
real UID/GID: 996/992
effective UID/GID: 996/992
```

(continues on next page)

(continued from previous page)

```
user name: tnsr
home dir: /var/lib/tnsr
shell: /bin/bash
```

6.4.5 Shell Alias

For convenience, the command to launch the TNSR CLI can be added to an alias in the shell. For example, the following line can be added to `~/.bashrc` to run TNSR as the current user:

```
alias tnsrcli='/usr/bin/clixon_cli'
```

Note: The changes to `~/.bashrc` will not take effect immediately. Either logout and login again, or source the file by running `source ~/.bashrc` or `. ~/.bashrc`.

Then the TNSR CLI may be accessed using the alias from the shell, `tnsrcli`.

6.5 Configuration Database

TNSR maintains three separate configuration databases: startup, candidate, and running. These files are stored as XML in plain text files.

startup The configuration loaded when the host boots up.

Note: A restart of TNSR services is not the same as a reboot. If, for example, the clixon services are restarted, TNSR will still be using the running database.

candidate An in-process potential configuration that exists while the TNSR configuration is being actively edited. When committed, this configuration will be accepted as the running configuration by TNSR if it is free of errors.

running The active running configuration, which reflects the current state of TNSR.

Note: These databases are located in `/var/tnsr/` on the host, but these files are not intended to be accessed outside of TNSR.

The configuration database is managed using the `configuration` command from within `config` mode.

6.5.1 Saving the Configuration

For changes to persist between reboots of the TNSR host, the running configuration must be copied to the startup configuration as shown in this example:

```
tnsr# configure
tnsr(config)# configuration copy running startup
```

6.5.2 Viewing the Configuration

To view the configuration databases, use the `show configuration` command followed by the database name, for example:

```
tnsr# show configuration running
```

or:

```
tnsr# show conf run
```

The default output is XML, but the configuration may also be printed in json format by adding `json` to the end of the command.

6.5.3 Reverting to the Startup Configuration

TNSR can also revert to the previously saved startup configuration to remove undesirable changes to the running configuration, should a regression in behavior occur.

For example:

```
tnsr# configure
tnsr(config)# configuration copy startup candidate
tnsr(config)# configuration candidate commit
tnsr(config)# exit
```

Warning: It is not possible to copy the startup configuration directly to the running configuration as that will not result in the settings being active. The configuration must be committed after copying to the candidate.

6.5.4 Configuration Database Commands

These brief examples show other available configuration database management commands.

Delete the candidate database entirely, which if committed will leave TNSR with an empty configuration:

```
tnsr(config)# configuration candidate clear
```

Commit changes made to the candidate database, which if successful will become the running database:

```
tnsr(config)# configuration candidate commit
```

Discard the current candidate database to remove a change that has failed to validate, returning to the running configuration without the attempted changes:

```
tnsr(config)# configuration candidate discard
```

Attempt to validate the current candidate configuration to locate errors:

```
tnsr(config)# configuration candidate validate
```

Load a file from the host into the candidate database. The contents of the file can replace the candidate entirely, or merge a new section into an existing configuration. After loading, the candidate must be committed manually.

```
tnsr(config)# configuration candidate load <filename> [(replace|merge)]
```

Copy the candidate configuration to the startup configuration:

```
tnsr(config)# configuration copy candidate startup
```

Copy the running configuration to either the candidate or startup configuration:

```
tnsr(config)# configuration copy running (candidate|startup)
```

Copy the startup configuration to the candidate configuration:

```
tnsr(config)# configuration copy startup candidate
```

Save either the candidate or running configuration to a file on the host.

```
tnsr(config)# configuration save (candidate|running) <filename>
```

While not a configuration database command directly, the TNSR CLI automatically discards the candidate database if it fails to validate. This behavior can be changed using the following command:

```
tnsr(config)# no cli option auto-discard
```

6.6 Configuration Mode

After starting the TNSR CLI, the administrator is in basic mode and not configuration mode. To enter configuration mode, enter the `configure` command. This command may be abbreviated to `config` and it is also acceptable to write `terminal` after, as a convenience for administrators familiar with IOS who type it out of habit.

All of the following commands are equivalent:

```
tnsr# configure
tnsr# configure terminal
tnsr# config
tnsr# conf t
```

After entering any **one** of the above commands, the prompt changes to reflect the new configuration mode:

```
tnsr# configure
tnsr(config)#
```

After entering other configuration commands, the new configuration is stored in the candidate database (*Configuration Database*). A candidate database may be committed either when all of the required information is present, or when exiting the current context. Some commands are committed immediately.

6.6.1 Navigating Configuration Modes

Certain commands in configuration mode enter other modes, for example, the `interface` command will enter `config-interface` mode when used on an existing interface:

```
tnsr(config)# interface GigabitEthernet3/0/0
tnsr(config-interface)#
```

To leave a mode, use the `exit` command. This will return to the previous, lower mode:

```
tnsr(config-interface)# exit
tnsr(config)#
```

From config mode, using `exit` will return to basic mode:

```
tnsr(config)# exit
tnsr#
```

From any mode, the `exit` command may be repeated until the prompt returns to basic mode.

At that point, if no errors have been encountered by TNSR, all changes will have been committed to the running database. If an error occurs, TNSR will print a message indicating the problem. Solving such problems is covered in [Troubleshooting](#) later in this section.

6.6.2 Removing Configuration Items

Items are removed or negated using `no`, followed by the option to remove. For example, to remove an interface description:

```
tnsr(config)# interface GigabitEthernet0/14/1
tnsr(config-interface)# no description
```

In this case, since there is only one description, removing the the description does not require the existing content of that option. In most cases, the `no` command only requires enough parameters to uniquely identify an entry to be removed or negated.

In certain cases, a partial command may remove multiple items or may be used as a shorthand method of removing a longer entry when the details do not uniquely identify an entry.

For example, this command removes one input ACL from an interface:

```
tnsr(config-interface)# no access-list input acl idsblock
```

Where this shorter version will remove all input ACL entries on an interface:

```
tnsr(config-interface)# no access-list input acl
```

Finally, this form would remove all ACLs of any type from an interface:

```
tnsr(config-interface)# no access-list
```

The `? help` command ([Finding Help](#)) is useful in determining when these actions are possible. For example, the CLI will show `<cr>` (“carriage return”) as an available keyword when testing a command:

```
tnsr(config-interface)# no access-list ?
<cr>
acl          ACL Rule
input        ACL applies to ingress packets
macip        MACIP Rule
output       ACL applies to egress packets
```

Since the help request printed `<cr>` among the choices, the command may be completed by pressing `Enter`.

Interactive Large Delete Confirmation

When performing a delete operation in sensitive areas, TNSR checks the size of the pending change when exiting a mode or committing configuration changes. This feature prevents accidental removal of significant sections of the router configuration.

If TNSR considers a change too large to happen automatically, such as removing the entire OSPF configuration, then TNSR will prompt for confirmation before proceeding.

This feature is disabled by default, but may be enabled as follows:

```
tnsr# configure
tnsr(config)# cli option check-delete-thresholds
```

To disable the feature, precede it with `no`:

```
tnsr# configure
tnsr(config)# no cli option check-delete-thresholds
```

For example, with the feature enabled, attempting to make a large change results in a confirmation prompt:

```
tnsr(config)# route dynamic ospf
tnsr(config-frr-ospf)# no server
Really delete that? [yes/no]: y
tnsr(config-frr-ospf)#
```

6.6.3 Troubleshooting

If a change to the candidate database fails a validation check or application of the change to the system fails for some reason, it is discarded automatically by default. TNSR resets the candidate database to the current contents of the running database to avoid further attempts to apply the faulty configuration contained in the candidate database.

This automatic behavior can be changed, however, in cases where power users want more control to troubleshoot failed configuration transactions:

```
tnsr# configure
tnsr(config)# no cli option auto-discard
```

When auto-discard is disabled, if a configuration commit fails the candidate database retains the faulty configuration data. Further configuration commands may apply additional changes to the candidate database. However, until the configuration data which caused the failure is removed or set to a value which can be successfully applied, no further commit will succeed.

Disabling the auto-discard feature only persists for the duration of the current CLI session in which it was disabled. At the start of a new CLI session, auto-discard will again be enabled by default.

To view the status of the `auto-discard` option, use `show cli`:

```
tnsr# show cli
Discard erred candidate database: true
```

A faulty candidate can be viewed with the `show configuration candidate` command, as described in [Configuration Database](#)

There are three approaches to rectify this situation:

- Issue alternate commands that directly correct the faulty configuration.
- Abandon the attempted configuration:

```
tnsr# configure
tnsr(config)# configuration candidate discard
```

- Remove the fault from the candidate configuration by reverting to the running configuration:

```
tnsr# configure
tnsr(config)# configuration copy running candidate
tnsr(config)# configuration candidate commit
```

6.7 Configuration Backups

The candidate and running databases can be saved to or loaded from files in the host OS. This can be used to make backups, copy configurations to other routers, or similar purposes.

Warning: This procedure only backs up the configuration database. It does not back up other important files such as PKI entries (Certificate Authorities, Certificates, Keys) from `/etc/pki/tls/tnsr/`. Copy those files, and any other modified files such as custom scripts and shell configuration files, outside of TNSR as described in *Backups Outside of TNSR*.

The filenames can take an absolute path and filename, or the path may be omitted to save the file in the directory from which the TNSR CLI was invoked by the administrator. When saving, this file must be writable by the TNSR backend daemon. When loading, this file must be readable by the TNSR backend daemon.

Tip: The best practice is to store backup configuration files in a secure location to prevent unauthorized access to sensitive information.

Saving the running configuration as a backup:

```
tnsr# config
tnsr(config)# configuration save running backup.xml
```

Loading a configuration file from a backup:

```
tnsr# config
tnsr(config)# configuration candidate load backup.xml
tnsr(config)# configuration candidate commit
```

6.7.1 Backups Outside of TNSR

The previous procedure creates and restores the configuration from within the TNSR CLI. In certain cases that method may not be viable, such as when the configuration from an older version of TNSR must be updated (*Updating the Configuration Database*).

In these cases, the TNSR configuration database files in `/var/tnsr` may be accessed directly.

Warning: Unlike operations performed within TNSR, these actions must be performed with elevated privileges, either by the `root` account or using `sudo`.

Warning: This procedure only backs up the configuration database. It does not back up other important files such as PKI entries (Certificate Authorities, Certificates, Keys) from `/etc/pki/tls/tnsr/`. Copy those files, and any other modified files such as custom scripts and shell configuration files, using the same method described here.

To make a configuration backup of the running database:

```
$ sudo cp -p /var/tnsr/running_db ~/backup.xml
```

Warning: The configuration database files may be read while TNSR is running, but TNSR must be stopped when making changes.

To restore a backup to the running and startup databases:

```
$ sudo systemctl stop vpp clixon-restconf
$ sudo cp -p ~/backup.xml /var/tnsr/running_db
$ sudo cp -p /var/tnsr/running_db /var/tnsr/startup_db
$ sudo systemctl start clixon-backend clixon-restconf
```

6.8 Viewing Status Information

Status information can be viewed using the `show` command from either basic or configuration mode.

For a full list of possible `show` commands, enter `show ?`:

```
tnsr# show ?
acl                Access Control Lists
bfd                Bidirectional Forwarding Detection
cli                State of per-session CLI options
clock              Show the current system date and time
configuration      Config DB configuration state
dslite             DS-Lite
gre                GRE tunnels
history-config     Show history configuration
host               Host information
http               HTTP
interface          Interface details
ipsec              IPsec
kea                Kea/DHCP
macip              MACIP Access Control Lists
map                MAP-E/MAP-T
nacm               NACM data
nat                Network Address Translation
neighbor           Neighbors (ARP/NDP)
ntp                NTP
packet-counters    Packet statistic and error counters
route              Show routing info.
span               SPAN mirrors
sysctl             Sysctl parameters
system             System information
unbound            Unbound DNS
version            Show version of system components
vxlan              VXLAN tunnels
```

(continues on next page)

(continued from previous page)

```
tnsr# show version
Version: tnsr-v19.02-1
Build timestamp: Thu Feb 21 17:12:00 2019 CST
Git Commit: 0x40204091
```

6.9 Service Control

Services controlled directly by TNSR can be restarted from within the TNSR CLI in configuration mode.

To control a service, use the `service` command as follows:

```
tnsr# configure
tnsr(config)# service <name> <action>
```

The service name must be one of the following:

- backend** Configuration backend (`clixon_backend`)
- bgp** BGP routing (`bgpd`, `zebra`)
- dataplane** Dataplane (`vpp`)
- dhcp** DHCP (`kea`)
- http** HTTP for RESTCONF API (`nginx`)
- ntp** Time service (`ntpd`)
- ospf** OSPF Routing (`ospfd`, `zebra`)
- ospf6** OSPF6 Routing (`ospf6d`, `zebra`)
- restconf** RESTCONF API (`clixon_restconf`)
- rip** RIP Routing (`ripd`, `zebra`)
- snmp** SNMP Server (`snmpd`)
- unbound** DNS Resolver (`unbound`)

The following `action` types are available:

- start** Start the service if it is not already running.
- stop** Stop the service if it is currently running.
- restart** Stop and restart the service, or start the service if it is not running. This action is not available for the *dhcp* service.
- reload** Reload the service configuration without restarting. This action is available for the *dhcp* and *unbound* services.
- status** Show the current status of the service daemon(s) and the last few log entries.

6.10 Diagnostic Utilities

The TNSR CLI includes convenience utilities for testing connectivity.

6.10.1 Diagnostic Routing Behavior

The utilities in this section behave the same with regard to routing. These utilities will send traffic using the host OS routing table by default unless a specific source address is passed to the command.

6.10.2 Ping

To perform a basic ICMP echo request, use the `ping` command:

```
tnsr# ping <destination host> source <interface IP address>
```

TNSR will send 10 ICMP echo requests to the destination host, waiting a maximum of 12 seconds for a reply. The source address would be a TNSR interface address, which will allow `ping` to send its request using the routing table in TNSR.

The `ping` command supports a number of additional parameters which alter its behavior:

```
tnsr# ping (<dest-host>|<dest-ip>) [ipv4|ipv6] [interface <if-name>] [source <src-  
→addr>]  
      [count <count>] [packet-size <bytes>] [ttl <ttl-hops>] [timeout <wait-sec>]
```

dest-host|dest-ip The target of the ICMP echo request. This may be a hostname, IPv4 IP address, or IPv6 IP address.

ipv4|ipv6 When a hostname is used for the destination, this parameter controls the address family used for the ICMP echo request when the DNS response for the hostname contains both IPv4 (A) and IPv6 (AAAA) records.

interface The TNSR interface from which the ICMP echo requests will originate.

source The source IP address for the ICMP echo requests. This is required to initiate an ICMP echo request using the routing table in TNSR. If omitted, the ICMP echo request will use the host OS routing table.

count The number of ICMP echo requests to send. Default value is 10.

packet-size The size of the ICMP echo request payload, not counting header information. Default value is 56.

ttl The Time To Live/Hop Limit value for ICMP echo requests, which can limit how far they may travel across the network. Default value is 121 hops.

timeout The total time to wait for the command to complete.

6.10.3 Traceroute

To perform a network routing trace to a destination host, use the `traceroute` command:

```
tnsr# traceroute <destination host> source <interface IP address>
```

The source address would be a TNSR interface address, which will allow `traceroute` to send its request using the routing table in TNSR.

As with the `ping` command, there several additional parameters to change the behavior of the trace:

```
tnsr# traceroute (<dest-host>|<dest-ip>) [ipv4|ipv6] [interface <if-name>] [source
↪<src-addr>]
      [packet-size <bytes>] [no-dns] [timeout <seconds>] [ttl <ttl-hos>] [waittime
↪<wait-sec>]
```

Most parameters are the same as those found in `ping` (*Ping*). Only the items that differ are listed here:

no-dns Do not attempt to use DNS to reverse resolve hosts that respond to probes.

waittime Amount of time the command will wait for individual probe responses to return.

Warning: The `traceroute` command requires `/usr/bin/traceroute` to be present in the base operating system. The TNSR package set includes a dependency which will automatically install a package for `traceroute`. It may also be installed manually using `sudo yum install -y traceroute` or a similar command, depending on the host OS package management configuration.

6.11 Basic System Information

The TNSR CLI can set several basic elements about the system itself, which also serves as a good introduction to making changes on TNSR. These settings are made in `config` mode.

Tip: These values are also propagated to SNMP, if configured. See *Simple Network Management Protocol* for information on setting up SNMP.

The following parameters are available:

system contact <text> System contact information, such as an e-mail address or telephone number (`sysContact` in SNMP).

system description <text> A brief description of this TNSR instance, for example its role or other identifying information (`sysDescr` in SNMP).

system location <text> The location of this TNSR instance, for example a physical location such as a building name, room number, rack number/position, or VM host (`sysLocation` in SNMP).

system name <text> The hostname of this TNSR instance (`sysName` in SNMP).

Warning: This setting also changes the hostname in the host operating system to match, replacing any previously configured hostname.

This example shows how to set the above parameters, starting from `master` mode:

```
gw tnsr# configure
gw tnsr(config)# system contact support@example.com
gw tnsr(config)# system description TNSR Lab Router
gw tnsr(config)# system location HQ MDF/Rack 2 Top
gw tnsr(config)# system name labrtr01
labrtr01 tnsr(config)# exit
```

To view the values of these parameters, along with uptime and memory usage, use the `show system` command from either `master` or `config` mode:

```
labrtr01 tnsr# show system
  description: TNSR Lab Router
  contact: support@example.com
  name: labrtr01
  location: HQ MDF/Rack 2 Top
System Parameters:
  object-id: 1.3.6.1.4.1.13644
  uptime: 1303615 seconds
  total-ram: 8004488 KiB
  free-ram: 3236820 KiB
  total-swap: 2932732 KiB
  free-swap: 2932732 KiB

Platform:
  os-name: Linux
  os-release: 3.10.0-957.21.3.el7.x86_64
  os-version: CentOS Linux release 7.6.1810 (Core)
  machine: x86_64

Product:
  product-vendor: Netgate
  product-name: TNSR
  product-model: x
  product-serial: 0
```

BASIC CONFIGURATION

Now that TNSR is installed, it needs additional manual setup.

Note: This section assumes TNSR was installed as described in [Installation](#). Devices pre-loaded with TNSR by Netgate do not require these extra steps.

This section contains information for a manual setup of interfaces. It can also serve as a reference for activating additional hardware added to an existing installation.

7.1 Setup Interfaces

TNSR requires complete control of the network interfaces that it will use. This means that the host operating system must not be attempting to use or control them in any way. The device ID of the interface(s) also must be obtained to inform VPP and TNSR what interfaces to use. The interface link can be tuned through VPP and configured through TNSR.

Warning: The host management interface must remain under the control of the host operating system. It must not be configured as an interface to be controlled by TNSR.

Network interfaces not configured in the installer will be disabled in CentOS during the installation process. The interfaces will need to be re-enabled in TNSR. For a fresh installation of TNSR, skip ahead to [Setup NICs in Dataplane](#).

Interfaces added to the TNSR instance after the initial setup will need to be disabled using the following procedure.

7.1.1 Identify NICs to use with TNSR

To start, locate the network interfaces in use by the host operating system. View a list of network interfaces known to the host OS with this command:

```
$ ip link
```

To determine if a network interface is in use by the host OS, run the following command:

```
$ ip link show up
```

If an interface shows in that list, and its name does not start with `vpp`, then it is under control of the host.

Note: The TNSR installer will automatically mark any interface not configured in the installer for use by TNSR.

Make a note of the network interfaces and their purpose. Note which interface will be used for host management, and which interfaces will be used by TNSR. The host management interface will be left under the control of the operating system, while the remaining interfaces may be used by TNSR. In this example, the host contains four network interfaces: `enp0s20f0`, `enp0s20f1`, `enp0s20f2`, and `enp0s20f3` and TNSR will use `enp0s20f1` and `enp0s20f2`.

7.2 Disable Host OS NICs for TNSR

In order for TNSR to control network interfaces, they must be disabled in the host OS. In most cases this is not necessary, as network interfaces not configured in the installer will be automatically disabled in CentOS during the installation process. For a fresh installation of TNSR, skip ahead to [Setup NICs in Dataplane](#). This section remains to explain how to change interfaces added after initial installation, or for installations which do not contain whitelisted network interfaces.

This is a two-step process. First, the link must be forced down, and then the network interface must be disabled in Network Manager.

Warning: The host management interface must remain under the control of the host operating system. It must not be configured as an interface to be controlled by TNSR. Do not disable the management interface during this step.

For each of the interfaces noted in the last section, manually force the link down:

```
$ sudo ip link set <interface name> down
```

For example:

```
$ sudo ip link set enp0s20f1 down
$ sudo ip link set enp0s20f2 down
```

Next, disable these network interfaces in Network Manager. For each of these interfaces, edit the corresponding startup script:

```
$ sudo vi /etc/sysconfig/network-scripts/ifcfg-<interface name>
```

In each of these files, ensure the following values are set. Add lines if they are not already present in the file:

```
ONBOOT=no
NM_CONTROLLED=no
```

Note: To change an interface from being usable by TNSR to back under host OS control, see [Remove TNSR NIC for Host Use](#).

7.3 Setup NICs in Dataplane

Next, determine the device ID for the interfaces. Start the CLI (*Entering the TNSR CLI*) and run the following command to output the device IDs as seen by the dataplane:

```
tnsr# configure
tnsr(config)# dataplane dpdk dev ?
  0000:02:01.0      Ethernet controller: Intel Corporation 82545EM Gigabit
→Ethernet
  Controller (Copper) (rev 01) ( Active Interface eth0 )
  0000:02:02.0      Ethernet controller: Intel Corporation 82545EM Gigabit
→Ethernet
  Controller (Copper) (rev 01)
  0000:02:03.0      Ethernet controller: Intel Corporation 82545EM Gigabit
→Ethernet
  Controller (Copper) (rev 01)
```

Interfaces under host control will be noted in the output with **Active Interface**. Other listed interfaces are usable by TNSR.

For a fresh installation of TNSR, skip ahead to *Configuring Interfaces for TNSR*, otherwise continue on to identify host interfaces added after TNSR was installed.

7.3.1 Host Interface Name to Dataplane ID Mapping

The output of the `dataplane dpdk dev ?` command includes the device IDs in the first column. The device IDs will map to the network cards in a way that is typically easy to determine. For example:

Table 1: Interface Identifiers

Interface	Identifier
enp0s20f0	0000:00:14.0
enp0s20f1	0000:00:14.1
enp0s20f2	0000:00:14.2
enp0s20f3	0000:00:14.3
enp3s0	0000:03:00.0
enp4s0	0000:04:00.0

The host OS interface name and VPP identifiers contain the same information represented in different ways. They both reference the PCI bus number, slot number, and function number. The Interface name contains the values in decimal while the identifier shown in VPP uses hexadecimal.

Deconstructing the first interface name, it contains the following:

Table 2: Interface Name Components

Component	Interface Value	VPP ID Value
Device Type	en (Ethernet)	n/a
PCI Bus	p0	00
Bus Slot	s20	14 (Decimal 20 in Hex)
Function	f0	.0

Using this pattern, make a note of the VPP identifiers for the next step. In this example, since `enp0s20f1` and `enp0s20f2` are the interfaces to use, the corresponding VPP IDs are `0000:00:14.1` and `0000:00:14.2`.

7.3.2 Configuring Interfaces for TNSR

Next, edit the dataplane configuration. Start the CLI (*Entering the TNSR CLI*) and enter configuration mode:

```
tnsr# configure
tnsr(config)#
```

Add the device IDs of the interfaces to be used by the VPP dataplane, determined above:

```
tnsr(config)# dataplane dpdk dev 0000:00:14.1 network
tnsr(config)# dataplane dpdk dev 0000:00:14.2 network
```

Then commit the configuration:

```
tnsr(config)# configuration candidate commit
```

Restart the VPP dataplane:

```
tnsr(config)# service dataplane restart
tnsr(config)# exit
```

The interfaces will now be available for TNSR. Start the CLI again and run `show interface` and verify that the interfaces appear in the output. The output example below has been shortened for brevity:

```
tnsr# show interface
Interface: GigabitEthernet0/14/1
[...]
Interface: GigabitEthernet0/14/2
[...]
Interface: local0
[...]
```

The TNSR interface name also reflects the type, followed by the PCI Bus/Slot/Function ID of each interface, using the same hexadecimal notation as VPP.

Note: Once TNSR attaches to interfaces in this way, they will no longer be shown as devices in the host OS. To return a network interface back to host OS control, see [Remove TNSR NIC for Host Use](#).

One exception to this behavior is Mellanox network interfaces as they use the same driver for both host OS and DPDK, they still appear in the host OS.

Customizing Interface Names

The default interface names, such as `GigabitEthernet0/14/1`, may be customized by an administrator. To customize the names, the PCI ID of the device must be known. The custom names can be used anywhere that an interface name is necessary in TNSR.

Note: Only dataplane hardware interface names may be customized in this way. Interfaces from virtual sources such as loopback, IPsec, and GRE cannot be renamed.

The command to rename interfaces is `dataplane dpdk dev <pci-id> network name <name>`. To activate the change, the dataplane must be restarted after making the name change.

This example changes the name of `GigabitEthernet0/14/1`, PCI ID `0000:00:14.1`, to `DMZ`:

First, look at the list of interfaces. Note that the interface is in the list with its original name:

```
tnsr# show interface
Interface: GigabitEthernet0/14/1
[...]
Interface: GigabitEthernet0/14/2
[...]
Interface: local0
[...]
```

Next, remove any references to the interface from TNSR, and then remove the interface configuration entirely:

```
tnsr(config)# no interface GigabitEthernet0/14/1
```

Now set the name of the device, then restart the dataplane:

```
tnsr(config)# dataplane dpdk dev 0000:00:14.1 network name DMZ
tnsr(config)# service dataplane restart
```

After the dataplane restarts, the interface will appear in the list with its new name:

```
tnsr# show interface
Interface: DMZ
[...]
Interface: GigabitEthernet0/14/2
[...]
Interface: local0
[...]
```

To change the name back at a later time, all references to the interface must first be removed, and then the name can be reset:

```
tnsr(config)# no interface DMZ
tnsr(config)# no dataplane dpdk dev 0000:00:14.1 name
tnsr(config)# service dataplane restart
```

7.3.3 Troubleshooting

If the interfaces do not appear in the `show interface` output, the default driver did not attach to those interfaces and they may require a different driver instead. To see a list of available drivers, use the following command from config mode:

```
tnsr(config)# dataplane dpdk uio-driver ?
  igb_uio           UIO igb driver
  uio_pci_generic   Generic UIO driver
  vfio-pci          VFIO driver
```

To enable a different driver, complete the command using the chosen driver name, then commit the configuration and restart the dataplane.

Note: Ethernet 700 Series Network Adapters based on the Intel Ethernet Controller X710/XL710/XXV710 and Intel Ethernet Connection X722 are not compatible with the `uio_pci_generic` DPDK driver. For these devices, use the `igb_uio` driver instead.

Note: Mellanox devices use RDMA and not UIO, so changing this driver will not have any effect on their behavior. If a Mellanox device does not appear automatically, TNSR may not support that device.

```
tnsr(config)# dataplane dpdk uio-driver igb_uio
tnsr(config)# configuration candidate commit
tnsr(config)# service dataplane restart
tnsr(config)# exit
```

Then attempt to view the interfaces with `show interface` again. If they are listed, then the correct driver is now active.

Warning: When using the `vfio-pci` driver, the DPDK IOVA mode must be explicitly set to `pa`. See [DPDK Configuration](#) for more details.

7.4 Setup QAT Compatible Hardware

TNSR Supports hardware compatible with Intel® QuickAssist Technology, also known as QAT, for accelerating cryptographic and compression operations.

This hardware can be found in CPIC cards as well as many C3000 and Skylake Xeon systems. Netgate [XG-1541](#) and [XG-1537](#) hardware has an add-on option for a CPIC card.

7.4.1 Setup Process

Enable SR-IOV in the BIOS

SR-IOV is required for QAT to function in TNSR. SR-IOV enables Virtual Functions which are required for binding by crypto devices.

The procedure to enable SR-IOV varies by platform. Generally this involves rebooting the hardware and entering the BIOS setup, making the change, and then saving and rebooting. The exact location of the SR-IOV option also varies in different BIOS implementations.

Note: Netgate devices which ship with a CPIC card preinstalled will have this step completed at the factory, but double check the BIOS to ensure it is set as expected.

Disable VT-d in the BIOS

Certain combinations of hardware may experience problems with QAT when VT-d is enabled in the BIOS. As such, we recommend disabling VT-t in the BIOS for the best possible experience with QAT. Netgate [XG-1537](#) and [XG-1541](#) devices with a DH895xcc QAT CPIC card installed are known to have this limitation.

Note: One exception to this is the [SG-5100](#) device, which is capable of using QAT while VT-d is active in its BIOS.

The procedure to disable VT-d varies by platform. The setting is typically located under **Advanced > Chipset Configuration > North Bridge > IIO > VT-d** or along a similar path.

If VT-d and QAT are incompatible, the problem can manifest in a few different ways, including:

- IPsec tunnels may come up but drop packets or otherwise fail to pass traffic.
- Errors may appear on the console when the dataplane tries to send buffers to the QAT device:

```
[110772.063766] DMAR: [DMA Read] Request device [04:01.0] fault addr 406482000_
↪[fault
   reason 06] PTE Read access is not set
[110773.059440] DMAR: DRHD: handling fault status reg 102
```

- The number of used “in flight” buffer resources will continually increase as traffic attempts to pass through IPsec, as observed under `used_resources` in the output of the shell `sudo vppctl show dpdk crypto devices` CLI command.

Enable IOMMU in grub

IOMMU (Input–Output Memory Management Unit), which in this context is also known as Intel VT-d, must be enabled in `grub` for QAT to function. It functions similar to PCI passthrough, allowing the dataplane to access the QAT device.

To enable IOMMU in `grub`:

- Open `/etc/default/grub` in a text editor (as root or with `sudo`)
- Locate the line starting with `GRUB_CMDLINE_LINUX`
- Check if that line includes `intel_iommu=on iommu=pt`
- If those parameters are not included on the line, append them to the end, before the end quote.
- Save and exit the text editor
- Run one following commands (depending on how the device boots):
 - Legacy: `sudo grub2-mkconfig -o /boot/grub2/grub.cfg`
 - UEFI: `sudo grub2-mkconfig -o /boot/efi/EFI/centos/grub.cfg`
- Reboot the device

Change the uio driver to igb_uio

Next, change the TNSR dataplane `uio` driver to `igb_uio`:

```
tnsr# configure
tnsr(config)# dataplane dpdk uio-driver igb_uio
```

Note: As of TNSR 20.02, `igb_uio` is the default driver.

Configure the QAT PCI device in TNSR

Next, configure the QAT device in TNSR.

To configure this device, first locate its PCI ID. TNSR will print the PCI ID when viewing possible parameters for dataplane devices

```
tnsr(config)# dataplane dpdk dev ?
0000:03:00.0      Ethernet controller: Intel Corporation Ethernet Connection_
↪X552 10 GbE SFP+
0000:03:00.1      Ethernet controller: Intel Corporation Ethernet Connection_
↪X552 10 GbE SFP+
0000:04:00.0      Co-processor: Intel Corporation DH895XCC Series QAT
0000:05:00.0      Ethernet controller: Intel Corporation I350 Gigabit Network_
↪Connection (rev 01) ( Active Interface enol )
0000:05:00.1      Ethernet controller: Intel Corporation I350 Gigabit Network_
↪Connection (rev 01)
```

In this instance, the following line from the output is for the QAT device:

```
0000:04:00.0 Co-processor: Intel Corporation DH895XCC Series QAT
```

The first value printed on the line is the PCI ID, 0000:04:00.0.

Now, tell TNSR the device at that address is a crypto device:

```
tnsr(config)# dataplane dpdk dev 0000:04:00.0 crypto
```

If TNSR is running in a virtual machine and the QAT device is passed through from the hypervisor host system using SR-IOV, use `crypto-vf` at the end of the command instead. When the device is defined with `crypto-vf`, the dataplane uses the Virtual Function (VF) instead of the Physical Function (PF), since the PF is not directly available in a virtual machine.

Note: Typically a VF can be identified by the string `Virtual Function` printed in the device description listed by `dataplane dpdk dev ?`. Some platforms may not make this distinction visible to TNSR, so the general guideline is to use `crypto-vf` when running in a virtual machine and `crypto` otherwise.

Note: TNSR will only display device types which are usable by the dataplane. This means:

- If a PF is available, it is usable by the dataplane and will appear in the device list.
 - If a VF is available without a corresponding PF, the VF is usable by the dataplane and will appear in the device list.
 - If both a VF and corresponding PF are available, only the PF is usable by the dataplane and thus only the PF will appear in the device list.
-

Activate and check the settings

When viewing the XML configuration with `show configuration running`, it will contain settings similar to the following example. Note that if other dataplane options are present in the configuration, those will also be visible. Here is how it looks once configured:

```
<dataplane-config>
  <dpdk>
    <dev>
      <id>0000:04:00.0</id>
      <device-type>crypto</device-type>
    </dev>
    <uio-driver>igb_uio</uio-driver>
```

(continues on next page)

(continued from previous page)

```
</dpdk>
</dataplane-config>
```

After configuring the `crypto` device and `uio` driver, TNSR will commit the settings to the dataplane configuration.

To activate the new settings, restart the dataplane.

```
tnsr(config)# service dataplane restart
tnsr(config)# exit
tnsr#
```

Lastly, using the shell command, verify that VPP can see the crypto device:

```
tnsr# shell sudo vppctl show dpdk crypto devices
0000:04:00.0_qat_sym      crypto_qat      up
  numa_node 0, max_queues 2
  free_resources 0, used_resources 1
  SYMMETRIC_CRYPTO, SYM_OPERATION_CHAINING, HW_ACCELERATED, IN_PLACE_SGL, OOP_SGL_IN_
→SGL_OUT, OOP_SGL_IN_LB_OUT, OOP_LB_IN_SGL_OUT, OOP_LB_IN_LB_OUT
  Cipher: none, aes-cbc-128, aes-cbc-192, aes-cbc-256, aes-ctr-128, aes-ctr-192, aes-
→ctr-256, aes-gcm-128, aes-gcm-192, aes-gcm-256
  Auth: none, md5-96, sha1-96, sha-256-96, sha-256-128, sha-384-192, sha-512-256
```

7.4.2 Troubleshooting

If the QAT device does not appear in the `show dpdk crypto devices` output, or it only shows an AES-NI device, then VPP can not see the crypto device. To correct this, first verify the QAT drivers are loaded, VFs exist for the QAT device, and grub `BOOT_IMAGE` is passing the necessary iommu parameters.

Verify IOMMU parameters:

```
$ dmesg | grep iommu
```

The following parameters should appear somewhere on the `BOOT_IMAGE` line in the `dmesg` output:

```
intel_iommu=on iommu=pt
```

Verify that the QAT drivers are loaded in the operating system:

```
$ lsmod | grep qat
qat_dh895xccvf      13281  0
qat_dh895xcc        13510  0
intel_qat           141755  2 qat_dh895xccvf,qat_dh895xcc
dh_generic          13286  1 intel_qat
rsa_generic         18819  1 intel_qat
authenc             17776  1 intel_qat
```

Verify Virtual Functions (VFs) exist for the QAT device:

```
$ lspci | grep QAT | wc -l
```

The number of listings are dependent on how many threads VPP uses to process packets. At minimum there will be at least three entries, but there may be many more. The lines will look similar to this example:

```
04:00.0 Co-processor: Intel Corporation DH895XCC Series QAT
04:01.0 Co-processor: Intel Corporation DH895XCC Series QAT Virtual Function
04:01.1 Co-processor: Intel Corporation DH895XCC Series QAT Virtual Function
```

TNSR stores the device Physical Function (PF), 04:00.0 for example, in its configuration because the VFs do not yet exist at boot time. They are created by `clixon-backend` when it processes the `crypto` device. Then, the allocated VFs on the PF have their addresses written to `startup.conf`.

The VFs are bound to `igb_uio` because `igb_uio` is a driver which allows a userspace process to do RDMA on buffers that are used by a PCI device.

If the drivers are loaded and the VFs show under `lspci`, then verify `/etc/vpp/startup.conf` has the appropriate `dppdk` settings. The `igb_uio` driver must be present and the PCI IDs of TNSR interfaces along with one of the VFs for the QAT device:

```
dppdk {
    uio-driver igb_uio
    dev 0000:04:01.0
    dev 0000:05:00.1
    dev 0000:03:00.0
    dev 0000:03:00.1
}
```

If that looks correct, verify `igb_uio` is being used by the QAT VF and interfaces:

```
$ sudo vppctl show pci all | grep igb_uio
0000:03:00.0  0  8086:15ac  2.5 GT/s x1  igb_uio
0000:03:00.1  0  8086:15ac  2.5 GT/s x1  igb_uio
0000:04:01.0  0  8086:0443  unknown      igb_uio
0000:05:00.1  0  8086:1521  5.0 GT/s x4  igb_uio
```

Physical TNSR interfaces will display there in addition to the QAT VF ID, which matches the QAT VF ID configured for `dppdk` in `/etc/vpp/startup.conf`.

If any of those tests do not provide the expected output, then reboot the system and check again. Ensure the TNSR services and VPP are running, and then check the VPP QAT status again.

```
$ sudo vppctl show dppdk crypto devices
```

If there is still no output, verify the PCI ID for the crypto device specified in TNSR is accurate. It must be the first PCI ID displayed by `lspci | grep qat`. Then verify the PCI ID of the next listing in that output (first VF device) is specified in `/etc/vpp/startup.conf` properly and also the same PCI ID seen by VPP when running:

```
$ sudo vppctl show pci all | grep igb_uio
```

7.5 Remove TNSR NIC for Host Use

If TNSR is controlling a network interface that should be used by the host OS, it can be returned to host OS control in a few steps.

7.5.1 Locate the Interface

First, identify the interface in question. The PCI ID and Linux interface name are required to proceed, and [Host Interface Name to Dataplane ID Mapping](#) explains the relationship between these interface names and IDs.

In this example, the TNSR interface `GigabitEthernet0/14/3` will be returned to the host OS. Based on the name, the PCI ID is `0000:00:14.3`, and converting from hexadecimal to decimal yields the Linux interface name `enp0s20f3`. This is determined based on PCI bus 0, Bus slot 20 (decimal), function 3.

7.5.2 Remove the Interface from TNSR

First, remove any configuration items using the interface. The interface could be present in several places, so inspect the entire running configuration for references to this interface and then remove them.

Next, remove the interface configuration itself:

```
tnsr# configure
tnsr(config)# no interface GigabitEthernet0/14/3
```

If the interface was manually specified in the dataplane by PCI ID as mentioned in [Configuring Interfaces for TNSR](#), that must be also be removed. This will be present in the running configuration inside the `<dataplane>` section, if one exists. To remove the configuration, follow this example using the correct PCI ID:

```
tnsr(config)# no dataplane dpdk dev 0000:00:14.3
```

Save the configuration after making these changes, as the next steps will involve actions that may result in the startup configuration being used by TNSR:

```
tnsr(config)# configuration copy running startup
```

Exit the TNSR CLI.

7.5.3 Edit the Host Interface Configuration

The network manager interface configuration scripts are located in `/etc/sysconfig/network-scripts/`. This directory will contain an interface configuration script for the Linux interface name determined above, in the form of `ifcfg-<name>`. In this example, this is `ifcfg-enp0s20f3`.

From a shell on the host OS, edit the file for this interface using `sudo`, for example:

```
$ sudo vi /etc/sysconfig/network-scripts/ifcfg-enp0s20f3
```

Inside that file change `ONBOOT` to `yes`:

```
ONBOOT=yes
```

Remove the `NM_CONTROLLED` line, if one is present.

7.5.4 Reactivate the Host Interface

At this point, the interface is ready to return to host OS control. There are two methods to complete the process: Reboot the host, or manually reactivate the interface.

Reboot

The fastest and easiest option is to **reboot the host**. This will allow the host to naturally locate and resume control of the device.

Warning: All traffic processing by TNSR will stop while the host is rebooting!

Reboot the host from the shell as follows:

```
$ sudo shutdown -r
```

Manually Reactivate

Warning: The following procedure is advanced and we do not recommend using this method. We strongly advise rebooting the host instead.

There is also a manual method which may be used if a reboot is not feasible.

First, stop the dataplane and related services:

Warning: All traffic processing by TNSR will stop while this service is stopped!

```
$ sudo systemctl stop vpp
```

Next, start a root shell and unbind the device from the current driver (TNSR):

```
$ sudo -s
# echo '0000:00:14.3' > '/sys/bus/pci/devices/0000:00:14.3/driver/unbind'
```

Warning: Note the use of the PCI ID in both locations in the command, and the use of quotes around parameters.

That leaves the device unbound. Now it must be returned to a host kernel driver. The name of this driver depends on the hardware. For most Netgate TNSR devices this will be `igb`, as in the following example.

Still using the root shell from the previous command, bind the interface to the driver as follows:

```
# echo '0000:00:14.3' > '/sys/bus/pci/drivers/igb/bind'
```

Lastly, start the dataplane and related services:

```
$ sudo systemctl start clixon-backend
```

7.5.5 Configure the Host Interface

At this point the interface is now under host OS control and will be listed in the output of `ip` and similar commands.

```
$ ip addr show dev enp0s20f3
5: enp0s20f3: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN group_
↳ default qlen 1000
    link/ether 00:08:a2:09:95:b4 brd ff:ff:ff:ff:ff:ff
```


The interface configuration in the host OS can be used to change the interface behavior as needed. The default behavior is to act as a DHCP client. This can be changed by editing the interface configuration file noted in [Edit the Host Interface Configuration](#). Consult the CentOS documentation for additional details.

UPDATES AND PACKAGES

TNSR software updates are available to download over the Internet using Linux package management tools (RPM, yum). The settings required to communicate with the software repository containing TNSR updates are preconfigured on TNSR. Connections to the Netgate TNSR repository must be authenticated using a valid signed client certificate.

Warning: Trial versions of TNSR cannot be updated. Reinstall with a full version of TNSR or install a new trial version.

Note: All versions of TNSR, including trial versions, can update the operating system packages even without the update certificate in place. Only TNSR-related packages require authentication to update.

This guide explains how to obtain and install the required client certificate on a TNSR instance.

Warning: Portions of this process are not final and may change.

Commands must be executed on the TNSR instance to generate an X.509 certificate signing request. The request must then be submitted to Netgate for signing. Once the request has been signed and a certificate has been generated, the certificate must be downloaded and installed in TNSR.

Note: While it is possible to create the certificate outside of TNSR and import it afterward, this guide only demonstrates using TNSR directly. See [Public Key Infrastructure](#) for more details about creating and importing certificates.

At a high level, the steps involved in the process can be summarized as:

8.1 Generate a Key Pair

This guide uses the TNSR CLI `pki` commands documented in [Public Key Infrastructure](#) to generate cryptographic keys that can be used for secure communications and authentication.

Warning: When creating keys and certificates for updates, the name of each component **must** be `tnsr-updates`, which is the name required by the software repository configuration.

The first step is to generate a set of cryptographic keys:

```
tnsr# pki private-key tnsr-updates generate
-----BEGIN PRIVATE KEY-----
[...]
-----END PRIVATE KEY-----
tnsr#
```

Note: This command can be run only once successfully. Subsequent attempts will result in an error unless the existing key is deleted.

This new `tnsr-updates` key object contains the private key, which is secret, and a public key, which is included in the certificate.

The same key pair can be used as the basis for multiple certificate signing requests. If a certificate expires, is accidentally deleted, or needs to be replaced for any other reason other than the keys being compromised, generate a new signing request using the existing key pair.

8.2 Generate a Certificate Signing Request

The Certificate Signing Request (CSR) contains a public key derived from the key pair generated in the previous step, plus attributes that uniquely identify the requester. A CSR is signed by a Certificate Authority to generate a certificate.

To generate a CSR, first set values which identify this TNSR instance:

```
tnsr# pki signing-request set common-name tnsr-example.netgate.com
tnsr# pki signing-request set country US
tnsr# pki signing-request set state Texas
tnsr# pki signing-request set city Austin
tnsr# pki signing-request set org Netgate
tnsr# pki signing-request set org-unit Engineering Testing 1 2 3
```

For the **Common Name**, enter the fully qualified domain name or Public IP address of the TNSR instance. For the other fields, enter information about the name and location of the organization controlling this TNSR instance.

A **Digest Algorithm** is also required to sign the request:

```
tnsr# pki signing-request set digest sha256
```

View the values that have been set before generating the request:

```
tnsr# pki signing-request settings show
Certificate signing request fields:
  common-name: tnsr-example.netgate.com
  country: US
  state: Texas
  city: Austin
  org: Netgate
  org-unit: Engineering Testing 1 2 3
  digest: sha256
```

Any typos can be corrected by re-running the appropriate `set` commands.

When all values are correct, generate the request:

Warning: As with the key pair, the request must have the name `tnsr-updates`.

```
tnsr# pki signing-request tnsr-updates generate
-----BEGIN CERTIFICATE REQUEST-----
MIICzTCCAbUCAQAwgYcxITAfBgNVBAMMGHRuc3ItZXhhbXBsZS5uZXRnYXR1LmNv
bTELMakGA1UEBhMCVVMxZjAMBgNVBAgMBVRleGFzMQ8wDQYDVQQHDAZBdXN0aW4x
EDAOBgNVBAoMB05ldGdhGUxIjAgBgNVBAsMGUVuZ2luZWVyaW5nIFRlc3Rpbmcg
MSAyIDMwgGgEiMAOGCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQDAUxpX5KYNnult
7xIKV5ES6kPMDtBHqXB7d2fywtqfI/UVvV9+LhCHLL0z8ovqq/GcHioddCBQH63a
+Uqh0cMIZVOWRQhe7eYMO3GmHMyuxz6P5eW03E9d/3sT0rL+fUDH8CVWwjmwX0tC
ldP3PADH4ennxqaWk0+lHga0Dm93hrErX5crzJMyZpGZ/BXfDY0+0uxktZOHIsSb
9gDtEN2534I2wk0hm6mFashDWxmYpcb8ventcVwtEOQGABYnsCg8z3VwcPQY6x9k
YIKFuQM3U8hZ2y6oEjjPqfsc+GnZ6b+7bWnck7tITqz6FQwnSW3sKvXkwsyeDnEa
3eyIjSrFagMBAAGgADANBgkqhkiG9w0BAQsFAAOCAQEAAetjRqn6IoekxZErrPvZf
encbvedPUTLSEbGF923PmpmH5KBAOe4QMT2wEA7dWd5Geu0EA5+6/QlvQh3kl1yU
bzDqRASjl67cKFxp6fL2iDkvoaGf+PusLGM3eQthGzF6t7q6cH1500ANVbrLZws2
qu09evqHgPCJkOhcmPLXSGgitMJwH7EBSmySsZPuEyUCsozA8YLSdLM0dxU5PQnX
XesDhG0AMcFhu34nmsUrCqJwi3CM4ruLTlYseVvyZDyjhTEWuCp9lZf7jzRl2qEF
afis853CjtURIEkfzeKIqqacr1Y0XXt119DtKDz19Z4sWu3C1PsdciOgalCnSVHh
5g==
-----END CERTIFICATE REQUEST-----
```

TNSR will print the CSR data to the terminal, as shown above. Copy the text, including the lines containing BEGIN CERTIFICATE REQUEST and END CERTIFICATE REQUEST, and save it to a file.

8.3 Submit the Certificate Signing Request

To generate a signed certificate, the signing request must be submitted to Netgate. Netgate will sign the request with a Certificate Authority key trusted by the TNSR update repository servers.

8.3.1 Required Customer Information

The certificate signing request must be accompanied by information Netgate can use to identify the customer and validate the request. This information varies by platform.

TNSR Device or ISO Install

For customers using a device preloaded with TNSR or installing TNSR from an ISO image, the certificate signing support request must be accompanied by information that Netgate can use to validate the request. Netgate must be able to determine that the request is being sent from an authorized user on an account that has an appropriate TNSR purchase.

For example, send the support request from the same e-mail address which was used when making the TNSR purchase and include an order number and other relevant information in the support request when submitting the CSR.

TNSR in AWS

For AWS customers, two additional pieces of information are necessary to validate the status of customer accounts before Netgate can sign a certificate:

- The AWS Customer ID

- The AWS Instance ID

Note: When registering a TNSR instance to obtain a client certificate, Netgate must be able to prove that this instance of TNSR is a valid instance of the currently published AWS image. To do this, Netgate utilizes the AWS API that indicates which TNSR image the specified instance ID is an instance of. This is the only use of the customer instance ID, which is not stored or retained in any way.

The **AWS Customer ID** can be found using the instructions at <https://docs.aws.amazon.com/general/latest/gr/acct-identifiers.html>

The **AWS Instance ID** can be retrieved from the EC2 Web Console:

1. Navigate to <https://console.aws.amazon.com/ec2/>
2. Click **Instances**
3. Click the box next to the TNSR instance to select it
4. The **AWS Instance ID** is displayed at the bottom of the page under the **Description** tab

8.3.2 Create a Support Request for the CSR

Using the CSR and customer information, submit a request on the Netgate Support Portal.

Warning: The following steps are still under design and development and may change at any time.

1. Navigate to <https://go.netgate.com/support/login>
2. Log in with an existing account using an email address and password, or register a new account using the **Sign Up** button and following the prompts
3. Create a new support request with the following properties:
 - Department** Select Netgate Global Support
 - Software Product** Select the matching purchased TNSR product, either TNSR Business or TNSR Enterprise
 - Platform** Choose the value that matches where TNSR is running, for example TNSR in AWS, Netgate XG-1541 1U, or Whitebox / Other
 - General Problem Description** Select TNSR Certificate Authorization
 - Support Level** Choose the support level that matches the purchased TNSR product, TNSR Business, TNSR Business Plus, or TNSR Enterprise
 - AWS Instance ID** For TNSR on AWS customers only, The ID for this TNSR instance located previously
 - AWS Customer ID** For TNSR on AWS customers only, the AWS Customer ID located previously
 - Order Number** For device and ISO customers, the order number of the TNSR purchase for this device
4. Include any other necessary identifying information in the **Description** field
5. Click **Attach file** and attach the file containing the CSR text
6. Submit the support request

8.4 Retrieve the signed certificate

Warning: The following steps are still under design and development and may change at any time.

Once the certificate signing request has been signed by Netgate, the status of the support request will be updated to reflect that the certificate is ready.

When this occurs, download the signed certificate:

1. Navigate to <https://go.netgate.com/support/login>
2. Locate the support request
3. Download the attached signed certificate file

8.5 Install the certificate

With the signed certificate in hand, it can now be installed on the TNSR instance:

Warning: As with the key and CSR, the name of the certificate must be `tnsr-updates`.

```
tnsr# pki certificate tnsr-updates enter
Type or paste a PEM-encoded certificate.
Include the lines containing 'BEGIN CERTIFICATE' and 'END CERTIFICATE'
-----BEGIN CERTIFICATE-----
MIIE7DCCAtSgAwIBAgIJANbZBxsCVDpvMA0GCSqGSIb3DQEBCwUAMHxwCzAJBgNV
BAYTA1VTMQ4wDAYDVQQIDAVUZXBhczEPMA0GA1UEBwwGQXVzdGluMRAwDgYDVQQK
DAdOZXRNyYXRlMRgwFgYDVQLDA9OZXRnYXRlIFROU1IgQ0ExGDAWBgNVBAMMD051
dGdhdGUgVE5TUjBDQTAeFw0xODA0MzAxNTE1MDFaFw0xODA1MzAxNTE1MDFaMIGH
MSEwHwYDVQQDBh0bnNyLWV4YVw1bGUubmV0Z2F0ZS5jb20xCzAJBgNVBAYTA1VT
MQ4wDAYDVQQIDAVUZXBhczEPMA0GA1UEBwwGQXVzdGluMRAwDgYDVQQKDAdOZXRN
YXRlMSIwIAYDVQQQLDBlFbmdpbmVlcmduZyBUZXR0aW50aW50aW50aW50aW50aW50
hkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAwFMAV+SmDZ7tbe8SCleREupDzA7QR6lw
e3dn8sLanyP1Fblffi4Qhy9M/KL6qvxnB4qHXQgUB+t2v1KodHDCGVTsEUIXu3m
DDtxphzMrsc+j+XlJtXPxf97E9Ky/n1Ax/AlVsI5sF9LQpXT9zwAx+Hp58amlpNP
pr4GtA5vd4axK1+XK8yTmMaRmfV3w2KPtLsZLWThyLEm/YA7RDdud+CNsJNIZup
hWRIQ1sZmKXG/L3p7XfCLRDkBgAcjbAoPM91cHD0GOsfZGCChbkDN1PIWdsuqBI4
z6n7HPhp2em/u21p3JO7SE6s+hUMJ0lt7Cr15MLMng5xGt3siI0qxQIDAQABO20w
azAJBgNVHRMEAjAAMBEGCWCsSAGG+EIBAQQEAWFoDAdBgNVHQ4EFgQUXp0sedA8
QS34KxEmzZJInKwJZKQwHwYDVR0jBBgwFoAU8CpQYHQGB9CuwnHWU0lUnf7WE50w
CwYDVR0PBAQDAgXgMA0GCSqGSIb3DQEBCwUAA4ICAQC+6M81sTW9c/NL1Ls1ziQ
LWWd0L3qc7Q1R6r+HdouU2R//+gP2ylHJe1CM9kjCqHSQos5y+BDJ1/cbrV5JR5U
cnA2s54uePzGZGk89vZHCcUkuXDIGloU8q+p6e7pIyLoJxRU99psj8gT4nUBcczD
W+Vb7x4fotekPwXNWohsRsAXSPqEKbwuf03H4ntfmXLMHSq/qWmv1/g2nH79DRRN
M+AIsEyKL1XwGljY4mjblsOV8PY42LAjnsf7x+LZXnLSYL+9jZgt1A3U8FnQn4Wd
cSEUDDPE5yAj7xye96AAE7ayHtrBLKqbrVQXzVUX8xYQKroXyt1WabMnTdhZxu7K
ZM92H2OglSW2VO1ABjzBIIPPJ2pvCZWvt4XM1krmyTJEsem+U3oByY/wGp93DN0e
S0sM7GMBEj8+aYNgEYIrVcX63VKy3dCLWjZpldwH1v8BNwJn/npWP0MbIh0EIE7/
WeqGTJu86UVKzuezi1sPkUjgP0cdGJHHMrGB8Q8uJ4ReHdRLs7Rs6CK00F2v68iQ
MyILSwy3cnlsxDnsm3JGIhXkm5aVCKLhBV0EM8GXJtW49ftP9ts0DKM3DWLLe82p
CG4IiLHO/nlVMEe0Hn5xE05r+GjYy8vDLJvAukDaet9li3ZaPAOFHHzgLxNhWaPF5
jiSpPVrJiAlsJCv6Fy2FvA==
```

(continues on next page)

(continued from previous page)

```
-----END CERTIFICATE-----
tnsr#
```

After successfully installing the certificate, TNSR can now download software updates from the repository.

8.6 Package Management

The package management commands allow the operator to install new software packages as well as discover and perform updates for installed packages.

8.7 Package Information Commands

There are three commands which query the package database.

A `<pkg-glob>` is a simple regular expression. It consists of a string of alphanumeric characters which is optionally prefixed or suffixed with a `*` character. The `*` character indicates zero or more characters.

For example:

abc	matches only the package abc and would not match abcd.
*abc	matches abc or zabc and would not match abcz.
abc*	matches abc or abcz and would not match zabc.
abc	matches any package with abc contained anywhere in its name.
*	matches any package.

Tip: Do not escape or quote the glob as would typically be required by a Unix shell. The glob `abc*` is **not** the same as `abc*`.

The first two commands have qualifiers that limit the scope of the packages to all, installed, or updatable packages. These limitations are optional, and if not specified then it defaults to all packages in the database.

To display detailed information on packages:

```
tnsr# package info [ available | installed | updates ] <pkg-glob>
```

Warning: package information is limited to the first 25 packages found. If a query returns more items, a more specific `pkg-glob` must be used to narrow the search.

To display a simple listing of package names and versions for all matching packages:

```
tnsr# package list [ available | installed | updates ] <pkg-glob>
```

The `search` command searches for a string in either the package name or description. The output includes the package name and description of the package. The search term is literal, it is not a regular expression or glob:

```
tnsr# package search <term>
```

8.8 Package Installation

Warning: Recommended procedure is to reboot the router after any package install, remove, or upgrade operation.

To install a package and its required dependencies:

```
package install <pkg-glob>
```

To remove a package:

```
package remove <pkg-glob>
```

To upgrade a package:

```
package upgrade [ <pkg-glob> ]
```

8.9 Updating TNSR

Warning: Trial versions of TNSR packages cannot be updated. Reinstall with a full version of TNSR or install a new trial version. The operating system may be updated, but not TNSR.

With a signed client certificate from Netgate in place, TNSR has access to the Netgate software repositories which contain important updates to TNSR. These updates can be retrieved using the `package` command in the TNSR CLI, or `yum` in the host OS shell.

Note: Updating TNSR will also update the operating system. Even when there are no TNSR updates available, it is a good practice to periodically perform an update to obtain important operating system updates such as security vulnerability mitigations.

See also:

Most of this document covers in-place updates. For information on updating by redeploying/reinstalling, see [Upgrading by Redeploying TNSR](#).

8.9.1 Pre-Upgrade Tasks

Before updating TNSR, perform the following tasks:

- Make sure the signed certificate is in place ([Install the certificate](#))
- Make sure the TNSR instance has working Internet connectivity through the host OS management interface
- Take a backup of the running and startup configurations ([Configuration Backups](#))

8.9.2 Updating via the TNSR CLI

The easiest way to update TNSR is from within the TNSR CLI itself.


```
tnsr# package upgrade
```

That command will download and apply all available updates. Afterward, exit the CLI and start it again.

Note: There will be no output from this command until the process completely finishes, which may take a few minutes for larger updates.

8.9.3 Updating via the shell

TNSR can also be updated from the command line using the host OS package management commands, in this case, yum:

```
$ sudo yum clean all
$ sudo yum clean expire-cache
$ sudo yum -y upgrade
```

Update Script

The following shell script may be used to keep TNSR and CentOS updated. In addition to the updates it also makes a local backup before performing the update.

Listing 1: Download: updatetnsr.sh

```
1  #!/bin/sh
2
3  # Stop existing services
4  sudo systemctl stop strongswan frr vpp clixon-restconf
5
6  # Time to make the backups
7  mkdir -p ~/tnsr-config-backup
8  sudo cp -p /var/tnsr/running_db ~/tnsr-config-backup/running_db-`date +%Y%m%d%H%M%S`.
   ↪.xml
9  sudo cp -p /var/tnsr/startup_db ~/tnsr-config-backup/startup_db-`date +%Y%m%d%H%M%S`.
   ↪.xml
10
11 # Update all RPMs
12 sudo yum clean all
13 sudo yum clean expire-cache
14 sudo yum -y upgrade
15
16 # Ensure services are stopped, in case some automatically started after update.
17 sudo systemctl stop strongswan frr vpp clixon-restconf
18
19 # Run config database upgrade script
20 sudo /usr/bin/tnsr-db-update
21 sudo cp -p /var/tnsr/running_db /var/tnsr/startup_db
22
23 # Start services
24 sudo systemctl start clixon-backend clixon-restconf
```

8.9.4 Updating the Configuration Database

Automatic Configuration Update

If TNSR fails to start due to a change in the configuration, use the included configuration database migration utility, `/usr/bin/tnsr-db-update`, to check for and correct potential problems.

Warning: TNSR must be stopped before attempting to run this script, which modifies the running database at `/var/tnsr/running_db`.

The `/usr/bin/tnsr-db-update` script must be run as root either directly or via `sudo`:

```
$ sudo /usr/bin/tnsr-db-update
```

The utility will make a backup of the configuration before making alterations. This backup is placed in `/var/lib/tnsr/db-backups/tnsr-<version>/running_db-<timestamp>`.

Warning: This utility currently only alters the running configuration database and not the startup database. After running the script and validating its output, copy the running database to the startup database:

```
$ sudo cp -p /var/tnsr/running_db /var/tnsr/startup_db
```

If TNSR still will not start after running this utility, then the configuration database may contain errors which were not able to be corrected by the script.

Manual Configuration Update

Any errors which could not be corrected by the automatic script must be corrected by hand. Alternately, the configuration databases may be erased and recreated from scratch.

To attempt manual corrections, check the system logs after attempting to start TNSR for information about which configuration entries are causing the failure. View the logs with `sudo systemctl status clixon-backend.service`, `sudo journalctl -xelu clixon-backend.service` and `sudo journalctl -xe`.

A log entry for a configuration problem could look like the following example:

```
clixon_backend: startup_failsafe: 297: Database error: Startup failed and no
  Failsafe database found, exiting
clixon_backend: <rpc-reply><rpc-error>
  <error-type>application</error-type>
  <error-tag>unknown-element</error-tag><error-info>
  <bad-element>someinvalidtag</bad-element></error-info>
  <error-severity>error</error-severity>
  <error-message>namespace is: urn:ietf:params:xml:ns:netconf:base:1.0</error-message>
</rpc-error></rpc-reply>
```

To correct such problems, edit the configuration in `/var/tnsr/running_db` (e.g. `sudo vi /var/tnsr/running_db`), erase or adjust the offending tag or configuration section, copy the repaired configuration to `/var/tnsr/startup_db`, and attempt to start TNSR again. Repeat until no errors are reported and TNSR starts normally.

To erase the configuration database, remove its files from `/var/tnsr`:

```
$ sudo rm /var/tnsr/*_db
```

After removing the configuration and starting TNSR, the TNSR configuration will need to be created again manually from scratch using the CLI or RESTCONF. Open the contents of a configuration backup in a text editor to use as a guide.

8.9.5 Reboot / Update Verification

If a TNSR update installed a new kernel or drivers, a reboot of the router is necessary to activate those changes.

Additionally, after performing a TNSR update and updating the configuration, administrators may wish to perform a reboot of the router to ensure it starts up correctly with the expected configuration.

Though not every TNSR update will require a reboot, the best practice is to test that the router performs as expected at startup during an upgrade maintenance window.

8.9.6 Update Troubleshooting

If the TNSR CLI method does not work, use the shell method instead.

If either method prints an error referring to a broken package database, recover it as follows:

```
$ mkdir -p ~/tmp/  
$ sudo mv /var/lib/rpm/__db* ~/tmp/  
$ sudo rpm --rebuilddb  
$ sudo yum clean all
```

8.9.7 Upgrading by Redeploying TNSR

Rather than performing an in-place update of a TNSR installation, administrators may instead choose to deploy a fresh instance of TNSR using the new version. This practice is typical of environments such as cloud providers or virtual machines, but may be performed for ISO installations and others as well.

In those cases, follow this general procedure:

- Take a backup of the configuration
- Deploy a new instance of TNSR using the installation instructions for the chosen platform
- Restore the configuration
- Update the configuration (*Updating the Configuration Database*)

See *Configuration Backups* for details on saving and restoring configuration backups, and review *Updating the Configuration Database* for important information about updating the configuration for a new version of TNSR.

INTERFACES

An interface must exist in TNSR before it is available for configuration. For hardware interfaces this is handled by the procedure in *Setup Interfaces*. To create additional types of interfaces, see *Types of Interfaces* later in this chapter.

Once interfaces are present in TNSR, they can be configured to perform routing and other related tasks.

See also:

For information on interface status, see *Monitoring Interfaces*.

9.1 Locate Interfaces

The next step is to decide the purpose for which TNSR will use each interface.

First, look at the list of interfaces:

```
tnsr# show interface
Interface: GigabitEthernet0/14/1
[...]
Interface: GigabitEthernet0/14/2
[...]
Interface: local0
[...]
```

In the above shortened output, there are two viable interfaces, GigabitEthernet0/14/1 and GigabitEthernet0/14/2. These can be used for any purpose, so map them as needed for the design of the network for which TNSR will be routing.

The example configuration for this network is:

Table 1: Example Configuration

Interface	Function	IP Address	Gateway
GigabitEthernet0/14/1	WAN	203.0.113.2/24 2001:db8:0:2::2/64	203.0.113.1 2001:db8:0:2::1
GigabitEthernet0/14/2	LAN	10.2.0.1/24 2001:db8:1::1/64	n/a

Connect the interfaces on the router hardware to the appropriate networks at layer 1 and layer 2, for example by plugging the WAN into an Internet circuit and the LAN into a local switch. If TNSR is plugged into a managed switch, ensure that its ports are configured for the appropriate VLANs.

Tip: These interface names can be set to custom values. See [Customizing Interface Names](#) for details.

9.2 Configure Interfaces

With the configuration data in hand, it is now possible to configure TNSR interfaces for basic IP level connectivity.

From within the TNSR CLI (*Entering the TNSR CLI*), enter configuration mode and setup the interfaces using this example as a guide:

```
tnsr# configure terminal
tnsr(config)# interface GigabitEthernet0/14/1
tnsr(config-interface)# description WAN
tnsr(config-interface)# ip address 203.0.113.2/24
tnsr(config-interface)# ipv6 address 2001:db8:0:2::2/64
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# interface GigabitEthernet0/14/2
tnsr(config-interface)# description LAN
tnsr(config-interface)# ip address 10.2.0.1/24
tnsr(config-interface)# ipv6 address 2001:db8:1::1/64
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# exit
tnsr#
```

In this sample session, both interfaces were configured with an appropriate description for reference purposes, an IP address/subnet mask, and then placed into an enabled state.

If other hosts are present and active on the connected interfaces, it will now be possible to ping to/from TNSR to these networks.

Tip: After making changes, don't forget to save them to ensure they persist for the next startup by issuing the `configuration copy running startup` command from within `config` mode. See [Saving the Configuration](#) for more information.

9.2.1 Interface Command

The `interface` command can configure existing interfaces and create new interfaces.

Configure an existing interface:

```
tnsr(config)# interface <name>
tnsr(config-interface)#
```

This command enters `config-interface` mode

Note: The maximum interface name length is 63 characters.

Create a new interface:

```
tnsr(config)# interface <type> <options>
```

The mode entered by this command depends upon the type of interface it creates. For more information on interface types and how to configure them, see *Types of Interfaces*.

Print a list of available interfaces and types:

```
tnsr(config)# interface ?
```

9.2.2 Interface Configuration Options

The following commands are available when configuring an interface (config-interface mode):

access-list (input|output) acl <acl-name> Access Control Lists which apply to packets on this interface in the given direction (*Standard ACLs*).

access-list macip <macip-name> MACIP Access Control Lists which apply to packets on this interface (*MACIP ACLs*).

bond <id> Set this interface as a part of the given bonding group (*Bonding Interfaces*).

bridge domain <id> Set this interface as a member of the given bridge domain (*Bridge Interfaces*).

description Set the interface description.

detailed-stats (enable|disable) Enable or disable the collection of detailed packet statistics which individually track received and transmitted unicast, multicast, and broadcast packets. Disabled by default. Disabling these counters for an interface will not clear the values, it only stops new data collection.

dhcp client [ipv4] Configures this interface to obtain its IPv4 address using Dynamic Host Configuration Protocol.

dhcp client ipv4 hostname <host-name> Sets the hostname sent with DHCP client requests.

disable Disable interface administratively.

enable Enable interface administratively.

ip address <ip-address> Sets the IPv4 address for this interface. May be repeated to add multiple addresses to an interface.

Note: TNSR 19.08 and later support multiple IP addresses in the same prefix. Older versions only allowed a single address per prefix.

ip nat (inside|outside|none) Configures this interface to be an inside or outside NAT interface (*Network Address Translation*). To stop an interface from participating in NAT, use either `no ip nat` or `ip nat none`.

ip reassembly enable Enables *IP Reassembly* for IPv4.

ip reassembly type (full|virtual) Sets the *type of IP Reassembly* to perform on this interface for IPv4 fragments.

ip route-table <route-table-name> Configures a specific IPv4 route table to be used for traffic exiting this interface.

ipv6 address <ip6-address> Sets the IPv6 address for this interface. May be repeated to add multiple addresses to an interface.

Note: TNSR 19.08 and later support multiple IP addresses in the same prefix. Older versions only allowed a single address per prefix.

ipv6 reassembly enable Enables *IP Reassembly* for IPv6.

ipv6 reassembly type (full|virtual) Sets the *type of IP Reassembly* to perform on this interface for IPv6 fragments.

ipv6 route-table <route-table-name> Configures a specific IPv6 route table to be used for traffic exiting this interface.

lldp LLDP options for this interface (*Link Layer Discovery Protocol*).

mac-address Configures an alternative MAC address for this interface.

Warning: Changing the MAC address on an active interface will result in unpredictable behavior. Packets already in transit addressed to the old MAC will be dropped, and it may take some time for other hosts and equipment on directly connected networks to update their ARP tables with the new MAC address.

The best practice is to set an interface administratively down (`disable`) before changing the MAC address, and then enable it again afterward.

map MAP-E/T options for this interface (*MAP (Mapping of Address and Port)*).

mtu <size> Sets the interface L2 Maximum Transmission Unit (MTU) size, in bytes.

rx-queue <queue_num> cpu <core-id> Pin a specific receive queue for this interface to a specific CPU core. Both the queue number and core ID must be valid and within range for the configured number of queues and cores as set with either `corelist-workers` or `coremask-workers`.

See also:

For more information on configuring interface queue sizes, see *DPDK Configuration*. To configure CPU core usage see *CPU Workers and Affinity*.

Warning: This option requires a list of cores configured for dataplane use by the either the `corelist-workers` or `coremask-workers` methods. RX queue core pinning is incompatible with the `workers` and `skip-list` methods of defining CPU cores available for use by the dataplane.

vlan tag-rewrite disable Disable tag rewriting for this interface

vlan tag-rewrite pop-1 Remove one level of VLAN tags from packets on this interface.

vlan tag-rewrite pop-2 Remove two level of VLAN tags from packets on this interface.

vlan tag-rewrite push-1 (dot1ad|dot1q) <tag 1> Add a new layer of VLAN tagging to frames on this interface using the provided VLAN tag.

vlan tag-rewrite push-2 (dot1adldot1q) <tag 1> <tag 2> Add two new layers of VLAN tagging to frames on this interface using the provided VLAN tags.

vlan tag-rewrite translate-1-1 (dot1adldot1q) <tag 1> Replace one layer of VLAN tags with the a different VLAN ID.

vlan tag-rewrite translate-1-2 (dot1adldot1q) <tag 1> <tag 2> Replace one layer of VLAN tags with two layers of tagging using the provided VLAN IDs.

vlan tag-rewrite translate-2-1 (dot1adldot1q) <tag 1> Replace two layers of VLAN tags with one layer of tagging using the provided VLAN ID.

vlan tag-rewrite translate-2-2 (dot1adldot1q) <tag 1> <tag 2> Replace two layers of VLAN tags with two different layers of tagging using the provided VLAN IDs.

9.2.3 DHCP Client Example

The previous example was for a static IP address deployment.

To configure a TNSR interface to obtain its IP address via DHCP as a client, follow this example instead:

```
tnsr# configure terminal
tnsr(config)# interface GigabitEthernet3/0/0
tnsr(config-interface)# dhcp client ipv4
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# exit
```

9.3 Types of Interfaces

Regular Interfaces Typically these are hardware interfaces on the host, or virtualized by the hypervisor in a virtual machine environment. These are made available to TNSR through VPP, as described in *Setup Interfaces*.

VLAN Subinterfaces VLAN interfaces are configured on top of regular interfaces. They send and receive traffic tagged with 802.1q VLAN identifiers, allowing multiple discrete networks to be used when connected to a managed switch performing VLAN trunking or tagging.

memif Shared memory packet interfaces (**memif**) are virtual interfaces which connect between TNSR and other applications on the same host.

tap Virtual network TAP interfaces which are available for use by host applications.

ipip Interfaces created and used by *IPsec* tunnels.

Loopback Local loopback interfaces used for a variety of reasons, including management and routing so that the address on the interface is always available, no matter the status of a physical interface.

GRE Generic Routing Encapsulation, an unencrypted tunneling interface which can be used to route traffic to remote hosts over a virtual point-to-point interface connection.

SPAN Switch Port Analyzer, copies packets from one interface to another for traffic analysis.

Bond Bonded interfaces, aggregate links to switches or other devices employing a load balancing or failover protocol such as LACP.

Bridge Bridges connect interfaces together bidirectionally, linking the networks on bridge members together into a single bridge domain. The net effect is similar to the members being connected to the same layer 2 or switch.

VXLAN Interfaces Virtual Extensible LAN (VXLAN) is a similar concept to VLANs, but it encapsulates Layer 2 traffic in UDP, which can be transported across other IP networks. This enables L2 connectivity between physically separated networks in a scalable fashion.

Host Interfaces Host interfaces exist outside TNSR, in the operating system. These are used primarily for host OS management.

9.3.1 VLAN Subinterfaces

VLANs enable a device to carry multiple discrete broadcast domains, allowing a single switch to function as if it were multiple switches. VLANs are commonly used for network segmentation in the same way that multiple switches can be used: To place hosts on a specific segment, isolated from other segments. Where trunking is employed between switches, devices on the same segment need not reside on the same switch. Devices that support trunking can also communicate on multiple VLANs through a single physical port.

TNSR supports VLANs primarily through subinterfaces, though a variety of VLAN tag rewriting options are available directly on interfaces (*Configure Interfaces*). Using subinterfaces, TNSR can send and receive VLAN tagged traffic on one or more interfaces. The device to which TNSR is connected must also tag traffic in the same way as TNSR.

TNSR also supports multiple levels of VLAN tagged subinterfaces, commonly known as QinQ or 802.1ad. This is used to transport multiple VLANs inside another VLAN-tagged outer frame. Intermediate equipment only sees the outer tag, and the receiving end can pop off the outer tag and use the multiple networks inside independently as if it had a direct layer 2 connection to those networks. In this way, providers can isolate multiple tenants on the same equipment, allowing each tenant to use whichever VLAN tags they require, or achieve other goals such as using greater than the default limit of 4096 VLANs.

VLAN Subinterface Configuration

A few pieces of information are necessary to create a VLAN subinterface (“subif”):

- The parent interface which will carry the tagged traffic, e.g. `GigabitEthernet3/0/0`
- The subinterface ID number, which is a positive integer that uniquely identifies this subif on the parent interface. It is commonly set to the same value as the VLAN tag
- The VLAN tag used by the subif to tag outgoing traffic, and to use for identifying incoming traffic bound for this subif. This is an integer in the range 1–4095, inclusive. This VLAN must also be tagged on the corresponding switch configuration for the port used by the parent interface.

Creating a VLAN Subinterface

The `interface subif <parent> <subinterface id>` command creates a new subif object with the given identifier. This command enters `config-subif` mode. That mode contains the following commands:

default Default subinterface, will match any traffic that does not match another subinterface on the same parent interface.

untagged This subinterface will match frames without any VLAN tags.

exact-match Specifies whether to exactly match the VLAN ID and the number of defined VLAN IDs. When this is not set, frames with more VLAN tags will also be matched. Layer 3/routed interfaces must use `exact-match`, it is optional for unrouted/L2 interfaces.

dot1q (<vlan-id>|any) The VLAN tag to match for this subinterface.

inner-dot1q (<vlan-id>|any) An inner 802.1q VLAN tag for use with QinQ

outer-dot1ad (<vlan-id>|any) An outer 802.1ad VLAN tag for use with QinQ

outer-dot1q (<vlan-id>|any) An outer 802.1q VLAN tag for use with QinQ

vlan <vlan-id> VLAN ID for tag rewriting

Note: Where multiple similar options are present, generally this is for compatibility with other equipment that requires using those specific options. Consult the documentation for the peer device to find out which options it prefers.

After creating the interface, it will be available in TNSR. The name of this interface is composed of the parent interface name and the subif id, joined by a .. For example, TenGigabitEthernet6/0/0.70.

VLAN Subinterface Examples

VLAN Example

First, create a new subif object. In this example, both the subif id and the 802.1q VLAN tag are the same, 70:

```
tnsr(config)# interface subif TenGigabitEthernet6/0/0 70
tnsr(config-subif)# dot1q 70
tnsr(config-subif)# exact-match
tnsr(config-subif)# exit
```

Upon commit, this creates a corresponding subif interface which appears with the parent interface name and the subif id, joined by a ..

```
tnsr(config)# interface TenGigabitEthernet6/0/0.70
tnsr(config-interface)#
```

At this point, it behaves identically to regular interface in that it may have an IP address, routing, and so on.

QinQ Example

This example creates a QinQ subinterface with an inner tag of 100 and an outer tag of 200. The subinterface ID number can be any arbitrary unsigned 32-bit integer, but in this case it makes the purpose more clear to have it match the outer and inner VLAN tags of the subinterface:

```
tnsr(config)# subif GigabitEthernet0/b/0 200100
tnsr(config-subif)# inner-dot1q 100
tnsr(config-subif)# outer-dot1q 200
tnsr(config-subif)# exit
tnsr(config)# exit
```

9.3.2 Shared Memory Packet Interfaces (memif)

A Shared Memory Packet Interface (**memif**) has two components: A socket and an interface. A memif also requires a role, either master or slave. In most TNSR applications, it will be the master and the other endpoint will be a slave. A single socket may only be associated with one role type.

Memif Configuration

Creating a memif Socket

The interface `memif socket` command requires an identifier number and a filename, both of which must be unique to this socket. The full form of the command is: `interface memif socket id <id> filename <socket-filename>`

In this command, the available parameters are:

- id** A required identifier unique to this memif instance. This is an integer in the range 1 . . 4294967294.
- socket-filename** The full path to a socket file used for establishing memif connections. A socket can be used for either master or slave interfaces, but not both. A socket can have more than one master, or it can have more than one slave.

Creating a memif interface

Next, the interface `memif interface <id>` command creates a memif object. This command requires its own interface identifier, and it must be tied to the socket using the same ID from the previous command.

This command enters `config-memif` mode, where the following commands are available:

- socket-id <id>** The socket ID for the associated memif socket created previously. This value is required.
- buffer-size <size>** The size of the buffer allocated for each ring entry. Default 2048.
- mac-address <mac>** MAC address for the memif interface.
- mode <mode>** Sets the mode for the memif interface. Mode must be one of:
 - ethernet** Ethernet (L2) mode.

Note: When `ethernet` mode is active and a `mac-address` is not set, TNSR will generate a random MAC for the interface.

- ip** IP (L3) mode.
- punt/inject** Reserved for future use. Not yet implemented.
- ring-size <size>** Number of entries in receive and transmit rings. Value is 8 . . 32 and is used as a power of 2. Default value is 10 for 1024 (2^{10}) entries.
- role <role> [<options>]** Sets the role of the memif interface. The default role is `master` and this is the most common role for TNSR. The following modes and options are available:
 - master** Master role. The master does not expose its memory to the slave peer.
 - slave [(rx-queues/tx-queues) <num-queues>]** Slave role. Allocates and shares memory with the master to transfer data. When operating in slave mode, the number of receive or transmit queues may be set as an option:
 - rx-queues <n-rx-qs>** Number of receive queues. May be between 1 . . 255.
 - tx-queues <n-tx-qs>** Number of transmit queues. May be between 1 . . 255.
- secret <sec-str>** A quoted secret string, up to 24 characters.

After creating the interface, it will be available in TNSR. The name of this interface is composed of the socket ID and the interface ID: `interface memif<socket id>/<interface id>`.

Memif Example

First, create a socket with an ID of 23, using a socket file of `/tmp/memif23.sock`:

```
tnsr(config)# interface memif socket id 23 filename /tmp/memif23.sock
```

Next, run commands to create a memif interface with an interface ID of 100 taking on the role master on the socket created previously:

```
tnsr(config)# interface memif interface 100
tnsr(config-memif)# socket-id 23
tnsr(config-memif)# role master
tnsr(config-memif)# exit
```

Now the interface will be available to TNSR. In this example with a socket ID of 23 and an interface ID of 100, the full interface name is `memif23/100`.

Memif status

For a list of all current memif entries, along with their names and configuration, use the `show interface memif` command:

```
tnsr# show interface memif

Socket Id  Filename
-----
0          /run/vpp/memif.sock
23         /tmp/memif23.sock

memif id: 100
  Memif name: memif23/100
  Interface: memif23/100
  Role: master
  Mode: ethernet
  MAC address: 02:fe:8c:e5:ce:06
  Socket id: 23
  Ring size: 0
  Buffer size: 0
  Admin up: false
  Link up: false
```

9.3.3 Tap Interfaces

Virtual network tap interfaces give daemons and clients in the host operating system access to send and receive network traffic through TNSR to other networks. A tap interface can carry layer 2 and layer 3 frames between the host OS and TNSR, and be a bridge member.

Tap Configuration

The `interface tap <name>` command creates a tap object with the given name. This name is also used to create the tap interface in the host OS. For example, if a tap object was created with `interface tap mytap`, then the interface in the host OS is named `mytap`.

This command enters `config-tap` mode, which contains the following commands:

instance <instance> Required instance identifier for the tap interface. A tap interface appears in TNSR using the `tap` prefix followed by the chosen identifier number. For example, with an identifier number of 1, the TNSR interface will be `tap1`.

mac-address <mac> The MAC address for the TNSR side of the tap interface.

(rx-ring-size|tx-ring-size) <size> Configures the receive (rx) or transmit (tx) ring buffer size.

Note: Default ring size is 256. The value must be a power of 2 and must be less than or equal to 32768.

host bridge <bridge-name> Configure the tap as part of a host bridge.

Note: A tap object cannot have both an IP address and a bridge name set.

host (ipv4|ipv6) gateway <ip-addr> Configure a gateway for the host tap interface.

host (ipv4|ipv6) prefix <ip-addr> Configures the host IPv4 or IPv6 address for the tap interface.

host mac-address <mac> The MAC address for the host side of the tap interface.

host namespace <ns> Configure a namespace inside which the tap will be created on the host.

TAP Examples

Example tap Interface

The following commands create a tap object named `mytap` with an instance id of 1:

```
tnsr(config)# interface tap mytap
tnsr(config-tap)# instance 1
```

At this point, the TNSR and host OS interfaces exist but contain no configuration:

In TNSR:

```
tnsr# show interface tap1
Interface: tap1
  Admin status: down
  Link up, unknown, unknown duplex
  Link MTU: 9216 bytes
  MAC address: 02:fe:77:d9:be:1e
  IPv4 Route Table: ipv4-VRF:0
  IPv6 Route Table: ipv6-VRF:0
```

In the host OS:

```
$ ip address show mytap
300: mytap: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UNKNOWN
↪group
default qlen 1000
    link/ether 42:5a:f0:6f:d9:77 brd ff:ff:ff:ff:ff:ff
    inet6 fe80::405a:f0ff:fe6f:d977/64 scope link
        valid_lft forever preferred_lft forever
```

Example Tap Interface Addresses

Configuring addresses for tap interfaces depends on the location of the interface.

For the interface visible in TNSR, configure it in the same manner as other TNSR interfaces:

```
tnsr# configure
tnsr(config)# int tap1
tnsr(config-interface)# ip address 10.2.99.2/24
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# exit
tnsr#
```

The MAC address of the tap interface may also be set on the tap object:

```
tnsr# configure
tnsr(config)# interface tap mytap
tnsr(config-tap)# mac-address 02:fe:77:d9:be:ae
tnsr(config-tap)# exit
tnsr(config)# exit
tnsr#
```

The address for the host OS interface is configured by the host command under the tap object instance:

```
tnsr# configure
tnsr(config)# interface tap mytap
tnsr(config-tap)# host ipv4 prefix 10.2.99.1/24
tnsr(config-tap)# exit
tnsr(config)# exit
tnsr#
```

At this point, the interfaces will show the configured addresses:

In TNSR:

```
tnsr# show interface tap1
Interface: tap1
  Admin status: up
  Link up, unknown, unknown duplex
  Link MTU: 9216 bytes
  MAC address: 02:fe:77:d9:be:ae
  IPv4 Route Table: ipv4-VRF:0
  IPv4 addresses:
    10.2.99.2/24
  IPv6 Route Table: ipv6-VRF:0
```

In the host OS:

```
$ ip address show mytap
308: mytap: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UNKNOWN
↳group
default qlen 1000
    link/ether 02:fe:77:d9:be:ae brd ff:ff:ff:ff:ff:ff
    inet 10.2.99.1/24 scope global mytap
        valid_lft forever preferred_lft forever
    inet6 fe80::02fe:77d9:beae/64 scope link
        valid_lft forever preferred_lft forever
```

The host <family> prefix <address> syntax works similarly for IPv6 with an appropriate address.

9.3.4 Loopback Interfaces

Loopback interfaces are internal interfaces available for use in TNSR for routing and other internal traffic handling purposes such as acting as a bridged virtual interface (*Bridge Interfaces*).

Loopback Configuration

Before a loopback interface can be configured, a loopback instance must be created by the `interface loopback <name>` command. This command enters `config-loopback` mode. The loopback must be given a unique name and a positive numeric instance identifier.

The following commands are available in `config-loopback` mode:

instance A required instance identifier. This value is used to generate the loopback interface name in TNSR in the form of `loop<id>`. For example, with an `id` of 1, the loopback interface name is `loop1`.

description A brief text description of this loopback instance.

mac-address An optional MAC address to use for the loopback interface. If omitted, TNSR will generate a MAC in the form of `de:ad:00:00:00:<id>`.

Loopback Example

This example creates a new loopback object named `mgmtloop` with an instance identifier of 1:

```
tnsr(config)# interface loopback mgmtloop
tnsr(config-loopback)# instance 1
tnsr(config-loopback)# exit
```

Upon commit, the new interface will be available for use by TNSR. The interface will be designated `loop<instance id>`, in this case, `loop1`. It can then be configured in the same manner as other interfaces:

```
tnsr(config)# interface loop1
tnsr(config-interface)# ip address 10.25.254.1/24
tnsr(config-interface)# exit
```

9.3.5 GRE Interfaces

A Generic Routing Encapsulation (GRE) interface enables direct routing to a peer that does not need to be directly connected, similar to a VPN tunnel, but without encryption. GRE is frequently combined with an encrypted transport to enable routing or other features not possible with the encrypted transport on its own. GRE interfaces can be combined with dynamic routing protocols such as BGP, or use static routing.

GRE Configuration

To create a GRE object, TNSR requires an object name, positive integer instance ID, source IP address, and destination IP address. The first step is to run the `gre <object-name>` command, which enters `config-gre` mode. Inside `config-gre` mode, the following commands are available:

instance <id> Required instance identifier. This value is used to generate the GRE interface name in TNSR in the form of `gre<id>`. For example, with an `id` of 1, the GRE interface name is `gre1`.

source <ip-address> Required IP address on TNSR to use as a source for GRE traffic associated with this instance. Can be an IPv4 or IPv6 address.

destination <ip-address> Required IP address of the remote GRE peer, which is the destination for GRE traffic associated with this instance. Can be an IPv4 or IPv6 address, but the address family must match that of the `source` IP address.

encapsulation route-table <route-table> This option controls which route table is used by the GRE object, for traffic utilizing the GRE interface. The default behavior is to use the default routing table.

tunnel-type <type> TNSR supports multiple GRE tunnel types, where `<type>` is one of the following:

l3 Layer 3 encapsulation, the default type of GRE tunnel, which can carry layer 3 IP traffic and above.

erspan session-id <id> Encapsulated Remote Switched Port Analyzer (ERSPAN). This requires a session ID number, which is an integer in the range 0..1023. When combined with *Switch Port Analyzer (SPAN) Interfaces*, ERSPAN can deliver copies of local packets to a remote host for inspection. Explained in detail in *GRE ERSPAN Example Use Case*.

teb Transparent Ethernet Bridging (TEB)

GRE Examples

This example creates a new GRE object named `test1`, with an instance `id` of 1, and the source and destination addresses shown:

```
tnsr(config)# gre test1
tnsr(config-gre)# instance 1
tnsr(config-gre)# source 203.0.113.2
tnsr(config-gre)# destination 203.0.113.25
tnsr(config-gre)# exit
```

Upon commit, the new GRE interface will be available for use by TNSR. The name of the GRE interface is `gre<instance id>`, which in this case results in `gre1`. The GRE interface can then be configured similar to other interfaces (*Configure Interfaces*):

```
tnsr(config)# interface gre1
tnsr(config-interface)# ip address 10.2.123.1/30
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# exit
```

See also:

For an example ERSPAN configuration, see *GRE ERSPAN Example Use Case*

GRE Status

To view a list of current GRE objects, use `show gre`:


```
tnsr# show gre
```

Name	Instance	Type	Source IP	Dest IP	Encap Rt	Session Id
test1	1	L3	203.0.113.2	203.0.113.25	ipv4-VRF:0	0

This command prints a list of all GRE objects and a summary of their configuration.

9.3.6 Switch Port Analyzer (SPAN) Interfaces

A SPAN interface ties two interfaces together such that packets from one interface (the source) are directly copied to another (the destination). This feature is also known as a “mirror port” on some platforms. SPAN ports are commonly used with IDS/IPS, monitoring systems, and traffic logging/statistical systems. The target interface is typically monitored by a traffic analyzer, such as snort, that receives and processes the packets.

A SPAN port mirrors traffic to another interface which is typically a local receiver. To send SPAN packets to a remote destination, see [GRE ERSPAN Example Use Case](#) which can carry mirrored packets across GRE.

SPAN Configuration

SPAN instances are configured from `config` mode using the `span <source-interface>` command. That command enters `config-span` mode. Inside `config-span` mode, the following commands are available:

onto <destination-interface> <layer> <state> Specifies a destination for SPAN traffic. May be repeated for multiple destinations. This interface may not be the same as the `<source-interface>` given to create the span instance.

The available parameters include:

destination-interface The interface which will receive copies of packets from the source interface. The destination interface can be any interface available to TNSR except for the `<source-interface>` given to create the span instance.

layer Sets the layer above which packet information is forwarded to the destination. Can be one of the following choices:

- hw** Mirror hardware layer packets.
- l2** Mirror Layer 2 packets.

state Can be one of the following choices:

- rx** Enables receive packets
- tx** Enables transmit packets
- both** Enables both transmit and receive packets
- disabled** Disables both transmit and receive

Note: When removing a `span` instance, the state does not need to be present on the command, and will be ignored.

SPAN Example

This example creates a new span that copies all packets sent and received on `GigabitEthernet0/14/0` to `memif1/1`. The packet copies include hardware level information and above.

```
tnsr(config)# span GigabitEthernet0/14/0
tnsr(config-span)# onto memif1/1 hw both
tnsr(config-span)# exit
```

See also:

For an example ERSPAN configuration that combines GRE in ERSPAN mode with a span instance, see [GRE ERSPAN Example Use Case](#).

9.3.7 Bonding Interfaces

TNSR supports bonding multiple interfaces together for link aggregation and/or redundancy. Several bonding methods are supported, including Link Aggregation Control Protocol (LACP, 802.3ad). These types of interfaces may also be called LAG or LAGG on other platforms and switches.

Bond Configuration

A bond instance has two main components on TNSR: The bond itself, and the interfaces which are a member of the bond. Beyond that, the device to which the bonded interfaces connect, typically a switch, must also support the same bonding protocol and it must also have ports with an appropriately matching configuration.

Warning: Bonds may only be created between hardware interfaces. Virtual interfaces such as Tap interfaces, loopback interfaces, subinterfaces, and other bond interfaces cannot be added to a bond.

Creating a bond

The `interface bond <instance>` command in `config` mode enters `config-bond` mode. An instance number, such as 0, must be manually specified to create a new bond interface.

`config-bond` mode contains the following commands:

load-balance {12|123|134} Configures the load balancing hash for the bonded interface. This setting determines how traffic will be balanced between ports. Traffic matching a single source and destination pair for the configured hash value will flow over a single link. Using higher level hashing will balance loads more evenly in the majority of cases, depending on the environment, but requires additional resources to handle.

This `load-balance` configuration is only available in `lacp` and `xor` modes.

This should be set to match the switch configuration for the ports.

- 12** Layer 2 (MAC address) hashing only. Any traffic to/from a specific pair of MAC addresses will flow over a single link. This method is the most common, and may be the only method supported by the other end of the bonded link.

Note: If the bonded interface only transmits traffic to a single peer, such as an upstream gateway, then all traffic will flow over a single link. The bond still has redundancy, but does not take advantage of load balancing.

- 123** Layer 2 (MAC address) and Layer 3 (IP address) hashing. For non-IP traffic, acts the same as 12.

- 134** Layer 3 (IP address) and Layer 4 (Port, when available) hashing. If no port information is present (or for fragments), acts the same as 123, and for non-IP traffic, acts the same as 12.

mode (round-robin|active-backup|xor|broadcast|lacp)

round-robin Load balances packets across all bonded interfaces by sending a packet out each interface sequentially. This does not require any cooperation from the peer, but can potentially lead to packets arriving at the peer out of order. This can only influence outgoing traffic, the behavior of return traffic is up to the peer.

active-backup Provides only redundancy. Uses a single interface of the bond, and will switch to another if the first interface fails. The switch can only see the MAC address of the active port.

xor Provides hashed load balancing of packet transmission. The transmit behavior is controlled by the `load-balance` option discussed previously. This mode is a step up from `round-robin`, but the behavior of return traffic is still up to the peer.

broadcast Provides only link redundancy by transmitting all packets on all links.

lacp Provides dynamic load balancing and redundancy using Link Aggregation Control Protocol (LACP, 802.3ad). In this mode, TNSR will negotiate an LACP link with an appropriately-configured switch, and monitors the links. This method is the most flexible and reliable, but requires active cooperation from a switch or suitable peer. The load balancing behavior can be controlled with the `load-balance` command discussed previously.

mac-address <mac-address> Optionally specifies a manually-configured MAC address to be used by all members of the bond, except in `active-backup` mode in which case it is only used by the active link.

Bond Interface Settings

Additionally, from within `config-interface` on an Ethernet interface, the following commands are available:

bond <instance> [long-timeout] [passive]

instance The instance ID of the bond to which this interface will belong.

long-timeout Uses a 90-second timeout instead of the default timeout of 3 seconds when monitoring bonding peers, such as with LACP.

passive This interface will be a member of the bond but will not initiate LACP negotiations.

Bond Example

This example sets up a basic LACP bond between two interfaces. The first step is to create the bond instance:

```
tnsr(config)# interface bond 0
tnsr(config-bond)# load-balance 12
tnsr(config-bond)# mode lacp
tnsr(config-bond)# mac-address 00:08:a2:09:95:99
tnsr(config-bond)# exit
```

Next, decided which TNSR interfaces will be members of the bond, and configure them to be a part of the bond instance. In this case, the example uses `GigabitEthernet0/14/2` and `GigabitEthernet0/14/3`:

```
tnsr(config)# int GigabitEthernet0/14/2
tnsr(config-interface)# bond 0
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# int GigabitEthernet0/14/3
tnsr(config-interface)# bond 0
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# exit
```

With that complete, TNSR will now have a new interface, BondEthernet0:

```
Interface: BondEthernet0
  Admin status: down
  Link up, unknown, unknown duplex
  Link MTU: 9216 bytes
  MAC address: 00:08:a2:09:95:99
  IPv4 Route Table: ipv4-VRF:0
  IPv6 Route Table: ipv6-VRF:0
  Slave interfaces:
    GigabitEthernet0/14/2
    GigabitEthernet0/14/3
  counters:
    received: 0 bytes, 0 packets, 0 errors
    transmitted: 0 bytes, 0 packets, 0 errors
    0 drops, 0 punts, 0 rx miss, 0 rx no buffer
```

Looking at the interfaces that are members of the bond, the BondEthernet0 membership is also reflected there:

```
Interface: GigabitEthernet0/14/2
  Admin status: up
  Link up, unknown, full duplex
  Link MTU: 9206 bytes
  MAC address: 00:08:a2:09:95:99
  IPv4 Route Table: ipv4-VRF:0
  IPv6 Route Table: ipv6-VRF:0
  Bond interface: BondEthernet0
  counters:
    received: 52575 bytes, 163 packets, 0 errors
    transmitted: 992 bytes, 8 packets, 19 errors
    31 drops, 0 punts, 0 rx miss, 0 rx no buffer

Interface: GigabitEthernet0/14/3
  Admin status: up
  Link up, unknown, full duplex
  Link MTU: 9206 bytes
  MAC address: 00:08:a2:09:95:99
  IPv4 Route Table: ipv4-VRF:0
  IPv6 Route Table: ipv6-VRF:0
  Bond interface: BondEthernet0
  counters:
    received: 4006 bytes, 37 packets, 0 errors
    transmitted: 620 bytes, 5 packets, 13 errors
    20 drops, 0 punts, 0 rx miss, 0 rx no buffer
```

A configuration can now be applied to BondEthernet0:

```
tnsr(config)# interface BondEthernet0
tnsr(config-interface)# ip address 10.2.3.1/24
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# exit
```

Finally, look at the completed interface configuration:

```
tnsr# show interface BondEthernet0

Interface: BondEthernet0
  Admin status: up
  Link up, unknown, unknown duplex
  Link MTU: 9216 bytes
  MAC address: 00:08:a2:09:95:99
  IPv4 Route Table: ipv4-VRF:0
  IPv4 addresses:
    10.2.3.1/24
  IPv6 Route Table: ipv6-VRF:0
  Slave interfaces:
    GigabitEthernet0/14/2
    GigabitEthernet0/14/3
  counters:
    received: 0 bytes, 0 packets, 0 errors
    transmitted: 806 bytes, 9 packets, 0 errors
    2366 drops, 0 punts, 0 rx miss, 9 rx no buffer
```

For information on the LACP state, use `show interface lacp`:

```
tnsr# show interface lacp
Interface name: GigabitEthernet0/14/2
  Bond name: BondEthernet0
  RX-state: CURRENT
  TX-state: TRANSMIT
  MUX-state: COLLECTING_DISTRIBUTING
  PTX-state: PERIODIC_TX

Interface name: GigabitEthernet0/14/3
  Bond name: BondEthernet0
  RX-state: CURRENT
  TX-state: TRANSMIT
  MUX-state: COLLECTING_DISTRIBUTING
  PTX-state: PERIODIC_TX
```

Bond Status

To view the bond configuration, use `show interface bond`. This will show the configured bond parameters and other information that does not appear on the interface output:

```
tnsr# show interface bond
Interface name: BondEthernet0
  Mode: lacp
  Load balance: 12
  Active slaves: 2
  Slaves: 2
  Slave interfaces:
```

(continues on next page)

(continued from previous page)

```
GigabitEthernet0/14/2
GigabitEthernet0/14/3
```

To view the bonding status of all interfaces, use `show interface bonding`:

```
tnsr# show interface bonding

Interface: BondEthernet0
  Admin status: up
  Slave interfaces:
    GigabitEthernet0/14/2
    GigabitEthernet0/14/3

Interface: GigabitEthernet0/14/0
  Description: Uplink
  Admin status: up

Interface: GigabitEthernet0/14/1
  Admin status: down

Interface: GigabitEthernet0/14/2
  Admin status: up
  Bond interface: BondEthernet0

Interface: GigabitEthernet0/14/3
  Admin status: up
  Bond interface: BondEthernet0

Interface: GigabitEthernet3/0/0
  Description: Local Network
  Admin status: up
```

To view the LACP status, use `show interface lacp [interface name]`:

```
tnsr# show interface lacp

Interface name: GigabitEthernet0/14/2
  Bond name: BondEthernet0
  RX-state: CURRENT
  TX-state: TRANSMIT
  MUX-state: COLLECTING_DISTRIBUTING
  PTX-state: PERIODIC_TX

Interface name: GigabitEthernet0/14/3
  Bond name: BondEthernet0
  RX-state: CURRENT
  TX-state: TRANSMIT
  MUX-state: COLLECTING_DISTRIBUTING
  PTX-state: PERIODIC_TX
```

9.3.8 Bridge Interfaces

Bridges connect multiple interfaces together bidirectionally, linking the networks on bridge members together into a single bridge domain. The net effect is similar to the members being connected to the same layer 2 or switch.

This is commonly used to connect interfaces across different types of links, such as Ethernet to VXLAN. Another common use is to enable filtering between two segments of the same network. It could also be used to allow individual

ports on TNSR to act in a manner similar to a switch, but unless filtering is required between the ports, this use case is not generally desirable.

Warning: Bridges connect together multiple layer 2 networks into a single larger network, thus it is easy to unintentionally create a layer 2 loop if two bridge members are already connected to the same layer 2. For example, the same switch and VLAN.

There are two components to a bridge: The bridge itself, and the interfaces which are members of the bridge.

Bridge Configuration

Creating a Bridge

A bridge is created by the `interface bridge domain <bdi>` command, available in `config` mode. This command enters `config-bridge` mode where the following options are available:

arp entry ip <ip-addr> mac <mac-addr> Configures a static ARP entry on the bridge. Entries present will be used directly, rather than having TNSR perform an ARP request flooded on all bridge ports to locate the target. Additionally, when a bridge is not set to learn MACs, these entries must be created manually to allow devices to communicate across the bridge.

arp term Boolean value that when present enables ARP termination on this bridge. When enabled, TNSR will terminate and respond to ARP requests on the bridge. Disabled by default.

flood Boolean value that when present enables Layer 2 flooding. When TNSR cannot locate the interface where a request should be directed on the bridge, it is flooded to all ports.

forward Boolean value that when present enables Layer 2 unicast forwarding. Allows unicast traffic to be forwarded across the bridge.

learn When present, enables Layer 2 learning on the bridge.

mac-age <minutes> When set, enables MAC aging on the bridge using the specified aging time.

uu-flood When present, enables Layer 2 unknown unicast flooding.

Warning: At least one of `flood`, `forward`, `learn`, or `uu-flood` **must** be enabled when creating a bridge for it to be valid.

Bridge Interface Settings

To add an interface to a bridge as a member, the following settings are available from within `config-interface` mode:

```
interface bridge domain <domain-id> [bvi] [shg <n>]
```

domain id Bridge Domain ID, corresponding to the ID given when creating the bridge interface previously.

bvi Boolean value that when present indicates that this is a Bridged Virtual Interface (BVI). A bridge connects multiple interfaces together but it does not connect them to TNSR. A BVI interface, typically a loopback, allows TNSR to participate in the bridge for routing and other purposes.

An L3 packet routed to the BVI will have L2 encapsulation added and then is handed off to the bridge domain. Once on the bridge domain, the packet may be flooded to all bridge member ports

or sent directly if the destination is known or static. A packet arriving from the bridge domain to a BVI will be routed as usual.

Note: A bridge domain may only contain one BVI member.

shg <n> A Split Horizon Group identifier, used with VXLAN interfaces. This number must be non-zero and the same number must be used on each VXLAN tunnel added to a bridge domain. This prevents packets from looping back across VXLAN interfaces which are meshed between peers.

Bridge Example

This example will setup a bridge between GigabitEthernet3/0/0 and GigabitEthernet0/14/1, joining them into one network. Further, a loopback interface is used to allow TNSR to act as a gateway for clients on these bridged interfaces.

First, create the bridge with the desired set of options:

```
tnsr(config)# interface bridge domain 10
tnsr(config-bridge)# flood
tnsr(config-bridge)# uu-flood
tnsr(config-bridge)# forward
tnsr(config-bridge)# learn
tnsr(config-bridge)# exit
```

Next, add both interfaces to the bridge:

```
tnsr(config)# int GigabitEthernet3/0/0
tnsr(config-interface)# bridge domain 10
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# int GigabitEthernet0/14/1
tnsr(config-interface)# bridge domain 10
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)# interface loopback bridgeloop
tnsr(config-loopback)# instance 1
tnsr(config-loopback)# exit
tnsr(config)# interface loop1
tnsr(config-interface)# ip address 10.25.254.1/24
tnsr(config-interface)# bridge domain 10 bvi
tnsr(config-interface)# enable
tnsr(config-interface)# exit
```

Bridge Status

To view the status of bridges, use the `show interface bridge domain [<id>]` command:

```
tnsr(config)# show interface bridge domain 10
Bridge Domain Id: 10
  flood: true
  uu-flood: true
  forward: true
  learn: true
  arp-term: false
```

(continues on next page)

(continued from previous page)

```

mac-age: 0
BVI IF: loop1
Domain Interface Members
  IF: GigabitEthernet0/14/1    SHG: 0
  IF: GigabitEthernet3/0/0    SHG: 0
  IF: local0                  SHG: 0
  IF: loop1                   SHG: 0
ARP Table Entries

```

If the `id` value is omitted, TNSR will print the status of all bridges.

9.3.9 VXLAN Interfaces

Virtual Extensible LAN, or VXLAN, interfaces can be used to encapsulate Layer 2 frames inside UDP, carrying traffic for multiple L2 networks across Layer 3 connections such as between routed areas of a datacenter, leased lines, or VPNs.

VXLAN tunnels are commonly used to bypass limitations of traditional VLANs on multi-tenant networks and other areas that require large scale L2 connectivity without direct connections.

There are two main components to a VXLAN tunnel: The VXLAN tunnel itself, and the bridge domain used to terminate the tunneled traffic to another local interface.

VXLAN Configuration

A new VXLAN tunnel is created with the `vxlan <if-id>` command in `config` mode, which then enters `config-vxlan` mode.

Note: An `<if-id>` is a string which starts with a letter (a-z or A-Z) or underscore followed by letters, digits (0-9), or any of the following allowed characters: `_`, `/`, `.`, and `-`. For VXLAN interfaces, the string may be at most 63 characters long.

In `config-vxlan` mode, the following commands are available:

instance <id> Required `instance` identifier configured on the VXLAN tunnel. Based on this, a new interface will be available in TNSR named `vxlan_tunnel<id>`. For example, with `instance 0` the interface is named `vxlan_tunnel0`.

vni <u24> Required VXLAN Network Identifier

source <ip-addr> Required source IP address on TNSR used to send VXLAN tunnel traffic.

destination <ip-addr> Required destination IP address for the far side of the tunnel. This can be a multicast address, but if it is, then the `multicast` interface must also be defined.

encapsulation route-table <rt-table-name> Routing table used for VXLAN encapsulation.

multicast interface <if-name> Interface used for multicast. Required if the `destination` address is a multicast address. If defined, the `destination` address must be multicast.

Note: The `source` IP address, `destination` IP address and `encapsulation route table` must all be of the same address family, either IPv4 or IPv6.

VXLAN Examples

The following examples demonstrate common ways that VXLAN interfaces can be used on TNSR.

VXLAN Bridging Example

VXLAN-Related Settings

When using VXLAN interfaces in combination with bridging, there are related settings in bridges and interfaces which supplement the settings placed directly on VXLAN interfaces.

In config-bridge mode, the `arp term` command to enable ARP termination is needed for bridges used with VXLAN tunnels.

In config-interface mode, when adding an interface to a bridge, the `shg` (Split Horizon Group) parameter is required for VXLAN tunnels. This number must be non-zero and the same number must be used on each VXLAN tunnel added to a bridge domain. This prevents packets from looping back across VXLAN interfaces which are meshed between peers.

VXLAN Bridge Configuration

First, create the bridge with the desired set of options:

```
tnsr(config)# interface bridge domain 10
tnsr(config-bridge)# arp term
tnsr(config-bridge)# flood
tnsr(config-bridge)# uu-flood
tnsr(config-bridge)# forward
tnsr(config-bridge)# learn
tnsr(config-bridge)# exit
```

Add host interface to bridge domain:

```
tnsr(config)# int GigabitEthernet3/0/0
tnsr(config-interface)# bridge domain 10 shg 1
tnsr(config-interface)# exit
```

Create the VXLAN tunnel:

```
tnsr(config)# vxlan xmpl
tnsr(config-vxlan)# instance 0
tnsr(config-vxlan)# vni 10
tnsr(config-vxlan)# source 203.0.110.2
tnsr(config-vxlan)# destination 203.0.110.25
tnsr(config-vxlan)# exit
```

Add the VXLAN tunnel to bridge domain:

```
tnsr(config)# int vxlan_tunnel0
tnsr(config-interface)# bridge domain 10 shg 1
tnsr(config-interface)# exit
```

VXLAN SPAN Example

VXLAN can be used to transport traffic in a manner similar to GRE, which can be useful in environments incompatible with GRE. For example, this type of setup can be used in place of the [ERSPAN/GRE recipe example](#) for use on Azure which does not allow GRE.

On TNSR, setup a VXLAN tunnel to the remote peer

```
tnsr(config)# vxlan vxlan1
tnsr(config-vxlan)# instance 1
tnsr(config-vxlan)# vni 13
tnsr(config-vxlan)# source 203.0.110.2
tnsr(config-vxlan)# destination 203.0.110.27
tnsr(config-vxlan)# exit
```

Now setup a SPAN on TNSR between a local interface and the newly created VXLAN

```
tnsr(config)# span GigabitEthernet3/0/0
tnsr(config-span)# onto vxlan_tunnell hw both
tnsr(config-span)# exit
```

On the remote peer, which in this example is a CentOS host acting as a VXLAN tunnel endpoint, configure a matching VXLAN interface:

```
$ sudo ip link add vxlan1 type vxlan id 13 dev ens192 remote 203.0.110.2 dstport 4789
$ sudo ip link set dev vxlan1 up
```

VXLAN Status

To view the status of VXLAN tunnels, use the `show vxlan` command:

```
tnsr# show vxlan
Name Instance Source IP Dest IP Encap Rt Decap Node IF Name Mcast IF
--
--
--
xmpl 0 203.0.110.2 203.0.110.25 ipv4-VRF:0 1 vxlan_tunnel0 10
```

9.3.10 Host Interfaces

Host interfaces are interfaces which have not been allocated to the dataplane. As such, these exist separate from other types of TNSR interfaces. As the name implies, they are available for use by the host operating system. These interfaces are primarily used for host OS management.

Host interfaces may be managed from TNSR as described in this section, or using another mechanism in the host OS, such as Network Manager.

Warning: To be used as a host interface, an interface must not be used by the dataplane. To return an interface from dataplane to host control, see [Remove TNSR NIC for Host Use](#).

Host Interface Configuration

To configure a host interface, from `config` mode, use the `host interface <name>` command to enter `config-host-if` mode. The `<name>` parameter is the name of the interface in the host operating system. To see a list of available interfaces, use `show host interface`.

`config-host-if` mode contains the following commands:

description <text> A brief text description of this interface, such as Management.

enable/disable Enables or disables the interface.

ip address <ipv4-prefix> Sets a static IPv4 address and CIDR mask to use on the interface.

ipv6 address <ipv6-prefix> Sets a static IPv6 address and prefix to use on the interface.

mtu <mtu-value> Sets the maximum transmission unit size for the interface.

Host Interface Example

This example configures the host OS interface `enp8s0f1` with an IP address of `10.2.178.2/24` and an MTU of 1500:

```
tnsr# configure
tnsr(config)# host int enp8s0f1
tnsr(config-host-if)# ip address 10.2.178.2/24
tnsr(config-host-if)# mtu 1500
tnsr(config-host-if)# enable
tnsr(config-host-if)# exit
tnsr(config)# exit
```

To confirm that the settings were applied to the interface, use `show host interface`:

```
tnsr# show host interface enp8s0f1
Interface: enp8s0f1
  Link up
  Link MTU: 1500 bytes
  MAC address: 00:90:0b:7a:8a:6a
  IPv4 addresses:
    10.2.178.2/24
```

As additional confirmation, check how the interface looks in the host operating system using a shell command:

```
tnsr# shell ip addr show enp8s0f1
7: enp8s0f1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group_
↪default qlen 1000
    link/ether 00:90:0b:7a:8a:6a brd ff:ff:ff:ff:ff:ff
    inet 10.2.178.2/24 scope global enp8s0f1
        valid_lft forever preferred_lft forever
```

Host Interface Status

The `show host interface (<name>|ipv4|ipv6|link)` command shows the current status of host interfaces. When run without parameters, `show host interface` will print the status of all host interfaces.

The command also supports the following parameters:

<name> The name of an interface. Restricts the output to only the single given interface.

ipv4 Restricts the output to include only interface IPv4 addresses.

ipv6 Restricts the output to include only interface IPv6 addresses.

link Restricts the output to include only interface link status information, including the MTU and MAC address.

Any subset of these parameters may be given in the same command to include the desired information.

STATIC ROUTING

A route is how TNSR decides where to deliver a packet. Each route is comprised of several components, including:

Route Table A discrete collection of routes to be consulted by TNSR or its services.

Destination The network/prefix to which clients or TNSR services will send packets.

Next Hop Address The neighboring router which can accept traffic for the destination network.

Next Hop Interface The interface through which TNSR can reach the neighboring router

10.1 Route Tables

TNSR is able to use multiple discrete route tables but these tables do not offer complete VRF-style isolation. When routing packets, TNSR consults the route tables present on the interface the packet enters (ingress) which match the address family of the packet (IPv4 or IPv6).

If an interface is not configured for a specific route table, TNSR uses the default table. For IPv4, the default routing table is `ipv4-VRF : 0`. For IPv6, the default is `ipv6-VRF : 0`. Custom routing tables may be given arbitrary names.

Warning: VRF is in the name of the default route tables, but TNSR does not offer full virtual routing and forwarding (VRF) features at this time.

Identical routes can have different destination paths in separate route tables, but identical networks **cannot** be directly connected to multiple interfaces.

10.2 Neighbors

For directly connected networks, TNSR will attempt to locate neighboring hosts via Address Resolution Protocol (ARP) for IPv4 or Neighbor Discover Protocol (NDP) for IPv6. In this way, TNSR can discover the hardware MAC address to which a packet will be delivered in these networks.

10.2.1 Static Neighbors

Static neighbor entries can override this dynamic behavior so that a specified IPv4 or IPv6 address is always associated with the same MAC address.

The command to specify a static neighbor takes the following form:

```
tnsr(config)# neighbor <interface> <ip-address> <mac-address>
[no-adj-route-table-entry]
```

The parameters for this command are:

- <interface>** The interface on which this static entry will be placed.
- <ip-address>** The IPv4 or IPv6 address for the static neighbor entry.
- <mac-address>** The MAC address to associate with the given IP address.
- no-adj-route-table-entry** Do not create an adjacency route table entry.

For example, to add a static entry to map 1.2.3.4 to a MAC address of 00:11:22:33:44:55 on the interface GigabitEthernet3/0/0, run this command from config mode:

```
tnsr(config)# neighbor GigabitEthernet3/0/0 1.2.3.4 00:11:22:33:44:55
```

10.2.2 View Neighbors

To see the current table of known IPv4 and IPv6 neighbors, use the `show neighbor [interface <if-name>]` command.

Note: In other products, this information may be referred to as the ARP table or NDP table.

```
tnsr# show neighbor
```

Interface	S/D	IP Address	MAC Address
GigabitEthernet0/14/0	D	203.0.113.1	00:90:0b:37:a3:24
GigabitEthernet0/14/0	D	203.0.113.14	00:0d:b9:33:0f:71
GigabitEthernet3/0/0	S	1.2.3.4	00:11:22:33:44:55
GigabitEthernet3/0/0	D	10.2.0.129	00:0c:29:4c:b3:9b

This output can optionally be filtered by interface name.

The S/D column shows if the entry is static (S) or dynamic (D).

10.3 Viewing Routes

To view the contents of all route tables:

```
tnsr# show route
```

To view the contents of a single route table:

```
tnsr# show route table <table name>
```

For example, to view the default IPv4 route table only, use:

```
tnsr# show route table ipv4-VRF:0
```

10.3.1 Route Flags

In the route display, the `flags` : row may contain the following:

- no flags** If the flags line is empty, this is a normal route with no special actions.
- local** This network is local to TNSR and packets to this destination will not leave the TNSR host.
- drop** Packets matching this route will be dropped by TNSR. Commonly seen with null routes for subnets or for traffic which must not leave a subnet.
- unreachable** Packets matching this route will be dropped by TNSR, and TNSR will send an ICMP “Destination unreachable” message back to the source address.
- prohibit** Packets matching this route will be dropped by TNSR, and TNSR will send an ICMP “Destination administratively prohibited” message back to the source address.

10.3.2 Common Routes

Routing tables on TNSR may include unexpected entries by default or even after adding and configuring interfaces and other services. The following list covers several of these types of routes that may be present and what they mean:

- 0.0.0.0/32 (drop)** Null route to drop traffic with an empty address.
- 0.0.0.0/0 or ::/0** Default route for packets that do not match any other route, such as for Internet hosts or other remote destinations.
- 224.0.0.0/4 (drop)** Multicast that must not be routed.
- 224.0.0.0/24** Local subnet multicast.
- 240.0.0.0/4 (drop)** Reserved network that must not be routed.
- 255.255.255.255/32 (local)** Special broadcast address for networks local to TNSR.
- fe80::/10** IPv6 link local.
- x.x.x.<first>/32 (drop)** Null route for subnet configured on an interface. Last octet will vary depending on subnet size and network address. For example, this is .0 in a /24 subnet.
- x.x.x.<last>/32 (drop)** Broadcast address for subnet configured on an interface. Last octet will vary depending on subnet size and network address. For example, this is .255 in a /24 subnet.
- x.x.x.x/32 (via x.x.x.x, local)** Internal route for an IPv4 address present on a TNSR interface.
- x::x::x/128 (via x::x::x, local)** Internal route for an IPv6 address present on a TNSR interface.

Routes can also be added to the table dynamically by other processes such as via BGP or if an interface is configured as a DHCP client. Check the status or other associated logs for configured features to find the origins of these routes.

10.4 Managing Routes

Routes are entered into TNSR using the `route (ipv4|ipv6) table <name>` command in configuration mode. When using the `route` command for this purpose, the address family and table name must be specified in order to establish the routing context. This command enters `config-route-table` mode. From there, individual routes can be managed.

Inside `config-route-table` mode, the following commands are available:

- description** Sets a description for the route table.

route <destination-prefix> Configures a route to the specified destination network. This enters `config-rttbl-next-hop` mode where the remaining parameters for the route are set.

Tip: For a single address, use a /32 mask for IPv4 or /128 for IPv6.

Inside `config-rttbl-next-hop` mode, the following commands are available:

description Sets a description for this route.

next-hop <hop-id> via <action|gateway> Configures how TNSR will handle traffic to this destination. This may be repeated multiple times with unique *hop-id* values to specify multiple destinations. The following parameters are available to control the route behavior:

hop-id The ID of the next hop. Must be unique between entries in the same route.

via <ip-address> Sets the next hop for this route as an IP address. Additional modifiers are possible for any `via` form using an IP address destination, see [Route modifiers](#).

via <ip-address> <interface> Configures both the IP address and interface for the next hop. This is the most commonly used form for routes. May use modifiers, see [Route modifiers](#).

via <ip-address> next-hop-table <route-table-name> Configures a recursive route lookup using a different route table. May use modifiers, see [Route modifiers](#).

via classify <classify-name> Reserved for future use.

via drop Drops traffic to this destination (null route).

via local The destination is local to TNSR, such as an interface address or loopback.

via null-send-prohibit Packets matching this route will be dropped by TNSR, and TNSR will send an ICMP “Destination administratively prohibited” message back to the source address.

via null-send-unreach Packets matching this route will be dropped by TNSR, and TNSR will send an ICMP “Destination unreachable” message back to the source address.

10.4.1 Route modifiers

For routes set with a next hop using `via <ip-address>`, additional modifiers control how TNSR resolves the route destination.

preference Sets the administrative distance preference. Helps to choose between multiple possible destinations when routing protocols are used. This is only a local value, and a lower value is taken as being more reliable (closer).

weight The weight of routes to the same destination. Acts as a ratio of packets to deliver to each next hop.

Tip: Equal weights will deliver the same amount of traffic to all next hops for this destination prefix, uneven weights will deliver more traffic via the higher weighted connection. If one path has a weight of 1, and the other has a weight of 3, then the first path will receive 25% ($1 / (1+3)$) of the traffic and the other will receive 75% ($3 / (1+3)$).

resolve-via-attached Sets a constraint on recursive route resolution via attached network. The next hop is unknown, but destinations in this prefix may be located via ARP.

resolve-via-host Sets a constraint on recursive route resolution via host. The next hop is known, but the interface is not.

Tip: Multiple modifiers may be used together, but when doing so, `weight` and `preference` must be set first.

Example

IPv4 example:

```
tnsr(config)# route ipv4 table ipv4-VRF:0
tnsr(config-route-table-v4)# route 10.2.10.0/24
tnsr(config-rttbl4-next-hop)# next-hop 0 via 10.2.0.2 GigabitEthernet0/14/2
```

IPv6 Example:

```
tnsr(config)# route ipv6 table ipv6-VRF:0
tnsr(config-route-table-v6)# route fc07:b337:c4f3::/48
tnsr(config-rttbl6-next-hop)# next-hop 0 via 2001:db8:1::2 GigabitEthernet0/14/2
```

Breaking down the examples above, first the route table is specified. Within that context a destination network route is given. The destination network establishes a sub-context for a specific route. From there, the next hop configuration is entered.

Note: When entering a next hop for a route in this way, **both** the IP address of the destination router **and** the interface must be given.

To specify more than one route, exit out of the `next-hop` context so that TNSR is in the correct context for the route table itself, then enter an additional destination and next-hop.

10.5 Default Route

In TNSR, the default route, sometimes called a default gateway, is the gateway of last resort. Meaning, traffic that is not local and does not have any other route specified will be sent using that route. There is no `default` keyword in TNSR; Instead, the special network `0.0.0.0/0` is used for IPv4 and `::/0` is used for IPv6.

In this example, the gateway from *Example Configuration* is added using the WAN interface:

IPv4 Default Route Example:

```
tnsr(config)# route ipv4 table ipv4-VRF:0
tnsr(config-route-table-v4)# route 0.0.0.0/0
tnsr(config-rttbl4-next-hop)# next-hop 0 via 203.0.113.1 GigabitEthernet0/14/1
tnsr(config-rttbl4-next-hop)# exit
tnsr(config-route-table-v4)# exit
```

IPv6 Default Route Example:

```
tnsr(config)# route ipv6 table ipv6-VRF:0
tnsr(config-route-table-v6)# route ::/0
tnsr(config-rttbl6-next-hop)# next-hop 0 via 2001:db8:0:2::1 GigabitEthernet0/14/1
tnsr(config-rttbl6-next-hop)# exit
tnsr(config-route-table-v6)# exit
```

DYNAMIC ROUTING

Dynamic routing refers to routes that are capable of changing, generally due to routing protocols exchanging routing information with neighboring routers.

Unlike static routes, dynamic routing does not require remote network destinations and gateways to be hardcoded in the configuration. Routes and gateways are automatically determined by the protocol instead.

Currently TNSR supports multiple dynamic routing protocols:

Border Gateway Protocol (BGP) BGP routes between autonomous systems, connecting to defined neighbors to exchange routing and path information. BGP supports IPv4 and IPv6.

Open Shortest Path First v2 (OSPF) OSPF is a link-state routing protocol that automatically locates neighboring IPv4 routers within an autonomous system, typically with multicast, and exchanges routing information for networks reachable through each neighbor. OSPF v2 only supports IPv4.

Open Shortest Path First v3 (OSPF6) Similar to OSPF v2, but for IPv6 networks.

Routing Information Protocol (RIP) A routing protocol where each router broadcasts its routing table to peers on connected segments. Simple and widely supported, but not as fast or efficient as other protocols.

Dynamic routing on TNSR is handled by [FRR](#).

11.1 Dynamic Routing Manager

The dynamic routing manager, currently the Zebra daemon from FRR, controls aspects of dynamic routing which are relevant to multiple types of dynamic routing. These include Access Lists, Prefix Lists, and Route Maps. These mechanisms allow for fine-tuning dynamic routing behavior.

11.1.1 Dynamic Routing Manager Configuration

Configuration of the dynamic routing manager itself is performed from within `config-route-dynamic-manager` mode, which is entered as follows:

```
tnsr(config)# route dynamic manager
tnsr(config-route-dynamic-manager)#
```

That mode offers logging and debugging commands, described next.

Logging

The dynamic routing manager daemon can send log messages to a file, via syslog, or both.

log file <filename> [<level>] Instructs the dynamic routing manager daemon to send log messages to the specified file. The optional `level` parameter determines the verbosity of the logged data. See [Log levels](#) for details.

log syslog [<level>] Instructs the dynamic routing manager daemon to send log messages to syslog. The optional `level` parameter determines the verbosity of the logged data. See [Log levels](#) for details.

Log levels

Log levels set the verbosity of the logging recorded by the dynamic routing manager. Each level includes messages from higher priority levels. The default level is `debugging`, which will log as much detail as possible.

Note: Even if the log level is set to `debugging`, actual debugging messages may not appear unless specific debug entries are set. See [Debugging](#) for details.

In order of verbosity, from low to high, the available `level` values are:

- `emergencies`
- `alerts`
- `critical`
- `errors`
- `warnings`
- `notifications`
- `informational`
- `debugging`

For example, if the log level is set to `errors`, then the logs will contain messages with a level of `emergencies`, `alerts`, `critical`, and `errors`, and will exclude the rest.

Debugging

The `debug` command controls which debugging messages will be logged by the dynamic routing manager. These include:

debug events General events.

debug fpm Forwarding Plane Manager events.

debug kernel Kernel messages.

debug kernel msgdump [send|receive] Raw netlink messages, optionally limited to `send` or `receive` messages.

debug nht Next-Hop tracking events

debug packet [send|receive] [detailed] Information about each packet seen by the dynamic routing manager. Optionally limited to `send` or `receive` packets. The `detailed` keyword will log additional information for each packet.

debug rib [detailed] Routing Information Base events, optionally with more `detailed` information.

Note: Debugging messages will only appear in logs if the logs are set to include debugging messages. See [Log levels](#) for details.

11.1.2 Dynamic Routing Access Lists

Access List entries determine if networks are allowed or denied in specific contexts used in various routing daemons. For example, an access list may be used to determine if a route is accepted or rejected, or for limiting routes distributed to neighbors.

The order of entries inside access lists is important, and this order is determined by a sequence number.

Access List Configuration

To create a new access list, use the `route dynamic access-list <name>` command, which enters `config-access-list` mode:

```
tnsr(config)# route dynamic access-list myacl
tnsr(config-access-list)#
```

`config-access-list` mode contains the following commands:

remark <text> A text comment to describe this access list.

sequence <sequence-number> (permit|deny) <ip-prefix> Creates a new rule with the specified sequence number to permit or deny a given prefix.

sequence <sequence-number> The sequence number for this rule, which controls the order in which rules are matched inside this access list. Each rule in an access list must have a unique sequence number. Best practice is to leave gaps in the sequence to allow for adding rules in the future. For example, use 10, 20, 30, rather than 1, 2, 3.

(permit|deny) The action to take for this rule, either `permit` or `deny`.

<ip-prefix> The IP prefix to match for this rule, given in network/prefix notation. For example, `192.168.0.0/16`.

Access List Example

For example, the following ACL would deny `192.168.0.0/16` but permit all other networks:

```
tnsr(config)# route dynamic access-list myacl
tnsr(config-access-list)# sequence 10 deny 192.168.0.0/16
tnsr(config-access-list)# sequence 20 permit 0.0.0.0/0
tnsr(config-access-list)# exit
tnsr(config)#
```

This access list would then be used in another context, such as with a route map, to match routes for anything except `192.168.0.0/16` when taking other actions.

Access List Status

To view access lists, use the `show route dynamic access-list [name]` command. Add the name of an access list to restrict the output to a single access list.

```
tnsr# show route dynamic access-list
```

```
Access List: myacl
```

```
Remark:
```

```
  Seq Action Prefix
  ---
  10  deny   192.168.0.0/16
  20  permit 0.0.0.0/0
```

11.1.3 Dynamic Routing Prefix Lists

Prefix List entries determine parts of networks which can be allowed or denied in specific contexts used in routing daemons. For example, a prefix list may be used to match specific routes in a route map.

The order of entries inside prefix lists is important, and this order is determined by a sequence number.

Prefix List Configuration

To create a new prefix list, use the `route dynamic prefix-list <name>` command, which enters `config-prefix-list` mode:

```
tnsr(config)# route dynamic prefix-list mypl
tnsr(config-prefix-list)#
```

`config-prefix-list` mode contains the following commands:

description <text> A text comment to describe this prefix list.

sequence <sequence-number> (permit|deny) <prefix> [ge <lower-bound>] [le <upper-bound>]

Creates a new rule with the specified sequence number to permit or deny a given prefix. This may optionally be bound by an upper or lower prefix size limit. When no upper or lower bound is set, the prefix will be matched only exactly as given. Setting bounds allows a prefix list to also match more specific routes which are a part of the specified network.

sequence <sequence-number> The sequence number for this rule, which controls the order in which rules are matched inside this prefix list. Each rule in a prefix list must have a unique sequence number. Best practice is to leave gaps in the sequence to allow for adding rules in the future. For example, use 10, 20, 30, rather than 1, 2, 3.

(permit|deny) The action to take for this rule, either permit or deny.

<ip-prefix> The IP prefix to match for this rule, given in network/prefix notation. For example, 192.168.0.0/16.

ge <lower-bound> Sets a lower bound for the prefix length. This must be greater than the prefix length given in <prefix>, and less than or equal to the value of le <upper-bound>, if present.

le <upper-bound> Sets an upper bound for the prefix length. This must be greater than the prefix length given in <prefix>, and greater than or equal to the value of ge <upper-bound>, if present.

Prefix List Examples

For example, the following prefix list will match any of the RFC1918 networks:

```
tnsr(config)# route dynamic prefix-list RFC1918
tnsr(config-prefix-list)# description List of RFC1918 private address space
tnsr(config-prefix-list)# sequence 10 permit 10.0.0.0/8 le 32
tnsr(config-prefix-list)# sequence 20 permit 172.16.0.0/12 le 32
tnsr(config-prefix-list)# sequence 30 permit 192.168.0.0/16 le 32
```

For each of these entries, the prefix list will match based on the bits specified in the prefix. A match will occur for any network included in the specified range. For example, 10.0.0.0/8 le 32 means a route for any smaller network inside 10.0.0.0/8 will also match, so long as the prefix length is less than 32. So 10.2.0.0/16 will also match this entry, as will 10.34.157.82/32. Taken as a whole, this prefix list will match not only the list of RFC1918 networks exactly, but any smaller network wholly contained inside.

As another example, consider this rule instead:

```
tnsr(config-prefix-list)# sequence 10 deny 10.0.0.0/8 ge 24 le 32
```

This matches routes for networks inside of 10.0.0.0/8 with a prefix length greater than or equal to 24 but less than or equal to 32. Meaning it will **not** match larger networks such as 10.2.0.0/16 but it will match more specific networks such as 10.2.56.128/29 anywhere inside the 10.0.0.0/8 address space. This type of rule can be used to exclude small prefixes from being matched by a route map, for example.

Prefix lists are then used in another context, such as with a route map, to match routes any of the specified networks when taking other actions.

Prefix List Status

To view prefix lists, use the `show route dynamic prefix-list [name]` command. Add the name of a prefix list to restrict the output to a single prefix list.

```
tnsr(config)# show route dynamic prefix-list

Prefix Name: RFC1918
Description: List of RFC1918 private address space
  Seq Action Prefix          LE Len GE Len
  ---
  10 permit 10.0.0.0/8      32
  20 permit 172.16.0.0/12   32
  30 permit 192.168.0.0/16 32

Prefix Name: mypl
Description:
  Seq Action Prefix          LE Len GE Len
  ---
  10 deny 192.168.0.0/16
```

11.1.4 Dynamic Routing Route Maps

Route maps are a powerful mechanism which can match or set various values for use by routing daemons, especially BGP. A route map can match based on criteria such as those set by *Dynamic Routing Access Lists* and *Dynamic Routing Prefix Lists*, among others. Route maps can control, for example, whether or not specific routes are accepted

from neighbors, or whether or not specific routes are distributed to neighbors. They can also adjust various properties of routes, which largely depends upon the context in which they are used, such as for BGP or OSPF.

Route Map Configuration

To create a new route map, use the `route dynamic route-map <route-map-name> (permit|deny) sequence <sequence>` command, which enters `config-rt-map` mode:

```
tnsr(config)# route dynamic route-map <route-map-name> (permit|deny) sequence
-><sequence>
tnsr(config-rt-map) #
```

In this command, the parameters are:

<route-map-name> The name of this route map.

(permit|deny) The action taken by this route map.

permit When an entry is matched and permitted, the *Route Map Set Operations* portions of the route map are carried out, if present, and then *Route Map Control Operations* entries, if present, are performed. The route will be allowed unless the control flow ultimately prevents that from happening.

deny When an entry is matched and denied, the route is not allowed.

sequence <sequence> The sequence number of this route map.

The `route dynamic route-map` command may be repeated with different sequence numbers to setup additional rule entries in the same route map.

`config-rt-map` mode offers a variety of commands, which have been broken up into sections.

Route Map General Parameters

description <string> A text description of this route map rule.

Route Map Matching Criteria

match as-path <as-path-name> Match based on *BGP AS Path Access Lists*.

match community <comm-list-name> [exact-match] Match based on *BGP Community Lists*.

match extcommunity <extcomm-list-name> Match based on Extended *BGP Community Lists*.

match interface <if-name> Match based on a specific interface name.

match ip address access-list <access-list-name> Match IPv4 route content based on *Dynamic Routing Access Lists*.

match ip address prefix-list <prefix-list-name> Match IPv4 route content based on *Dynamic Routing Prefix Lists*.

match ip next-hop access-list <access-list-name> Match the next-hop of IPv4 routes based on *Dynamic Routing Access Lists*.

match ip next-hop <ipv4-address> Match the next-hop of IPv4 routes based on IPv4 address.

match ip next-hop prefix-list <prefix-list-name> Match the next-hop of IPv4 routes based on *Dynamic Routing Prefix Lists*.

match ipv6 address access-list <access-list-name> Match IPv6 route content based on *Dynamic Routing Access Lists*.

match ipv6 address prefix-list <prefix-list-name> Match IPv6 route content based on *Dynamic Routing Prefix Lists*.

match large-community <large-comm-list-name> Match based on Large *BGP Community Lists*.

match local-preference <preference-uint32> Match based on configured local preference of a route.

match metric <metric-uint32> Match based on the metric of a route.

match origin (egpligplincomplete) Match based on the origin (source) of a route. It can be one of *egp* (exterior gateway protocols), *igp* (interior gateway protocols), or *incomplete*.

match peer <peer-ip-address> Match based on the IP address of the neighbor associated with a route.

match probability <percent> Match a subset of routes based on the given *percent* value. For example, a value of 60 would match 60% of routes.

match source-protocol <src-protocol> Matches based on the routing protocol for this route (For a list, see *Dynamic Routing Protocol Lists*.)

match tag <value> Match a tag value set by another route map rule. This value is an integer from 1–4294967295.

Route Map Set Operations

set aggregator as <asn> ip address <ipv4-address> Sets the AS of an aggregated route to the specified AS number and its origin to the specified IP address.

set as-path exclude <string-of-as-numbers> Excludes the specified AS numbers from the path of the route.

set as-path prepend <string-of-as-numbers> Prepends the specified AS numbers to the AS path

set as-path prepend last-as <asn> Prepends the specified AS number to the leftmost end of the path.

set atomic-aggregate Sets the BGP “atomic aggregate” attribute for the route. This informs BGP peers that some routing information may not be present due to route aggregation.

set community none Removes information about *BGP Community Lists* from the route.

set community <community-value> [additive] Sets the *BGP community* to the supplied list. The optional *additive* keyword causes the community value to be added to the route without replacing the existing values.

set comm-list <community-list-name> delete Removes specific values from *BGP Community Lists* lists.

set extcommunity rt <extcommunity-list-name> Sets the route target to the given extended community list.

set extcommunity soo <extcommunity-list-name> Sets the site of origin for the route to the given extended community list.

set forwarding-address <ipv6-address> Sets the OSPF forwarding address for this route to the given IPv6 address.

set ip next-hop <ipv4-address> Sets the next-hop for an IPv4 route to this specific address.

- set ip next-hop peer-address** For inbound IPv4 routes received from a neighbor, sets the next-hop to the address of the neighbor. For outgoing routes this is the local address used to establish an adjacency with the neighbor.
- set ip next-hop unchanged** Do not change the next-hop on the route.
- set ipv4 vpn next-hop (<ipv4-address>|<ipv6-address>)** Sets IPv4 VPN next-hop address to the given value.
- set ipv6 next-hop global <ipv6-address>** Sets IPv6 next-hop address to the given globally routable IPv6 address.
- set ipv6 next-hop local <ipv6-address>** Sets IPv6 next-hop address to the given link-local IPv6 address.
- set ipv6 next-hop peer-address** For inbound IPv6 routes received from a neighbor, sets the next-hop to the address of the neighbor. For outgoing routes this is the local address used to establish an adjacency with the neighbor.
- set ipv6 next-hop prefer-global** For inbound routes with both a global and link-local next-hop available, prefer to use the global address.
- set ipv6 vpn next-hop (<ipv4-address>|<ipv6-address>)** Sets IPv6 VPN next-hop address to the given value.
- set large-community none** Removes information about Large *BGP Community Lists* from the route.
- set large-community <large-community-value> [additive]** Sets the Large *BGP community* to the supplied list. The optional *additive* keyword causes the large community value to be added to the route without replacing the existing values.
- set large-comm-list <large-comm-list-name> delete** Removes specific values from Large *BGP Community Lists* lists.
- set local-preference <preference>** Sets the BGP local preference for the route to the supplied value.
- set metric [+]<metric>** Sets the MED value for routes. When this router has multiple links to the same AS, the MED value influences which path the router will prefer. The router will prefer to use links with a lower MED value. Adding a + before the metric value will result in a relative adjustment instead of setting an absolute value.
- set origin (egpligplunknown)** Sets the origin (source) of a route. It can be one of *egp* (exterior gateway protocols), *igp* (interior gateway protocols), or *incomplete*.
- set originator <ipv4-addr>** Sets the originator ID to the supplied address.
- set src <ip-address>** Sets the route source to the supplied address.
- set tag <tag>** Set a tag value to be matched by another route map rule. This value is an integer from 1–4294967295.
- set weight <weight>** Sets the weight of the route to the supplied value. When a remote AS is reachable via multiple paths through other intermediate AS neighbors, the router will prefer to use a higher weight path to reach it.

Route Map Control Operations

- call <rt-map-name>** Will immediately process the named route map. If the called route map returns *deny*, then processing is stopped and the route is denied.
- on-match next** Proceeds to the next rule in the route-map
- on-match goto <sequence>** Skips to the rule with the given sequence number in this route map.

Route Map Examples

This example creates a route map to control which routes will be sent to peers via BGP. The first rule prevents any route from sending if it matches entries in the RFC1918 prefix list. The second rule allows routes that match networks listed in the MY-ROUTES prefix list. This ensures that even if other mechanisms would try to export routes to peers, that no routes to private networks are leaked.

```
tnsr(config)# route dynamic route-map EBGp-OUT deny sequence 10
tnsr(config-route-map)# mat ip address prefix-list RFC1918
tnsr(config-route-map)# exit
tnsr(config)# route dynamic route-map EBGp-OUT permit sequence 30
tnsr(config-route-map)# mat ip address prefix-list MY-ROUTES
tnsr(config-route-map)# exit
```

This route map is to be used with incoming routes from peers. The first rule prevents routes for local networks from being received and processed. The second rule applies attributes to all other received routes.

```
tnsr(config)# route dynamic route-map PEERS-IN deny sequence 10
tnsr(config-route-map)# match ip address prefix-list RFC1918
tnsr(config-route-map)# exit
tnsr(config)# route dynamic route-map PEERS-IN permit sequence 20
tnsr(config-route-map)# set metric 5000
tnsr(config-route-map)# set local-preference 100
tnsr(config-route-map)# set community no-export
tnsr(config-route-map)# exit
```

See also:

For more examples, see the following recipes:

- [Service Provider Route Reflectors and Client for iBGP IPv4](#)
- [TNSR IPsec Hub for pfSense](#)

Route Map Status

To view route maps, use the `show route dynamic route-map [name]` command. Add the name of a route map to restrict the output to a route map.

```
tnsr(config)# show route dynamic route-map
route-map EBGp-OUT deny 10
    match ip address prefix-list RFC1918
route-map EBGp-OUT permit 30
    match ip address prefix-list MY-ROUTES
route-map PEERS-IN deny 10
    match ip address prefix-list RFC1918
route-map PEERS-IN permit 20
    set community no-export
    set local-preference 100
    set metric 5000
```

11.1.5 Dynamic Routing Manager Status

TNSR supports several commands to display information about the dynamic routing manager daemon configuration and its status.

See also:

For more specific dynamic routing daemon status information, see *BGP Status*, *OSPF Status*, and *OSPF6 Status*

Configuration Information

To view the current configuration file for the dynamic routing manager daemon, use `show route dynamic manager`:

```
tnsr# show route dynamic manager
debug zebra events
log file /tmp/zebra-crit.log critical
log syslog warnings
```

To view other individual sections of the configuration:

```
tnsr# show route dynamic access-list [<access-list-name>]
tnsr# show route dynamic prefix-list [<prefix-list-name>]
tnsr# show route dynamic route-map [<route-map-name>]
```

Additional Information

Additional status information can be obtained by using the `vttysh` program outside of TNSR.

The `vttysh` program must be run as root:

```
sudo vttysh
```

The `vttysh` interface offers numerous commands. Of particular interest for BGP status are the following:

show ip route The IP routing table managed by the FRR Zebra daemon, which marks the origin of routes to see which entries were obtained via BGP.

11.2 Border Gateway Protocol

Border Gateway Protocol (BGP) is a dynamic routing protocol used between network hosts. BGP routes between autonomous systems, connecting to defined neighbors to exchange routing information.

BGP can be used for exterior routing (ebgp) or interior routing (ibgp), routing across Internet circuits, private links, or segments of local networks.

11.2.1 BGP Required Information

Before starting, take the time to gather all of the information required to form a BGP adjacency to a neighbor. At a minimum, TNSR will need to know these items:

Local AS Number The autonomous system (AS) number for TNSR. This is typically assigned by an upstream source, an RIR, or mutually agreed upon by internal neighbors.

Local Router ID Typically the highest numbered local address on the firewall. This is also frequently set as the internal or LAN side IP address of a router. It does not matter what this ID is, so long as it is given in IPv4 address notation and does not conflict with any neighbors.

Local Network(s) The list of networks that are advertised over BGP as belonging to the Local AS. For external BGP, this is typically the IP address block allocated by the RIR. For internal BGP, this may be a list of local networks or a summarized block.

Neighbor AS Number The autonomous system number of the neighbor.

Neighbor IP Address The IP address of the neighboring router.

The example in this section uses the following values:

Table 1: Example BGP Configuration

Item	Value
Local AS Number	65002
Local Router ID	10.2.0.1
Local Network(s)	10.2.0.0/16
Neighbor AS Number	65005
Neighbor IP Address	203.0.113.14

Warning: If NAT is active on the same interface acting as a BGP peer, then NAT forwarding must also be enabled. See [NAT Forwarding](#).

11.2.2 BGP Example Configuration

The following example configures a BGP adjacency to a neighbor using the settings from [Example BGP Configuration](#):

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# server 65002
tnsr(config-bgp)# router-id 10.2.0.1
tnsr(config-bgp)# neighbor 203.0.113.14
tnsr(config-bgp-neighbor)# remote-as 65005
tnsr(config-bgp-neighbor)# enable
tnsr(config-bgp-neighbor)# exit
tnsr(config-bgp)# address-family ipv4 unicast
tnsr(config-bgp-ip4uni)# network 10.2.0.0/16
tnsr(config-bgp-ip4uni)# exit
tnsr(config-bgp)# exit
tnsr(config-frr-bgp)# enable
tnsr(config-frr-bgp)# exit
tnsr(config)# service bgp restart
```

BGP Example with Loopback

BGP on TNSR can also be used with loopback interfaces for more advanced routing scenarios. Using a loopback for a BGP update source allows the path to the routing peer to be handled in some other way. It may be static, or it may involve multiple paths to the peer, for example.

This scenario is based on the previous example, but uses a loopback interface for the update source.

Configure Loopback

First, setup the loopback interface and address:

```
tnsr(config)# interface loopback bgploop
tnsr(config-loopback)# instance 1
tnsr(config-loopback)# exit
```

(continues on next page)

(continued from previous page)

```
tnsr(config)# interface loop1
tnsr(config-interface)# ip address 10.5.222.1/32
tnsr(config-interface)# exit
```

Since the loopback is not on an interface, the 10.5.222.1 address must be routed to TNSR somehow. This could be an address in a routed block, or there could be another method of handling routes between the peers.

Route to Peer

Likewise, TNSR must know how to reach the remote peer, 10.5.222.2, which in this case the example also assumes is a loopback address configured in a similar manner. In this example, the peer is reachable at 203.0.113.14 which is in a network directly connected to TenGigabitEthernet6/0/0. For simplicity, this will only be a static route:

```
tnsr(config)# route ipv4 table ipv4-VRF:0
tnsr(config-route-table-v4)# route 10.5.222.2/32
tnsr(config-rttbl4-next-hop)# next-hop 0 via 203.0.113.14 TenGigabitEthernet6/0/0
```

Setup BGP with Loopback Address

Now setup the BGP service, using the new neighbor address and with the loopback address as an update source:

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# server 65002
tnsr(config-bgp)# router-id 10.2.0.1
tnsr(config-bgp)# neighbor 10.5.222.2
tnsr(config-bgp-neighbor)# remote-as 65005
tnsr(config-bgp-neighbor)# update-source 10.5.222.1
tnsr(config-bgp-neighbor)# enable
tnsr(config-bgp-neighbor)# exit
tnsr(config-bgp)# address-family ipv4 unicast
tnsr(config-bgp-ip4uni)# network 10.2.0.0/16
tnsr(config-bgp-ip4uni)# exit
tnsr(config-bgp)# exit
tnsr(config-frr-bgp)# enable
tnsr(config-frr-bgp)# exit
tnsr(config)# service bgp restart
```

11.2.3 BGP Configuration

The BGP service on TNSR contains numerous methods to configure and fine-tune BGP routing behavior. Due to this complexity, the topic has been split into several sections. Read through each section before attempting to create a new BGP configuration.

Enabling BGP

The BGP service has a master enable/disable toggle that must be set before BGP will operate. Enable BGP using the enable command in config-frr-bgp mode:

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# bgp enable
```

To disable the service, use `no enable` or `disable`.

The BGP service is managed as described in *Service Control*.

Warning: After starting or restarting TNSR, restart the BGP service from within the TNSR configuration mode CLI to ensure that the routes from BGP neighbors are fully populated throughout TNSR:

```
tnsr(config)# service bgp restart
```

BGP Router Configuration

This statement enters BGP Server mode and sets the autonomous system number for this router to 65002, and enters `config-bgp` mode.

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# server 65002
tnsr(config-bgp)#
```

BGP mode defines the main behaviors of the BGP daemon, as well as the neighbors and behavior of BGP for different address families, among other possibilities.

From within `config-bgp` mode, the following commands are available:

address-family (ipv4|ipv6) (unicast|multicast) Enter *BGP Address Family Configuration* mode.

always-compare-med Instructs the BGP daemon to always consult MED values in routes, no matter which AS the routes were received through.

bestpath as-path (confed|ignore|multipath-relax|as-set|no-as-set) Controls how the BGP daemon determines the best path to a destination. May be one of:

confed Considers the length of confederation path sets and sequences.

ignore Ignores AS path lengths when computing the route to a destination.

multipath-relax Consider paths of equal length when choosing between multiple paths to a destination, rather than looking for an exact match. This allows load sharing across different AS paths, so long as they are of equal length.

as-set For use with `multipath-relax`, it adds AS set information for aggregate routes.

no-as-set For use with `multipath-relax`, it prevents AS set generation.

bestpath compare-routerid Uses the router ID of peers (or originator ID, if present) to break ties when computing paths to a destination based on other information. A lower router ID will win in a tie.

bestpath med confed Compare confederation path MEDs

bestpath med missing-as-worst If a route is missing MED information, it will be considered least preferred.

client-to-client reflection Enables reflection of routes from one client to another client.

cluster-id (<ipv4>|<value>) Configures the BGP daemon to participate in route reflection with the given cluster ID. The ID may be given in IP address (dotted quad) notation or as an unsigned 32-bit integer (1-4294967295).

coalesce-time <value> Configures the Subgroup coalesce timer, in milliseconds (1-4294967295).

confederation identifier <ASN> Configures an AS number for the entire group of IBGP routers participating in confederation.

confederation peer <ASN> Configures the sub-AS number for the subset of peers inside a group of IBGP routers participating in confederation.

dampening [penalty <val> [reuse <val> [suppress <val> [maximum <max>]]]] This command enables BGP route flap dampening (RFC 2439) to prevent unstable routers from adversely affecting routing behavior.

penalty <penalty-val> The time duration during which the stability value will be reduced by half if the route is unreachable.

reuse <reuse-val> Stability threshold that must be crossed for a route to be reused.

suppress <suppress-val> Stability threshold that, when crossed, a route will be suppressed.

maximum <suppress-max> Maximum time to suppress a route considered stable.

deterministic-med Determine route selection locally, even when MED values are present. Picks the best MED path from neighbor advertisements.

disable-ebgp-connected-route-check Disable checking if nexthop is an eBGP session.

enforce-first-as When set, enforces the first AS for eBGP routes.

listen limit <value> Maximum number of dynamic neighbors from 1–5000.

listen range (<ip4-prefix>|<ip6-prefix>) peer-group <peer-group-name> Listen range for dynamic neighbors.

max-med administrative [<med>] Sends the defined MED value, or 4294967294 when unset, at all times.

max-med on-startup period <seconds> [<med>] Sends the defined MED value, or 4294967294 when unset, only at startup for the defined period in seconds, from 5–86400.

neighbor <peer> Enter *BGP Neighbor Configuration* mode.

network import-check Checks if a BGP network route exists in IGP before creating BGP table entries.

route-reflector allow-outbound-policy Allows attributes modified by route maps to be reflected.

router-id <A.B.C.D> Sets the router ID for the BGP daemon. This is typically set to an IP address unique to this router, and commonly is set to a local private address.

timers keep-alive <interval> hold-time <hold-time> Configures the intervals between keep alive messages and how long to wait for a response before considering the peer unreachable.

update-delay <delay> Keeps BGP in a read-only mode for the specified time after the daemon restarts or peers are cleared.

write-quanta <packets> Controls the size of peer update transmissions.

BGP Neighbor Configuration

From within `config-bgp` mode, the `neighbor` statement can take either an IP address to setup a single neighbor or it can take a name which configures a peer group. The command then changes to `config-bgp-neighbor` mode.

```
tnsr(config-bgp)# neighbor 203.0.113.14
tnsr(config-bgp-neighbor)#
```

Peer groups work nearly identical to neighbors, and they define options that are common to multiple neighbors.

Warning: A neighbor or peer group must first be defined here before it can be used inside an address family (*BGP Address Family Neighbor Configuration*).

`config-bgp-neighbor` mode contains the following commands:

advertisement-interval <interval-sec> Minimal time between sending routing updates to this neighbor. Expressed in seconds between 0–600.

bfd enabled (true|false) Enable Bidirectional Forwarding Detection for this BGP neighbor.

capability dynamic Enables negotiation of the `dynamic` capability with this neighbor or peer group.

capability extended-nexthop Enables negotiation of the `extended-nexthop` capability with this neighbor or peer group. This capability can set IPv6 next-hops for IPv4 routes when peering with IPv6 neighbors on interfaces without IPv4 connectivity. This is automatically enabled when peering with IPv6 link-local addresses.

disable-connected-check Disables a check that normally prevents peering with eBGP neighbors which are not directly connected. This enables using loopback interfaces to establish adjacencies with peers.

description <string> A brief text description of this neighbor.

dont-capability-negotiate Disables dynamic capability negotiation with the peer. When set, the router does not advertise capabilities, nor does it accept them. This results in using only locally configured capabilities.

ebgp-multihop [hop-maximum <hops>] The maximum allowed hops between this router and the neighbor, in the range 1–255. When enabled without a specific value, the default is 1. This value cannot be set if `ttl-security` is set.

(enable|disable) The default state of a neighbor is disabled. To enable the neighbor, use the `enable` command. To disable the neighbor, run `disable` or `no enable`.

local-as <asn> [no-prepend [replace-as]] Sets the local AS number sent to this neighbor, which replaces the AS number configured on the BGP server itself. By default, this value is prepended to the AS path for routes received from this neighbor or peer group, and is added to the AS path for routes sent to this neighbor or peer group after the AS number from the BGP server.

no-prepend Suppresses prepending this AS number to the AS path for received routes.

replace-as Suppresses prepending the BGP server AS to transmitted routes, so that only this value is present.

override-capability Ignores capabilities sent by the peer during negotiation and uses locally configured capabilities instead.

passive When set, this router will not issue requests to the neighbor on its own. The BGP daemon will only respond to remote requests from this neighbor.

password <line> A password used by BGP for TCP-MD5 (RFC 2385) authentication of communications with the neighbor, up to 64 characters in length.

peer-group [<peer-group-name>] Configure this neighbor as a member of the given peer group. Only valid for use in neighbors defined by address, not on peer groups.

port <port> An alternate port number used by this daemon for BGP messages, if it uses a value other than TCP port 179.

remote-as <asn> The remote AS number of this neighbor.

solo Instructs the router to prevent reflection of routes received from this neighbor back to this neighbor. This command is not useful in peer groups with multiple members.

strict-capability-match When set, enforces the comparison between the set of capabilities sent by the peer during negotiation and the set of capabilities present in the local configuration. If there is a mismatch, an error is transmitted to the peer.

timers keepalive <interval> holdtime <hold> Configures the intervals between keep alive messages and how long to wait for a response from this neighbor before considering the peer unreachable. This overrides the default values set on the BGP server itself. Both values must be in the range 0–65535, in seconds.

timers connect <seconds> The amount of time, in seconds from 1–65535, in which a connection to this peer must be established or else it is considered unsuccessful.

ttl-security hops <hops> Similar to `ebgp-multihop` but sets a specific hop count at which neighbors must be reached, rather than the maximum value set by `ebgp-multihop`. This command cannot be set if `ebgp-multihop` is set.

update-source (<ifname>|<ip-address>) Configures a specific interface or IP address to use when sending messages to this peer.

Note: Within BGP neighbor mode, the most important directives are `remote-as` to set the AS number of the neighbor and `enable`. The majority of other neighbor configuration is handled by the neighbor definition for a specific address family (*BGP Address Family Neighbor Configuration*).

BGP Address Family Configuration

The TNSR BGP implementation is capable of handling routing information for IPv4 and IPv6 independently, among other network layer protocols. The `address-family <family> <type>` command defines BGP behavior for each specific supported case. The most common address families are `ipv4 unicast` and `ipv6 unicast`. The other possible choices supported in this version are `ipv4 multicast` and `ipv6 multicast`.

The `address-family` command changes to BGP address family mode, which contains settings specific to each address family. The prefix for this mode varies depending on the address family command which entered the mode. For example, when configuring settings for the IPv4 unicast address family, the prompt indicates `config-bgp-ip4uni`.

```
tnsr(config-bgp)# address-family ipv4 unicast
tnsr(config-bgp-ip4uni)#
```

Each resulting mode, such as `config-bgp-ip4uni` or `config-bgp-ip6uni`, contains its own set of commands. As these may differ, they are split up in multiple sections here.

IPv4 or IPv6 Unicast

The following commands are available in `config-bgp-ip4uni` and `config-bgp-ip6uni` modes:

aggregate-address <ip-prefix> [as-set] [summary-only] This command configures route aggregation using the specified prefix. More specific routes contained within the specified prefix will be aggregated into the larger prefix, minimizing the set of networks advertised to peers.

as-set When present, routes for the specified prefix will include an AS set. An AS set is a collection of AS numbers for which routes have been aggregated. This allows peers to detect routing loops, duplicate routes, and so on.

summary-only When present, aggregated routes for this prefix will not be announced, so peers only see the aggregate prefix and not the component networks.

distance external <extern> internal <intern> local <local> Configures distance values which control how BGP will treat routes based on the length of their AS path.

external <extern> The distance at which routes are considered external, from 1–255.

internal <intern> The distance at which routes are considered internal, from 1–255.

local <local> The distance at which routes are considered local, from 1–255.

distance administrative <dist> prefix <ip-prefix> [access-list <access-list-name>] This command manually configures the administrative distance for a given prefix, with the following required parameters:

administrative <dist> The administrative distance for this prefix, from 1–255.

prefix <ip-prefix> The IP prefix to which this distance will be applied.

access-list <access-list-name> An access list which can be used to apply the distance to only a subset of the configured prefix.

maximum-paths <non-ibgp-paths> [igbp <ibgp-paths> [equal-cluster-length]] Configures the maximum number of paths for multi-path eBGP forwarding. This is enabled by default with a value of 64. This allows the router to utilize multiple equal identical paths via different routers.

Paths for prefixes advertised by multiple eBGP peers in the same AS are considered equal cost and result in a multi-path route.

Note: As this feature is enabled by default, to disable this behavior, set the value to 1 which limits routes to only a single path.

igbp <ibgp-paths> Configures a value for multi-path forwarding in iBGP roles.

equal-cluster-length Only consider paths as matching when cluster lengths are also equal.

neighbor <existing-neighbor> Specifies an existing neighbor address or peer group to use with this address family, and enters *BGP Address Family Neighbor Configuration* mode.

Warning: This command cannot define a new neighbor. A neighbor or peer group must first be defined using the `neighbor` command from within `config-bgp` mode before it can be used here.

network <ip-prefix> [route-map <route-map>] Configures a prefix to be advertised to peers in this address family.

route-map <route-map> Specifies a route map used to limit advertisements of this prefix.

redistribute <route-source> [metric <val>|route-map <route-map-name>] Enables redistribution of routes from another source. Available route sources are listed in *Dynamic Routing Protocol Lists*.

metric <val> A MED value to apply to redistributed routes.

route-map <route-map-name> Specifies a route map used to limit redistributed route advertisements.

redistribute ospf [metric <val>|route-map <route-map-name>] Configure redistribution of routes from OSPF.

metric <val> A MED value to apply to redistributed routes.

route-map <route-map-name> Specifies a route map used to limit redistributed route advertisements.

redistribute table id <kernel-table-id> [metric <val>|route-map <route-map-name>] Configure redistribution of routes only from a specific kernel routing table, rather than all tables.

metric <val> A MED value to apply to redistributed routes.

route-map <route-map-name> Specifies a route map used to limit redistributed route advertisements.

table-map <route-map-name> Uses the specified route map to control how routes received from BGP peers are passed to the dynamic routing manager process, and thus, into routing tables.

IPv4 Multicast

The following commands are available in `config-bgp-ip4multi` mode. See *IPv4 or IPv6 Unicast* for descriptions of the commands and parameters:

- `aggregate-address`
- `distance external`
- `distance administrative`
- `neighbor`
- `network`
- `table-map`

IPv6 Multicast

The following commands are available in `config-bgp-ip6multi` mode. See *IPv4 or IPv6 Unicast* for descriptions of the commands and parameters:

- `distance external`
- `distance administrative`
- `neighbor`
- `network`

BGP Address Family Neighbor Configuration

From within a BGP address family configuration mode, the `neighbor <existing-neighbor>` command specifies an existing neighbor defined in *BGP Neighbor Configuration* mode. This command then enters an address-family-specific neighbor mode. Like address families, the prefix for this mode varies based on the family and type of address family it is run within. For example, with IPv4 unicast mode, the prompt is `config-bgp-ip4uni-nbr`.

```
tnsr(config-bgp-ip4uni)# neighbor 203.0.113.14
tnsr(config-bgp-ip4uni-nbr)#
```

The following commands are available in `config-bgp-<familytype>-nbr` modes:

activate Activate this neighbor for use by BGP.

addpath-tx-all-paths Advertise all known paths to this peer, instead of only advertising the base path.

addpath-tx-bestpath-per-as Advertise only the best known base paths for each AS.

allowas-in [**<occurrence>|origin**] Allows routes to be received from this peer which are from the same AS of this router, but through a different path.

<occurrence> Allowed number of AS occurrences, from 1–10.

origin Accept the AS of this router in an AS-path if the route originated in the AS of this router.

as-override Override ASNs in outbound updates to this peer if the AS path is identical to the remote AS.

attribute-unchanged [**as-path|next-hop|med**] Propagates route attributes to this peer unchanged. This behavior can be optionally restricted to only specific attributes, including the `as-path`, `next-hop`, and `med` attributes.

capability orf prefix-list (send|receive|both) Advertise outbound route filtering capability to this peer. This behavior can be restricted by direction, `send`, `receive`, or `both`.

default-originate [**route-map <route-map>**] Enables advertisement of a default route to this peer.

route-map <route-map> Restricts this behavior based on the specified route map.

distribute-list <access-list-name> (in|out) Defines an access list which is used by BGP to filter route updates for this peer, in either the `in` or `out` direction.

filter-list <aspath-name> (in|out) Defines a list which is used by BGP to filter route updates by AS path, rather than prefix.

maximum-prefix [(limit|restart|threshold) <value>|warning-only] Defines the maximum number of prefixes this router will accept from the peer before tearing down the BGP session.

Note: This action is considered harsh and the best practice is to filter received prefixes by other mechanisms such as a `prefix-list` rather than to abruptly break contact in this way.

limit <val> The maximum number of prefixes to allow from the peer, from 1–4294967295.

restart <val> Restarts the connection after limits are exceeded. The restart is performed at the defined interval, in minutes, from 1–65535.

threshold <val> Warning message threshold, from 1–100.

warning-only Warn the peer when the limit is exceeded, rather than disconnecting.

next-hop-self [**force**] Uses the address of this router as the next-hop in routes announced to this peer if they are learned via eBGP.

force When present, also sets the next-hop to the address of this router on reflected routes.

prefix-list <prefix-list-name> (in|out) Defines a prefix list which is used by BGP to filter route updates for this peer, in either the `in` or `out` direction.

remove-private-AS [**all**] [**replace-AS**] Prevents the BGP daemon from sending routes with private AS numbers to this peer.

all When present, this action applies to all ASNs.

replace-AS When present, replaces private AS numbers with the AS number of this router.

route-map <name> (in|out) Defines a route map which is used by BGP to filter route updates for this peer, in either the `in` or `out` direction.

route-reflector-client Configures this peer as a route reflector client. This allows routes received from peers in the same AS or using iBGP to be reflected to other peers, avoiding the need for a full mesh configuration between all routing peers.

route-server-client Configures this peer as a route server client. This enables transparent mode, which retains attributes unmodified, and maintains a local RIB for this peer.

send-community (standard|large|extended) Sends the community attribute to this peer, limited to the specified type (standard, large, extended).

soft-reconfiguration inbound Allows the peer to send requests for soft reconfiguration, to apply changes to routes or new attributes without the need for a session reset.

unsuppress-map <route-map> Configures a route map which BGP can use to unsuppress routes that would otherwise be suppressed by other configuration settings.

weight <weight> Applies the given weight to routes received from this peer.

BGP AS Path Access Lists

AS Path access lists entries determine if networks are allowed or denied in specific BGP configuration contexts. They are primarily used in BGP route maps, but also can be used in other areas of BGP configuration which accept AS Path lists as parameters.

The order of entries inside an AS Path list is important, and this order is determined by a sequence number. As with other access lists, AS Path access lists implicitly deny anything not matched.

BGP AS Path Configuration

To create a new AS Path list, from `config-frr-bgp` mode, use the `as-path <name>` command, which enters `config-aspath` mode:

```
tnsr(config-frr-bgp)# as-path myasp
tnsr(config-aspath)#
```

`config-aspath` mode contains only the rule `<seq> (permit|deny) <pattern>` command which defines a new AS Path rule with the following parameters:

<seq> The sequence number for this rule, which controls the order in which rules are matched inside this AS Path list. Each rule must have a unique sequence number. Best practice is to leave gaps in the sequence to allow for adding rules in the future. For example, use 10, 20, 30, rather than 1, 2, 3.

(permit|deny) The action taken when this AS Path rule is matched, either `permit` or `deny`.

<pattern> A [regular expression](#) pattern which will match on the AS number.

Regular expression patterns support common pattern special characters for matching, but also a special `_` character. The `_` character matches common AS delimiters such as start of line, end of line, space, comma, braces, and parenthesis. The `_` character can be used on either side of an AS number to match it exactly, such as `_65534_`.

BGP AS Path Example

This AS Path could match an empty AS value or the specific value of 65002, and no others:

```
tnsr(config-frr-bgp)# as-path myasp
tnsr(config-aspath)# rule 10 permit ^$
tnsr(config-aspath)# rule 20 permit _65002_
tnsr(config-aspath)# exit
tnsr(config-frr-bgp)#
```

This AS Path will match only when the path being compared starts with 65500. This is a common way to ensure that routes from a peer contain the expected AS in the AS Path.

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# as-path R2-AS
tnsr(config-aspath)# rule 10 permit ^65005
tnsr(config-aspath)# exit
tnsr(config-frr-bgp)# exit
tnsr(config)# route dynamic route-map CHECK-R2-AS permit sequence 10
tnsr(config-route-map)# match as-path R2-AS
tnsr(config-route-map)# exit
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# server 65002
tnsr(config-bgp)# neighbor 10.2.222.2
tnsr(config-bgp-neighbor)# remote-as 65005
tnsr(config-bgp)# address-family ipv4 unicast
tnsr(config-bgp-ip4uni)# neighbor 10.2.222.2
tnsr(config-bgp-ip4uni-nbr)# route-map CHECK-R2-AS in
tnsr(config-bgp-ip4uni-nbr)# exit
tnsr(config-bgp-ip4uni)# exit
tnsr(config-bgp)# exit
tnsr(config-frr-bgp)# exit
tnsr(config)#
```

BGP AS Path Status

To view AS Path lists, use the `show route dynamic bgp as-path [<name>]` command. Add the name of an AS Path list to restrict the output to a single entry.

```
tnsr(config)# show route dynamic bgp as-path

Name  Seq Policy Pattern
-----
R2-AS  10 permit ^65005

myasp  10 permit ^$
myasp  20 permit _65002_
```

BGP Community Lists

A BGP community, as defined in [RFC 1997](#), is a group of destinations which share common properties. Community Lists define sets of community attributes which the BGP daemon can use to match or set community values in routing updates. BGP communities determine AS membership and priority values in BGP-specific contexts such as route-maps.

The order of entries inside a Community List is important, and this order is determined by a sequence number.

BGP Well-Known Communities

There are several “well-known” communities available for use in Community Lists. Each of these communities have special meanings:

- internet** A community value of 0, indicating the Internet as a destination.
- no-export** Routes received carrying this attribute value must not be exported to routers outside of the current confederation.
- no-advertise** Routes received carrying this attribute value must not be advertised to any other BGP peer.
- local-as** Also known as “No Export Subconfed”. Routes received carrying this attribute value must not be advertised to any external BGP peer, even those in the same confederation.
- blackhole** Routes received carrying this attribute should not be routed (e.g. null routed).
- graceful-shutdown** Indicates support for [RFC 8326](#) Graceful Shutdown, which allows BGP routers to indicate to peers that specific paths can be gracefully shut down rather than abruptly terminated when performing an intentional shutdown.
- no-peer** Indicates that routes with this community value should not be readvertised to peers ([RFC 3765](#)).

BGP Community List Configuration

To create a new Community List, from `config-frr-bgp` mode, use the `community-list <name> (standard|expanded) [normal|extended|large]` command, with the following parameters:

- <name>** The name of this BGP Community List.
- (standard|expanded)** The type of Community List, either `standard` or `expanded`:
 - standard** Matches based on specific values for community attributes.
 - expanded** Matches based on an ordered list using a regular expression. Due to the use of regular expression evaluation, these lists incur a performance penalty.
- [normal|extended|large]** The type of communities contained inside this Community List, either `normal`, `extended`, or `large`.
 - normal** Normal community values as described in [RFC 1997](#).
 - extended** Extended BGP communities specified using 8-octet values as described in [RFC 5668](#). These communities also allow for IPv4-based policies.
 - large** Large BGP communities specified using 12-octet values as described in [RFC 8092](#) and [RFC 8195](#).

The full `community-list` command enters `config-community-list` mode:

```
tnsr(config-frr-bgp)# community-list mycom standard normal
tnsr(config-community-list)#
```

`config-community-list` mode contains the following commands:

description

sequence <seq> (permit|deny) <community-value>

- <seq>** The sequence number for this rule, which controls the order in which rules are matched inside this Community List. Each rule must have a unique sequence number. Best practice is to leave gaps in the sequence to allow for adding rules in the future. For example, use 10, 20, 30, rather than 1, 2, 3.

(permit|deny) The action taken when this Community List rule is matched, either permit or deny.

<community-value> The value of the community to match.

Standard Community Lists This is a space-separated list of communities in AS:VAL format, or from the *BGP Well-Known Communities* list.

Expanded Community Lists A string containing a regular expression to match against.

Regular expression patterns support common pattern special characters for matching, but also a special `_` character. The `_` character matches common AS delimiters such as start of line, end of line, space, comma, braces, and parenthesis.

BGP Community List Example

This example sets up a Community List for the AS:VAL pair of AS 65002 and community value 10:

```
tnsr(config-frr-bgp)# community-list mycom standard normal
tnsr(config-community-list)# sequence 10 permit 65002:10
tnsr(config-community-list)# exit
tnsr(config-frr-bgp)#
```

This example sets up a Community List, used by a route map, to prevent distribution of routes marked with the well-known community `no-export`:

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# community-list POISON-ROUTES standard normal
tnsr(config-community-list)# sequence 10 permit no-export
tnsr(config-community-list)# exit
tnsr(config-frr-bgp)# exit
tnsr(config)# route dynamic route-map OUT deny sequence 10
tnsr(config-route-map)# match ip address prefix-list RFC1918
tnsr(config-route-map)# exit
tnsr(config)# route dynamic route-map OUT deny sequence 20
tnsr(config-route-map)# match community POISON-ROUTES
tnsr(config-route-map)# exit
tnsr(config)# route dynamic route-map OUT permit sequence 30
tnsr(config-route-map)# match ip address prefix-list MY-ROUTES
tnsr(config-route-map)# exit
tnsr(config)#
```

Note: In this example, note the use of `permit` in the Community List, which will succeed on a positive match. The route map then uses `deny` when a positive match is made on the community value.

BGP Community List Status

To view Community Lists, use the `show route dynamic bgp community-list [<name>]` command. Add the name of a Community List to restrict the output to a single entry.

```
tnsr(config)# show route dynamic bgp community-list
```

```

Name   Type      Size  Description
-----
POISON-ROUTES standard normal
      Seq Action Community
      ---
      10  permit no-export

mycom standard normal
      Seq Action Community
      ---
      10  permit 65002:10

```

11.2.4 BGP Status

TNSR supports several commands to display information about the BGP daemon configuration and its status.

See also:

For more general dynamic routing status information, see *Dynamic Routing Manager Status*

Configuration Information

To view the BGP configuration:

```
tnsr# show route dynamic bgp config [<as-number>]
```

To view other individual sections of the configuration:

```

tnsr# show route dynamic bgp as-path [<as-path-name>]
tnsr# show route dynamic bgp community-list [<community-list-name>]

```

Status Information

For a brief summary of BGP status information:

```
tnsr# show route dynamic bgp (ipv4|ipv6) summary
```

For a list of configured BGP Neighbors and their status details:

```

tnsr# show route dynamic bgp (ipv4|ipv6) neighbors [[<peer>] [advertised-routes|
      dampened-routes|flap-statistics|prefix-counts|received|received-
      ↪routes|routes]]

```

For information about a specific BGP peer group:

```
tnsr# show route dynamic bgp peer-group <peer-group-name>
```

For a list of valid BGP next hops:

```
tnsr# show route dynamic bgp nexthop [detail]
```

For details about an address or prefix in the BGP routing table:

```
tnsr# show route dynamic bgp (ipv4|ipv6) network <IP Address|Prefix>
```

BGP Active Session Control

The `session clear` command can be used to reset active BGP sessions. This command is available from within `config-frr-bgp` mode. The general form of the command is:

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# session clear (*|<peer>|<asn>) [soft]
```

The first parameter controls what will be cleared, and values may be completed automatically with `tab`:

- * Clears all open BGP sessions

- <peer>** Clears all sessions to a specific peer IP address or peer group name

- <asn>** Clears all sessions to a specific AS number

The second parameter, `soft` is optional and controls whether or not the command will trigger a soft reconfiguration.

Additional Information

Additional BGP status information can be obtained by using the `vttysh` program outside of TNSR.

The `vttysh` program must be run as root:

```
sudo vttysh
```

The `vttysh` interface offers numerous commands. Of particular interest for BGP status are the following:

- show bgp summary** A brief summary of BGP status information.

- show bgp neighbors** Lists configured BGP Neighbors and their status details.

- show ip bgp** A list of routes and paths for networks involved in BGP.

- show ip route** The IP routing table managed by the FRR Zebra daemon, which marks the origin of routes to see which entries were obtained via BGP.

11.2.5 Working with Large BGP Tables

When working with a large set of routes, roughly exceeding 30,000 route table entries, TNSR may require additional memory to be allocated for the VPP dataplane Forwarding Information Bases (FIB). Smaller routing tables do not require special configuration.

This memory allocation can be performed in configuration mode using one of the following commands:

For IPv4 (*Memory*):

```
tnsr# configure
tnsr(config)# dataplane ip heap-size <size>
```

For IPv6 (*Memory*):

```
tnsr# configure
tnsr(config)# dataplane ip6 heap-size <size>
```

The format of the size is <number> [KMG], for example: 512M or 1G for 512 Megabytes or 1 Gigabyte, respectively. Additionally, the statistics segment heap size may also need to be increased (*Statistics Segment*):

```
tnsr# configure
tnsr(config)# dataplane statseg heap-size <size>
```

Note: The default size for `dataplane statseg heap-size` is 96MB, which is sufficient for approximately one million routes when worker threads are not in use.

See also:

For more details about selecting an appropriate `heap-size` value, especially when worker threads are enabled, see *Statistics Segment*.

The VPP dataplane service requires a restart to enable these configuration changes. Restart VPP from the TNSR configuration mode CLI using the following command:

```
tnsr# configure
tnsr(config)# service dataplane restart
```

11.3 Open Shortest Path First v2 (OSPF)

Open Shortest Path First v2 (OSPF) is a link-state routing protocol defined by [RFC 2328](#). OSPF automatically locates neighboring IPv4 routers within an autonomous system, typically with multicast, and exchanges IPv4 routing information for networks reachable through each neighbor.

OSPF is an interior routing protocol (IGP), and facilitates routing between private links or segments of local networks.

11.3.1 OSPF Required Information

Before starting, take the time to gather all of the information required to form an OSPF adjacency to a neighbor. At a minimum, TNSR will need to know these items:

Local Router ID Typically the highest numbered local address on the firewall. This is also frequently set as the internal or LAN side IP address of a router. It does not matter what this ID is, so long as it is given in IPv4 address notation and does not conflict with any neighbors.

OSPF Area A designation for the set of networks to which this router belongs. Typically set to 0.0.0.0 for simple internal deployments, but can be any number capable of being expressed in dotted quad notation (IPv4 address) or as a 32-bit unsigned integer.

OSPF Active Interfaces The interfaces on this router upon which the OSPF daemon will advertise itself and look for neighbors. These interfaces are connected to network segments with other routers. They may be connected to local networks or remote point-to-point links. These interfaces must be configured with IP addresses.

Warning: Outside NAT interfaces (`ip nat outside`) cannot be used as active interfaces in OSPF! The presence of NAT prevents OSPF from properly communicating with neighbors to form a full adjacency.

OSPF Active Interface Cost Values OSPF calculates the most efficient way to route between networks based on the total cost of a path from source to destination. Less desirable links (e.g. wireless) can be given a higher cost so that paths over faster networks will be used by traffic unless the preferred path is unavailable. For single connections to other networks, this value is not necessary and may be omitted or set to a simple default such as 5 or 10.

OSPF Passive Interfaces These interfaces contain networks which should be advertised as reachable through this router, but do not contain other routers.

The example in this section uses the following values:

Table 2: Example OSPF Configuration

Item	Value
Local Router ID	10.2.0.1
OSPF Area	0.0.0.0
Active Interfaces (Cost)	TenGigabitEthernet6/0/0 (10)
Passive Interfaces	GigabitEthernet3/0/0

11.3.2 OSPF Example

This example configuration implements an OSPF setup using the required information from *Example OSPF Configuration*.

```
tnsr(config)# route dynamic ospf
tnsr(config-frr-ospf)# server
tnsr(config-ospf)# ospf router-id 10.2.0.1
tnsr(config-ospf)# passive-interface GigabitEthernet3/0/0
tnsr(config-ospf)# interface GigabitEthernet3/0/0
tnsr(config-ospf-if)# ip address * area 0.0.0.0
tnsr(config-ospf-if)# exit
tnsr(config-ospf)# interface TenGigabitEthernet6/0/0
tnsr(config-ospf-if)# ip address * cost 5
tnsr(config-ospf-if)# ip address * area 0.0.0.0
tnsr(config-ospf-if)# exit
tnsr(config-ospf)# exit
tnsr(config-frr-ospf)# enable
tnsr(config-frr-ospf)# exit
tnsr(config)#
```

A similar configuration may be applied to neighboring routers also connected to the same network as the TenGigabitEthernet6/0/0 interface. Adjust the router ID and interface names as needed.

For a simple configuration such as this, a single area for all routers is typical.

See also:

For a more complex example involving multiple areas, see *OSPF Router with Multiple Areas and Summarization*.

11.3.3 OSPF Configuration

OSPF configuration on TNSR, as shown in the example, can be fairly straightforward. That said, there are a number of ways to fine-tune the behavior and create complex OSPF routing configurations.

Read through each section before attempting to create a new OSPF configuration.

Enable OSPF

The OSPF service has a master enable/disable toggle that must be set before OSPF will operate. Enable OSPF using the `enable` command in `config-frr-ospf` mode:

```
tnsr(config)# route dynamic ospf
tnsr(config-frr-ospf)# enable
```

To disable the service, use `no enable` or `disable`.

The OSPF service is managed as described in [Service Control](#).

OSPF Server Configuration

To configure an OSPF server, start in `config-frr-ospf` mode and run the `server` command:

```
tnsr(config-frr-ospf)# server
tnsr(config-ospf)#
```

This changes into `config-ospf` mode, which contains the following commands:

area <area-id> Configures area-specific settings in [OSPF Area Configuration](#) mode.

auto-cost reference-bandwidth <bw> A base value, in Mbit/s, which is used when OSPF automatically calculates cost values. The default value is 100 which means that an interface with 100Mbit/s of bandwidth or greater will have a cost of 1, with lower bandwidth values incurring higher cost values.

All routers in the same area should use the same value, otherwise automatic cost calculations would fail to accurately represent total path costs between routers.

capability opaque-lsa Enables support for Opaque LSAs, as described in [RFC 2370](#).

compatible rfc-1583-compatibility Enables compatibility with the older OSPF standard from [RFC 1583](#), which has been obsoleted by the newer [RFC 2328](#). The specific change this option enables relates to external path preference calculation and routing loop prevention. See [RFC 2328](#) section [G.2](#) for specific details.

default-information originate (always|metric <val>|type <type>|route-map <map>) Enables origination of a Type 5 AS-External LSA containing default route information into all areas capable of external routing.

always Always advertise a default route, even when a default route is not present in the local routing table.

metric <val> Advertise the default route as having the given metric.

type <type> The type of metric, either 1 or 2. See [Metric Types](#) for details about each type operates.

route-map <map> Apply the given route map to the outbound route advertisement.

default-metric <val> Uses the given metric value as the default metric for OSPF routes when no other metric information is available.

distance [(external|inter-area|intra-area)] <dist> Sets an administrative distance for routes obtained via OSPF. This can be configured globally as well as for specific types of OSPF routes.

external <dist> Sets the administrative distance for external OSPF routes.

inter-area <dist> Sets the administrative distance for OSPF routes between areas.

intra-area <dist> Sets the administrative distance for OSPF routes inside an area.

distribution-list out <route-source> access-list <name> Applies the given access list <name> to routes redistributed from the specified <route-source>.

Available route sources are listed in *Dynamic Routing Protocol Lists*, with the exception of `ospf` which cannot be used with this command.

interface <if-name> Defines an interface for use with OSPF, and enters *OSPF Interface Configuration* mode.

log-adjacency-changes [detail] Instructs the OSPF daemon to log changes in neighbor adjacencies. This is useful for tracking changes to neighbor relationships, especially during initial configuration.

The optional `detail` parameter increases the verbosity of the resulting log messages.

max-metric router-lsa administrative Sets the administrative distance of routes through this router to infinity, so that other routers will avoid using this router to reach other networks. Networks on this router are still reachable. See [RFC 3137](#) for more information.

max-metric router-lsa (on-shutdown|on-startup) <seconds> Conditionally sets the administrative distance of routes through this router to infinity for a period of time after startup or shutdown. This allows other routers in the area to avoid using routes through this router until a full convergence is achieved.

neighbor <ip4-address> [(poll-interval <interval>|priority <prio>)] Configures per-neighbor settings for polling and priority for non-multicast neighbors.

poll-interval <interval> Time, in seconds, between sending OSPF Hello messages to neighbors in a down state.

priority <prio> A priority value applied to neighbors in a down state.

ospf abr-type (cisco|ibm|shortcut|standard) Controls the behavior of Area Border Router (ABR) functionality.

cisco|ibm The default behavior of OSPF on TNSR, discussed in [RFC 3509](#). This behavior allows an ABR without a backbone connection to act as an internal router for all connected areas.

shortcut Discussed in [draft-ietf-ospf-shortcut-abr-02](#), this behavior allows ABRs to consider summary LSAs from all attached areas, rather than being forced to route through a suboptimal path only because it is shorter.

standard The ABR behavior described in the original OSPF standard. When set, a router attached to multiple areas requires a connection to a backbone. If no backbone is available, traffic attempting to cross areas will be dropped.

ospf router-id <router-id> Sets the router ID for the OSPF daemon. This is typically set to an IP address unique to this router, and commonly is set to a local private address.

ospf write-multiplier <write> Number of interfaces processed per write operation, from 1–100. Default value is 20.

passive-interface <if-name> [<ip4-address>] Configures the specified interface as passive. This prevents the interface from actively participating in OSPF, while still allowing OSPF to operate on networks connected to that interface. This is commonly used for local interfaces without other routers attached. OSPF will announce networks attached to passive interfaces as stub links.

pce address (<ip4-address>|domain <asn>|flags <bits>|neighbor <asn>|scope <bits>) Configures [RFC 5088](#) Path Computation Element (PCE) Discovery for OSPF. When active, this router will advertise support for PCE to neighbors via router information (RI) announcements. Requires `router-info` as to also be enabled.

<ip4-address> The IP address used to reach the PCE

domain <asn> AS numbers of domains controlled by the PCE, meaning it can compute paths for the autonomous systems and has visibility into them.

flags <bits> Capability flags for the PCE, expressed as a bit pattern. The bits meanings are defined in [RFC 5088](#) section 4.5:

Table 3: PCE Capability Flags

Bit	Capability
0	Path computation with GMPLS link constraints
1	Bidirectional path computation
2	Diverse path computation
3	Load-balanced path computation
4	Synchronized path computation
5	Support for multiple objective functions
6	Support for additive path constraints (max hop count, etc.)
7	Support for request prioritization
8	Support for multiple requests per message

neighbor <asn> AS numbers of neighboring domains for which the PCE can compute paths.

scope <bits> Scope for path computation, such as intra-area, inter-area, inter-AS, or inter-layer, expressed as a bit mask. The bits meanings are defined in [RFC 5088](#) section 4.2:

Table 4: PCE Scope

Bit	Path Scope
0	L bit: Can compute intra-area paths.
1	R bit: Can act as PCE for inter-area TE LSP computation.
2	Rd bit: Can act as a default PCE for inter-area TE LSP computation.
3	S bit: Can act as PCE for inter-AS TE LSP computation.
4	Sd bit: Can act as a default PCE for inter-AS TE LSP computation.
5	Y bit: Can act as PCE for inter-layer TE LSP computation.

redistribute <route-source> [(metric <val>|route-map <map>|type <type>)] Enables redistribution of routes from another source. Available route sources are listed in [Dynamic Routing Protocol Lists](#).

metric <val> Advertise the default route as having the given metric.

type <type> The type of metric, either 1 or 2. See [Metric Types](#) for details about each type operates.

route-map <map> Apply the given route map to the redistributed route advertisements.

refresh timer <time> Time, in seconds from 10–1800, between refreshing LSA information. Default value is 10.

router-info as Enables advertisement of optional router capabilities to neighbors, as described in [RFC 4970](#). This adds information about enabled features, such as PCE, to Router Information (RI) LSA messages.

timers lsa min-arrival <min> The minimum time allowed between advertisements by neighbors, from 0–600000, in milliseconds. Default is 1000.

timers throttle lsa all <delay> Time between LSA transmissions from this router, in milliseconds, from 0–5000. Default is 5000.

timers throttle spf (delay|initial-hold|maximum-hold) <val> Controls timers that determine when the router will make SPF routing decisions.

delay <val> Minimum time after an event occurs before allowing SPF calculation. Lower values will react faster to changes, but can be less stable. Specified in milliseconds from 0–600000, with a default value of 0.

initial-hold <val> Lowest time allowed between SPF calculations. Specified in milliseconds from 0–600000, with a default value of 50.

maximum-hold <val> Highest time allowed between SPF calculations. Specified in milliseconds from 0–600000, with a default value of 5000.

SPF calculations are adaptive, and if a new event occurs which would otherwise trigger a calculation before the hold timer expires, then the hold is increased by the `initial-hold` value, up to the specified `maximum-hold`. This avoids excessive consecutive recalculations.

OSPF Interface Configuration

OSPF must use one or more interfaces to announce itself to neighbors and to receive announcements from neighbors. At least one interface must be configured and active in order to locate neighbors and form an adjacency.

Warning: Outside NAT interfaces (`ip nat outside`) cannot be used as active interfaces in OSPF! The presence of NAT prevents OSPF from properly communicating with neighbors to form a full adjacency.

To configure an interface for use with OSPF, start in `config-ospf` mode and use the `interface <if-name>` command to enter `config-ospf-if` mode.

```
tnsr(config-ospf)# interface <if-name>
tnsr(config-ospf-if)#
```

`config-ospf-if` mode contains the following commands:

bfd enabled (true|false) Enable Bidirectional Forwarding Detection for OSPF on this interface.

ip address (*|<ip4-address>) These commands specify how OSPF will behave for all addresses on an interface (*) for for a specific IPv4 address on an interface. In most cases, the * form will be used here, but when there are multiple addresses available on an interface, a specific choice may be necessary.

area <area-id> This command defines the interface as a member of the given area. This is required to activate an interface for use by OSPF.

authentication [message-digest|null] Configures authentication for OSPF neighbors on this interface. All routers connected to this interface must have identical authentication configurations. This can also be enabled in the area settings.

When run without parameters, simple password authentication is used.

message-digest When set, enables MD5 HMAC authentication for this interface.

null When set, no authentication is used by OSPF on this interface. This is the default behavior, but may be explicitly configured with this command to override the authentication configured for this area.

authentication-key <key> Configures a simple password to use for authentication when that type of authentication is active. This password may only have a maximum length of 8 characters.

Warning: This method of authentication is weak, and MD5 HMAC authentication should be used instead if it is supported by all connected routers.

cost <link-cost> A manual cost value to apply to this interface, rather than allowing automatic cost calculation to take place.

In situations where multiple paths are possible to the same destination, this allows OSPF to prefer one path over another when all else is equal.

dead-interval <time> Time, in seconds from 1–65535, without communication from a neighbor on this interface before considering it dead. This is also known as the RouterDeadInterval timer in OSPF. Default value is 40. This timer should be set to the same value for all routers.

dead-interval minimal hello <multiplier> When active, the dead-interval is forced to a value of 1 and OSPF will instead send <multiplier> number of Hello messages each second. This allows for faster convergence, but will consume more resources.

Note: When set, this overrides the values of both dead-interval and hello-interval. Custom values configured with those commands will be ignored by OSPF.

hello-interval <interval> The interval, in seconds from 1–65535, at which this router will send hello messages. This is also known as the HelloInterval timer in OSPF. Default value is 10. This timer should be set to the same value for all routers.

A lower value will result in faster convergence times, but will consume more resources.

message-digest-key key-id <id> md5-key <key> Configures MD5 HMAC authentication for use with message-digest type authentication.

key-id <id> An integer value from 1–255 which identifies the secret key. This value must be identical on all routers.

md5-key <key> The content of the secret key identified by key-id, which is used to generate the message digest. Given as an unencrypted string, similar to a password. The maximum length of the key is 16 characters.

mtu-ignore When present, OSPF will ignore the MTU advertised by neighbors and can still achieve a full adjacency when peers do not have matching MTU values.

retransmit-interval <interval> The interval, in seconds from 1–65535, at which this router will retransmit Link State Request and Database Description messages. This is also known as the RxmtInterval timer in OSPF. Default value is 5.

priority <priority> A priority value, from 0–255, assigned to this router. When determining which router will become the Designated Router (DR), the router with the highest priority is more likely to be elected as the DR.

The default value is 1. The value 0 is special and will prevent this router from being chosen as DR.

transmit-delay <delay> The interval, in seconds from 1–65535, at which this router will transmit LSA messages. This is also known as the `InfTransDelay` timer in OSPF. Default value is 1.

ip network (broadcast|non-broadcast|point-to-multipoint|point-to-point) Manually configures a specific type of network used on a given interface, rather than letting OSPF determine the type automatically. This controls how OSPF behaves and how it crafts messages when using an interface.

broadcast Broadcast networks, such as typical Ethernet networks, allow multiple routers on a segment and OSPF can use broadcast and multicast to send messages to multiple targets at once. OSPF assumes that all routers on broadcast networks are directly connected and can communicate without passing through other routers.

non-broadcast Non-broadcast networks support multiple routers but do not have broadcast or multicast capabilities. Due to this lack of support, neighbors must be manually configured using the `neighbor` command. When using this mode, OSPF simulates a broadcast network using Non-Broadcast Multi-Access (NBMA) mode, but transmits messages to known neighbors directly.

point-to-multipoint Similar to `non-broadcast` mode, but connections to manually configured neighbors are treated as a collection of point-to-point links rather than a shared network. Similar to a point-to-point network, OSPF disables DR election.

point-to-point A point-to-point network links a single pair of routers. The interface is still capable of broadcast, and OSPF will dynamically discover neighbors. With this type of network, OSPF disables election of a DR.

OSPF Area Configuration

To configure area-specific settings in OSPF, start in `config-ospf` mode and use the `area <area-id>` command to enter `config-ospf-area` mode.

```
tnsr(config-ospf)# area <area-id>
tnsr(config-ospf-area)#
```

`config-ospf-area` mode contains the following commands:

authentication Enables authentication for this area. Communication from peers must contain the expected authentication information to be accepted, and outgoing packets will have authentication information added.

When present on its own, the authentication mechanism used is simple passwords. Authentication passwords are configured in *OSPF Interface Configuration* mode using the `authentication-key` command.

message-digest When present, enables MD5 HMAC authentication for this area. Much stronger authentication than simple passwords. The key is configured in *OSPF Interface Configuration* mode using the `message-digest-key` command.

default-cost <cost> Sets the cost applied to default route summary LSA messages sent to stub areas.

export-list <acl-name> Uses the given ACL to limit Type 3 summary LSA messages for intra-area paths that would otherwise be advertised. This behavior only applies if this router is the ABR for the area in question.

filter-list (in|out) prefix-list <prefix-list-name> Similar to `export-list` and `import-list` but uses prefix lists instead of ACLs, and can work in either direction.

import-list <acl-name> Similar to `export-list`, but for routes announced by other routers into this area.

nssa [(no-summary|translate (always|candidate|never))] Configures this area as a Not-so-Stubby Area (NSSA), which does not contain external links but may contain static routes to non-OSPF destinations (See [Area Types](#) for more information on area types and behaviors).

no-summary When present, the area will instead of considered an NSSA Totally Stub area ([Area Types](#)).

translate (always|candidate|never) Configures NSSA-ABR translations, for converting between Type 5 and Type 7 LSAs.

always Always translate messages.

candidate Participate in NSSA-ABR candidate elections. Currently the default behavior.

never Never translate messages.

range <prefix> [cost <val>|not-advertise|substitute <sub-prefix>] Configure summarization of routes inside the given prefix. Instead of Type 1 (Router) and Type 2 (Network) LSAs, it creates Type 3 Summary LSAs instead.

cost <val> Apply the specified cost to summarized routes for this prefix.

not-advertise Disable advertisement for this prefix.

substitute <sub-prefix> Instead of advertising the first prefix, advertise this prefix instead.

shortcut (default|disable|enable) For use with `abr-type shortcut` ([OSPF Server Configuration](#)), this advertises the area as capable of supporting ABR shortcut behavior ([draft-ietf-ospf-shortcut-abr-02](#)).

stub [no-summary] Configure this area as a Stub Area ([Area Types](#)).

no-summary When present, the area will instead be considered a Totally Stub Area ([Area Types](#)).

virtual-link <router-id> Configures a virtual link in this area between this router and the specified router. Both this router and the target router must be ABRs, and both must have a link to this (non-backbone) area. Additionally, the virtual link must be added on both ends. This command enters `config-ospf-vlink` mode which has a subset of commands available similar to [OSPF Interface Configuration](#). The available commands are `authentication-key`, `dead-interval`, `hello-interval`, `message-digest-key`, `retransmit-interval`, and `transmit-delay`. The usage of these commands is explained in [OSPF Interface Configuration](#).

The virtual link is used to exchange routing information directly between the routers involved, and can be used to deliver traffic via the peer if necessary. Such a relationship may be necessary to nudge traffic from an ABR with a single undesirable link to another ABR with a faster link to a common remote destination, when the path would otherwise be selected because it is shorter.

OSPF Debugging Information

The following debugging commands are available in `config-ospf` mode. Messages will be logged in accordance with the settings in [Logging](#).

debug event Enable debugging information for OSPF events.

debug nssa Enable debugging information for OSPF Not-So-Stubby Area information.

debug sr Enable debugging information for OSPF Segment Routing information.

debug te Enable debugging information for OSPF Traffic Engineering information.

debug (ism|nsm) (events|status|timers) Enables State Machine debugging.

ism Enable debugging information for the Interface State Machine.

nsm Enable debugging information for the Neighbor State Machine.

For either of the above state machines, several types of debugging information are available:

events Enable event debugging for the chosen state machine.

status Enable status debugging for the chosen state machine.

timers Enable timer debugging for the chosen state machine.

debug lsa (flooding|generate|install|refresh) Enables Link State Advertisement debugging.

flooding Enables debugging for LSA flooding.

generate Enables debugging for LSA generation.

install Enables debugging for LSA installation and deletion.

refresh Enables debugging for LSA refresh.

debug packet (dd|hello|ls-acknowledgment|ls-request|ls-update) (send|recv) [detail] Enables packet-level debugging.

dd Debug database description packets.

hello Debug OSPF hello packets.

ls-acknowledgment Debug LSA acknowledgment packets.

ls-request Debug LSA request packets.

ls-update Debug LSA update packets.

Packet debugging entries are limited to a single direction:

send Debug packets sent by this router.

recv Debug packets received by this router.

Optionally, increased detail may be added to debugging messages by use of the `detail` parameter.

debug zebra (interface|redistribute) Enables OSPF-specific debugging for the dynamic routing manager daemon.

interface Debug dynamic routing manager interface information.

redistribute Debug dynamic routing manager route redistribution information.

11.3.4 OSPF Status

TNSR supports several commands to display information about the OSPF daemon configuration and its status.

See also:

For more general dynamic routing status information, see *Dynamic Routing Manager Status*

Configuration Information

To view the OSPF configuration:

```
tnsr(config)# show route dynamic ospf config
interface GigabitEthernet3/0/0
    ip ospf area 0.0.0.0
exit
interface TenGigabitEthernet6/0/0
    ip ospf area 0.0.0.0
    ip ospf cost 10
exit
router ospf
    ospf router-id 10.2.0.1
    passive-interface GigabitEthernet3/0/0
```

Status Information

To view the OSPF database:

```
tnsr(config)# show route dynamic ospf database

      OSPF Router with ID (10.2.0.1)

      Router Link States (Area 0.0.0.0)

Link ID        ADV Router      Age  Seq#           CkSum  Link count
10.2.0.1       10.2.0.1        129  0x80000005    0x6808  2
10.25.0.1      10.25.0.1       157  0x80000005    0x45ce  2

      Net Link States (Area 0.0.0.0)

Link ID        ADV Router      Age  Seq#           CkSum
203.0.113.25   10.25.0.1       158  0x80000001    0x2e80
```

Additional more specific information is available by adding a keyword onto the end of the command `show route dynamic ospf database <name>` where <name> is one of the following choices:

asbr-summary Autonomous System Boundary Router (ASBR) database summary.

external External link state information.

max-age Link State Advertisement (LSA) entries in MaxAge list.

network Network link states.

nssa-external Not-so-stubby-area external link states.

opaque-area Link area Opaque-LSA.

opaque-as Link AS Opaque-LSA.

opaque-link Link local Opaque-LSA.

router Router link states.

self-originate Link states originated from this router.

summary Network summary link states.

To view information about interfaces participating in OSPF:

```
tnsr(config)# show route dynamic ospf interface
GigabitEthernet3/0/0 is up
  ifindex 22, MTU 1500 bytes, BW 0 Mbit <UP,RUNNING>
  Internet Address 10.2.0.1/24, Broadcast 10.2.0.255, Area 0.0.0.0
  MTU mismatch detection: enabled
  Router ID 10.2.0.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DR, Priority 1
  No backup designated router on this network
  Multicast group memberships: <None>
  Timer intervals configured, Hello 10s, Dead 40s, Wait 40s, Retransmit 5
  No Hellos (Passive interface)
  Neighbor Count is 0, Adjacent neighbor count is 0
TenGigabitEthernet6/0/0 is up
  ifindex 23, MTU 1500 bytes, BW 0 Mbit <UP,RUNNING>
  Internet Address 203.0.113.2/24, Broadcast 203.0.113.255, Area 0.0.0.0
  MTU mismatch detection: enabled
  Router ID 10.2.0.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State Backup, Priority 1
  Backup Designated Router (ID) 10.2.0.1, Interface Address 203.0.113.2
  Multicast group memberships: OSPFAllRouters OSPFDesignatedRouters
  Timer intervals configured, Hello 10s, Dead 40s, Wait 40s, Retransmit 5
  Hello due in 8.281s
  Neighbor Count is 1, Adjacent neighbor count is 1
```

To view information about current OSPF neighbors and adjacencies:

```
tnsr(config)# show route dynamic ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
↔		RXmtL RqstL DBsmL			
10.25.0.1	1	Full/DR	39.774s	203.0.113.25	TenGigabitEthernet6/0/
↔0:203.0.113.2	0	0	0		

For more detailed neighbor information, use show route dynamic ospf neighbor detail.

To view information about current OSPF routes:

```
tnsr(config)# show route dynamic ospf route
===== OSPF network routing table =====
N    10.2.0.0/24          [10] area: 0.0.0.0
      directly attached to GigabitEthernet3/0/0
N    10.25.0.0/24        [20] area: 0.0.0.0
      via 203.0.113.25, TenGigabitEthernet6/0/0
N    203.0.113.0/24      [10] area: 0.0.0.0
      directly attached to TenGigabitEthernet6/0/0

===== OSPF router routing table =====
===== OSPF external routing table =====
```

To view information about this OSPF router:

```
tnsr(config)# show route dynamic ospf router-info
--- Router Information parameters ---
Router Capabilities: 0x10000000
```

To view information about all OSPF Area Border Routers (ABR) and Autonomous System Boundary Routers (ASBR):

```
tnsr(config)# show route dynamic ospf border-routers
===== OSPF router routing table =====
R      10.25.0.1          [10] area: 0.0.0.0, ABR
                        via 203.0.113.25, TenGigabitEthernet6/0/0
```

11.3.5 OSPF Terminology

OSPF has some common terms used throughout this section which can be confusing for those unfamiliar with the protocol.

Area A collection of routers inside an AS, each sharing the same area ID. An Area ID is typically formatted like an IP address in dotted quad notation, `nnn.nnn.nnn.nnn`, but can also be expressed as an unsigned 32-bit integer.

Area Border Router (ABR) A router connected to multiple areas.

Autonomous System Boundary Router (ASBR) A router connected to external networks (outside the area).

Backbone The central area of an AS, typically area `0.0.0.0`. All areas in the AS connect to the backbone through ABRs.

Cost A numeric value assigned to a link between networks, using by OSPF to calculate optimal paths to a destination. Typically higher bandwidth or higher quality circuits will be assigned a low cost, while circuits that are undesirable will be given a high cost. OSPF will prefer to use a route when it has the lowest total cost from a source to a destination.

Designated Router (DR) In a network with multiple routers, one of them will be elected as a Designated Router using Hello messages. The DR takes on the task of generating LSA messages for the network, among other special duties.

Flooding The mechanism by which OSPF routers distribute link state database information to neighbors.

Hello Special OSPF messages which introduce neighbors to each other. Using these messages, neighbors can discover each other and begin to form routing relationships.

Interior Gateway Protocol (IGP) A routing protocol, such as OSPF, which exchanges information about how to reach networks inside an autonomous system.

Link State Advertisement (LSA) Messages sent by OSPF routers which describe the state of network links, or the router itself, including information about its interfaces and other neighbors.

Link State Database (LSDB) A database containing the collected LSA messages of all routers and networks in the domain.

Link State Advertisement Message Types

LSA messages each have a type, indicating the information carried within. These types may be referenced throughout this section when describing routing behaviors.

Type 1 - Router LSA Sent by every router in an area. Contains a description of all links on the router, including their state and costs.

Type 2 - Network LSA Sent by the DR for a network. Contains a description of every router attached to the network, including the DR.

Type 3 - Network Summary-LSA Sent by ABRs. Contains a description of destinations outside the current area (inter-area) when the destination is an IP network.

Type 4 - ASBR Summary-LSA Similar to Type 3, but when sent when the destination is an ASBR.

Type 5 - AS-external LSA Sent by ASBRs. Contains a description of destinations outside of this AS. Typically each message only contains information about a single destination.

Type 6 - Multicast Group Membership LSA Not used.

Type 7 - NSSA External Link-State Advertisements Similar to Type 5, but are only exchanged inside an NSSA.

Type 8 - External attribute LSA Carry information from external routing protocols, such as BGP, when such destinations are announced with Type 5 LSAs.

Type 9 - Link Scope Opaque LSA Carries information intended for uses other than OSPF, such as available bandwidth. It is carried through to other routers without being processed by OSPF itself. Type 9 messages are for other routers on the same link.

Type 10 - Area Scope Opaque LSA Similar to Type 9, but flooded to all routers in an area.

Type 11 - AS Scope Opaque LSA Similar to Type 9, but flooded to all routers throughout the AS, except for special areas such as stubs.

Area Types

OSPF Areas can be one of several types which alter their behavior in important ways.

Normal A typical area in which all routers know all possible routes.

Stub Area An area with no external connections. Since traffic passing out of a stub area must pass through an ABR, it only needs to know about routes to the ABR, not beyond the ABR. Routers in a stub area do not receive Type 5 LSAs.

Totally Stub Area Similar to a stub area, but routers also do not receive summary LSA messages except for default route information. As such, they do not receive LSA messages of type 3, 4, or 5.

Not-so-Stubby-Area (NSSA) Similar to a Stub area but it may contain static routes to non-OSPF networks. Routers in an NSSA exchange external routing information in Type 7 LSAs instead of Type 5.

NSSA Totally Stub Area Similar to both NSSA and a Totally Stub area. As such, they do not receive LSA messages of type 3, 4, or 5.

Metric Types

Type 1 or E1 A Type 1 external metric, also known as E1, uses a similar cost calculation to typical link states, where internal and external costs are added together to find the total cost.

Type 2 or E2 A Type 2 external metric, also known as E2, only considers external costs and ignores internal costs.

11.4 Open Shortest Path First v3 (OSPF6)

Open Shortest Path First v3 (OSPF6) is defined by [RFC 5340](#) and is similar to OSPF v2, but operates with IPv6 networks. Thus, it is a link-state routing protocol that automatically locates neighboring IPv6 routers within an autonomous system, typically with multicast, and exchanges IPv6 routing information for networks each neighbor.

OSPF6 is an interior routing protocol (IGP), and facilitates routing between private links or segments of local networks.

Terms used in this section are shared with OSPF, and are covered in *OSPF Terminology*.

11.4.1 OSPF6 Required Information

Before starting, take the time to gather all of the information required to form an OSPF6 adjacency to a neighbor. This list is similar to that of OSPF. At a minimum, TNSR will need to know these items:

Local Router ID Typically the highest numbered local address on the firewall. This is also frequently set as the internal or LAN side IP address of a router. It does not matter what this ID is, so long as it is given in IPv4 address notation and does not conflict with any neighbors.

OSPF6 Area At present, the OSPF6 daemon only supports a single area. Use 0.0.0.0 when asked for an area.

OSPF6 Active Interfaces The interfaces on this router upon which the OSPF6 daemon will advertise itself and monitor for neighbors. These interfaces are connected to network segments with other routers. They may be connected to local networks or remote point-to-point links. These interfaces only require an IPv6 link local address.

OSPF6 Active Interface Cost Values OSPF6 calculates the most efficient way to route between networks based on the total cost of a path from source to destination. Less desirable links (e.g. wireless) can be given a higher cost so that paths over faster networks will be used by traffic unless the preferred path is unavailable. For single connections to other networks, this value is not necessary and may be omitted or set to a simple default such as 5 or 10.

OSPF6 Passive Interfaces These interfaces contain networks which TNSR will advertise as reachable through this router, but do not contain other routers.

The example in this section uses the following values:

Table 5: Example OSPF Configuration

Item	Value
Local Router ID	10.2.0.1
OSPF Area	0.0.0.0
Active Interfaces (Cost)	TenGigabitEthernet6/0/0 (10)
Passive Interfaces	GigabitEthernet3/0/0

11.4.2 OSPF6 Example

This example configuration implements an OSPF setup using the required information from *Example OSPF Configuration*.

```
tnsr(config)# route dynamic ospf6
tnsr(config-frr-ospf6)# server
tnsr(config-ospf6)# ospf router-id 10.2.0.1
tnsr(config-ospf6)# interface GigabitEthernet3/0/0
tnsr(config-ospf6-if)# passive
tnsr(config-ospf6-if)# area 0.0.0.0
tnsr(config-ospf6-if)# exit
tnsr(config-ospf6)# interface TenGigabitEthernet6/0/0
tnsr(config-ospf6-if)# cost outgoing 10
tnsr(config-ospf6-if)# area 0.0.0.0
tnsr(config-ospf6-if)# exit
tnsr(config-ospf6)# exit
tnsr(config-frr-ospf6)# enable
```

(continues on next page)

(continued from previous page)

```
tnsr(config-frr-ospf6) # exit
tnsr(config) #
```

A similar configuration may be applied to neighboring routers also connected to the same network as the TenGigabitEthernet6/0/0 interface. Adjust the router ID and interface names as needed.

Note: OSPF6 only supports a single area, 0.0.0.0, thus all interfaces and neighbors must use area 0.0.0.0.

11.4.3 OSPF6 Configuration

There are a number of ways to fine-tune the behavior and create complex OSPF6 routing configurations. The available configuration parameters are covered throughout this section.

Enable OSPF6

The OSPF6 service has a master enable/disable toggle that must be set before OSPF6 will operate. Enable OSPF6 using the `enable` command in `config-frr-ospf6` mode:

```
tnsr(config) # route dynamic ospf6
tnsr(config-frr-ospf6) # enable
```

To disable the service, use `no enable` or `disable`.

The OSPF6 service is managed as described in *Service Control*.

OSPF6 Server Configuration

To configure the OSPF6 server, start in `config-frr-ospf6` mode and run the `server` command:

```
tnsr(config-frr-ospf6) # server
tnsr(config-ospf6) #
```

This changes into `config-ospf6` mode, which contains the following commands:

area <area-id> Configures area-specific settings in *OSPF6 Area Configuration* mode.

Note: Currently OSPF6 only supports a single area, 0.0.0.0, but settings for that area are still configured using this command.

auto-cost reference-bandwidth <bw> A base value, in Mbit/s, which is used when OSPF6 automatically calculates cost values. The default value is 100 which means that an interface with 100Mbit/s of bandwidth or greater will have a cost of 1, with lower bandwidth values incurring higher cost values.

All routers in the same area should use the same value, otherwise automatic cost calculations would fail to accurately represent total path costs between routers.

distance [(external|inter-area|intra-area)] <dist> Sets an administrative distance for routes obtained via OSPF6. This can be configured globally as well as for specific types of OSPF6 routes.

external <dist> Sets the administrative distance for external OSPF6 routes.

inter-area <dist> Sets the administrative distance for OSPF6 routes between areas.

intra-area <dist> Sets the administrative distance for OSPF6 routes inside an area.

interface <if-name> Defines an interface for use with OSPF6, and enters *OSPF6 Interface Configuration* mode.

log-adjacency-changes [detail] Instructs the OSPF6 daemon to log changes in neighbor adjacencies. This is useful for tracking changes to neighbor relationships, especially during initial configuration.

The optional `detail` parameter increases the verbosity of the resulting log messages.

ospf router-id <router-id> Sets the router ID for the OSPF6 daemon. This is typically set to an IPv4 address unique to this router, and commonly is set to a local private address.

Note: Even though OSPF6 handles IPv6 routing, router IDs are still specified using IPv4 addresses in dotted quad notation.

redistribute <route-source> [route-map <map>] Enables redistribution of routes from another source. Available route sources are listed in *Dynamic Routing Protocol Lists*.

route-map <map> Apply the given route map to the redistributed route advertisements.

stub-router administrative Administratively declares this router as a stub router, having no external connections.

timers lsa min-arrival <min> The minimum time allowed between advertisements by neighbors, from 0–600000, in milliseconds. Default is 1000.

timers throttle spf (delayinitial-holdmaximum-hold) <val> Controls timers that determine when the router will make SPF routing decisions.

delay <val> Minimum time after an event occurs before allowing SPF calculation. Lower values will react faster to changes, but can be less stable. Specified in milliseconds from 0–600000, with a default value of 0.

initial-hold <val> Lowest time allowed between SPF calculations. Specified in milliseconds from 0–600000, with a default value of 50.

maximum-hold <val> Highest time allowed between SPF calculations. Specified in milliseconds from 0–600000, with a default value of 5000.

SPF calculations are adaptive, and if a new event occurs which would otherwise trigger a calculation before the hold timer expires, then the hold is increased by the `initial-hold` value, up to the specified `maximum-hold`. This avoids excessive consecutive recalculations.

OSPF6 Interface Configuration

OSPF6 must use one or more interfaces to announce itself to neighbors and to receive announcements from neighbors. At least one interface must be configured and active in order to locate neighbors and form an adjacency.

Warning: Outside NAT interfaces (`ip nat outside`) cannot be used as active interfaces in OSPF6! The presence of NAT prevents OSPF6 from properly communicating with neighbors to form a full adjacency.

To configure an interface for use with OSPF6, start in `config-ospf6` mode and use the `interface <if-name>` command to enter `config-ospf6-if` mode.

```
tnsr(config-ospf6)# interface <if-name>
tnsr(config-ospf6-if) #
```

config-ospf6-if mode contains the following commands:

advertise prefix-list <name> Filters route advertisements using the specified prefix list (*Dynamic Routing Prefix Lists*).

area <area-id> This command defines the interface as a member of the given area. This is required to activate an interface for use by OSPF6.

Note: OSPF6 currently only supports a single area, 0.0.0.0, thus all interfaces and neighbors must use area 0.0.0.0.

bfd enabled (true|false) Enable Bidirectional Forwarding Detection for OSPF6 on this interface.

cost outgoing <link-cost> A manual cost value to apply to this interface, rather than allowing automatic cost calculation to take place.

In situations where multiple paths are possible to the same destination, this allows OSPF6 to prefer one path over another when all else is equal.

dead-interval <time> Time, in seconds from 1-65535, without communication from a neighbor on this interface before considering it dead. This is also known as the RouterDeadInterval timer in OSPF6. Default value is 40. This timer should be set to the same value for all routers.

hello-interval <interval> The interval, in seconds from 1-65535, at which this router will send hello messages. This is also known as the HelloInterval timer in OSPF6. Default value is 10. This timer should be set to the same value for all routers.

A lower value will result in faster convergence times, but will consume more resources.

instance-id <value> An alternate OSPF6 instance identifier for this interface. Typically omitted or set to 0.

mtu <value> Explicitly configures an MTU value for this interface. This value will override the interface MTU determined automatically by the operating system. Useful in cases where the router is unable to determine the actual interface MTU, for example on virtual interfaces such as those used by IPsec.

mtu-ignore When present, OSPF6 will ignore the MTU advertised by neighbors and can still achieve a full adjacency when peers do not have matching MTU values.

network (broadcast|point-to-point) Manually configures a specific type of network used on a given interface, rather than letting OSPF6 determine the type automatically. This controls how OSPF6 behaves and how it crafts messages when using an interface.

broadcast Broadcast networks, such as typical Ethernet networks, allow multiple routers on a segment and OSPF6 can use multicast to send messages to multiple targets at once. OSPF6 assumes that all routers on broadcast networks are directly connected and can communicate without passing through other routers.

point-to-point A point-to-point network links a single pair of routers. The interface is still capable of broadcast, and OSPF6 will dynamically discover neighbors. With this type of network, OSPF6 disables election of a DR.

passive Configures this interface as passive. This prevents the interface from actively participating in OSPF6, while still allowing OSPF6 to operate on networks connected to that interface. This is commonly used for local interfaces without other routers attached. OSPF6 will announce networks attached to passive interfaces as stub links.

priority <priority> A priority value, from 0–255, assigned to this router. When determining which router will become the Designated Router (DR), the router with the highest priority is more likely to be elected as the DR.

The default value is 1. The value 0 is special and will prevent this router from being chosen as DR.

retransmit-interval <interval> The interval, in seconds from 1–65535, at which this router will retransmit Link State Request and Database Description messages. This is also known as the `RxmtInterval` timer in OSPF6. Default value is 5.

transmit-delay <delay> The interval, in seconds from 1–65535, at which this router will transmit LSA messages. This is also known as the `InfTransDelay` timer in OSPF6. Default value is 1.

OSPF6 Area Configuration

To configure area-specific settings in OSPF6, start in `config-ospf6` mode and use the `area <area-id>` command to enter `config-ospf6-area` mode.

Note: Currently OSPF6 only supports a single area, 0.0.0.0, but settings for that area are still configured using this command.

```
tnsr(config-ospf6)# area <area-id>
tnsr(config-ospf6-area)#
```

`config-ospf6-area` mode contains the following commands:

range <prefix> [cost <val>|not-advertise] Configure summarization of routes inside the given prefix. Instead of Type 1 (Router) and Type 2 (Network) LSAs, it creates Type 3 Summary LSAs instead.

cost <val> Apply the specified cost to summarized routes for this prefix.

not-advertise Disable advertisement for this prefix.

OSPF6 Debugging Information

The following debugging commands are available in `config-ospf6` mode. Messages will be logged in accordance with the settings in [Logging](#).

debug abr Enables debugging for Area Border Routers.

debug asbr Enables debugging for Autonomous System Boundary Routers.

debug flooding Enables debugging for Link State Advertisement flooding.

debug interface Enables debugging for OSPF6 interfaces.

debug border-routers (area <area-id>|router <router-id>) Enables debugging for specific border routers, either by area or router ID.

debug lsa <lsa-type> <event-type> Enables LSA message debugging.

LSA Type Specifies a type of LSA message to debug, which must be one of the following: `as-external`, `inter-prefix`, `inter-router`, `intra-prefix`, `link`, `network`, `router`, `unknown`. These message types are described further in [Link State Advertisement Message Types](#).

Event Type Specifies when to log debug information for the specified type of LSA message.

examine Enables debugging when examining LSA messages.

flooding Enables debugging when flooding LSA messages.

originate Enables debugging when originating LSA messages.

debug message (ddlhellolls-acknowledgmentlls-requestlls-update|unknown) (recv|send) Enables packet-level debugging of OSPF6 messages.

dd Debug database description packets.

hello Debug OSPF6 hello packets.

ls-acknowledgment Debug LSA acknowledgment packets.

ls-request Debug LSA request packets.

ls-update Debug LSA update packets.

unknown Debug OSPF6 messages of unknown types.

Message debugging entries are limited to a single direction:

send Debug messages sent by this router.

recv Debug messages received by this router.

debug neighbor event Enable debugging information for OSPF6 neighbor events.

debug neighbor state Enable debugging information for OSPF6 neighbor state changes.

debug route [(inter-areal|intra-areal|table)] Enables debugging for OSPF6 route calculations.

debug route memory Enables debugging for OSPF6 route table memory usage.

debug spf (database|process|time) Debug SPF calculations

database Enable debugging for LSA message counts during SPF calculation

process Enable detailed debugging of the SPF calculation process.

time Enable debugging for SPF calculation timing.

debug zebra [(recv|send)] Enables OSPF6-specific debugging for dynamic routing manager daemon messages, in the send or receive direction, or both when the direction is omitted.

11.4.4 OSPF6 Status

TNSR supports several commands to display information about the OSPF6 daemon configuration and its status.

See also:

For more general dynamic routing status information, see *Dynamic Routing Manager Status*

Configuration Information

To view the OSPF6 configuration:

```
tnsr# show route dynamic ospf6 config
interface GigabitEthernet3/0/0
    ipv6 ospf6 passive

interface TenGigabitEthernet6/0/0
    ipv6 ospf6 cost 10
```

(continues on next page)

(continued from previous page)

```

router ospf6
  ospf6 router-id 10.2.0.1
  interface GigabitEthernet3/0/0 area 0.0.0.0
  interface TenGigabitEthernet6/0/0 area 0.0.0.0

```

Status Information

To view the OSPF6 database:

```

tnsr# show route dynamic ospf6 database

      Area Scoped Link State Database (Area 0.0.0.0)

Type LSId          AdvRouter      Age   SeqNum          Payload
Rtr  0.0.0.0        10.2.0.1        146  80000002        10.27.0.1/0.0.0.13
Rtr  0.0.0.0        10.27.0.1       147  80000002        10.27.0.1/0.0.0.13
Net  0.0.0.13       10.27.0.1       147  80000001        10.27.0.1
Net  0.0.0.13       10.27.0.1       147  80000001        10.2.0.1
INP  0.0.0.0        10.2.0.1       146  80000003        2001:db8:f0::/64
INP  0.0.0.0        10.27.0.1      147  80000003        2001:db8:f2::/64
INP  0.0.0.13       10.27.0.1      147  80000001        2001:db8:0:2::/64

      I/F Scoped Link State Database (I/F GigabitEthernet3/0/0 in Area 0.0.0.0)

Type LSId          AdvRouter      Age   SeqNum          Payload
Lnk  0.0.0.14       10.2.0.1       187  80000001        fe80::290:bff:fe7a:8a65

      I/F Scoped Link State Database (I/F TenGigabitEthernet6/0/0 in Area 0.0.0.0)

Type LSId          AdvRouter      Age   SeqNum          Payload
Lnk  0.0.0.15       10.2.0.1       187  80000001        fe80::290:bff:fe7a:8a67
Lnk  0.0.0.13       10.27.0.1      192  80000001        fe80::290:bff:fe7a:87c1

      AS Scoped Link State Database

Type LSId          AdvRouter      Age   SeqNum          Payload

```

To view information about interfaces participating in OSPF6:

```

tnsr# show route dynamic ospf6 interface
GigabitEthernet3/0/0 is up, type BROADCAST
  Interface ID: 14
  Internet Address:
    inet : 10.2.0.1/24
    inet6: 2001:db8:f0::1/64
    inet6: fe80::290:bff:fe7a:8a65/128
  Instance ID 0, Interface MTU 1500 (autodetect: 1500)
  MTU mismatch detection: enabled
  Area ID 0.0.0.0, Cost 100
  State DR, Transmit Delay 1 sec, Priority 1
  Timer intervals configured:
    Hello 10, Dead 40, Retransmit 5
  DR: 10.2.0.1 BDR: 0.0.0.0
  Number of I/F scoped LSAs is 1

```

(continues on next page)

(continued from previous page)

```

    0 Pending LSAs for LSUpdate in Time 00:00:00 [thread off]
    0 Pending LSAs for LSAck in Time 00:00:00 [thread off]
TenGigabitEthernet6/0/0 is up, type BROADCAST
  Interface ID: 15
  Internet Address:
    inet : 203.0.113.2/24
    inet6: 2001:db8:0:2::2/64
    inet6: fe80::290:bff:fe7a:8a67/128
  Instance ID 0, Interface MTU 1500 (autodetect: 1500)
  MTU mismatch detection: enabled
  Area ID 0.0.0.0, Cost 100
  State BDR, Transmit Delay 1 sec, Priority 1
  Timer intervals configured:
    Hello 10, Dead 40, Retransmit 5
  DR: 10.27.0.1 BDR: 10.2.0.1
  Number of I/F scoped LSAs is 2
    0 Pending LSAs for LSUpdate in Time 00:00:00 [thread off]
    0 Pending LSAs for LSAck in Time 00:00:00 [thread off]
TenGigabitEthernet6/0/1 is down, type BROADCAST
  Interface ID: 16
  OSPF not enabled on this interface
TenGigabitEthernet8/0/0 is down, type BROADCAST
  Interface ID: 17
  OSPF not enabled on this interface

```

To view information about current OSPF neighbors and adjacencies:

```

tnsr# show route dynamic ospf6 neighbor
Neighbor ID      Pri    DeadTime    State/IfState    Duration I/F[State]
10.27.0.1        1      00:00:33     Full/DR          00:04:41 TenGigabitEthernet6/
↪0/0[BDR]

```

For more detailed neighbor information, use show route dynamic ospf6 neighbor detail.

```

tnsr# show route dynamic ospf6 neighbor detail
Neighbor 10.27.0.1%TenGigabitEthernet6/0/0
  Area 0.0.0.0 via interface TenGigabitEthernet6/0/0 (ifindex 15)
  His IfIndex: 13 Link-local address: fe80::290:bff:fe7a:87c1
  State Full for a duration of 00:04:58
  His choice of DR/BDR 10.27.0.1/10.2.0.1, Priority 1
  DbDesc status: Slave SeqNum: 0xb7380c00
  Summary-List: 0 LSAs
  Request-List: 0 LSAs
  Retrans-List: 0 LSAs
    0 Pending LSAs for DbDesc in Time 00:00:00 [thread off]
    0 Pending LSAs for LSReq in Time 00:00:00 [thread off]
    0 Pending LSAs for LSUpdate in Time 00:00:00 [thread off]
    0 Pending LSAs for LSAck in Time 00:00:00 [thread off]

```

To view information about current OSPF6 routes:

```

tnsr# show route dynamic ospf6 route-table
*N IA 2001:db8:0:2::/64      ::      TenGigabitEthernet_
↪00:05:37
*N IA 2001:db8:f0::/64      ::      GigabitEthernet3_
↪00:06:17
*N IA 2001:db8:f2::/64      fe80::290:bff:fe7a:87c1  TenGigabitEthernet_
↪00:05:32

```

(continues on next page)

(continued from previous page)

To view information about current OSPF6 border routers:

```
tnsr# show route dynamic ospf6 border-routers
Router-ID      Rtr-Bits Options      Path-Type  Area
```

To view information about the OSPF6 area:

```
tnsr# show route dynamic ospf6 area
+-10.2.0.1 [0]
  +-10.27.0.1 Net-ID: 0.0.0.13 [100]
    +-10.27.0.1 [100]
```

To view OSPF6 link state information:

```
tnsr# show route dynamic ospf6 linkstate

      SPF Result in Area 0.0.0.0

Destination: 10.2.0.1
Destination type: Linkstate
Installed Time: 00:07:10 ago
  Changed Time: 00:07:10 ago
Lock: 2 Flags: BA--
Memory: prev: (nil) this: 0x23fc980 next: 0x23fd140
Associated Area: 0.0.0.0
Path Type: Intra-Area
LS Origin: Router Id: 0.0.0.0 Adv: 10.2.0.1
Options: --|R|-|--|E|V6
Router Bits: -----
Prefix Options: xxx
Metric Type: 1
Metric: 0 (0)
Paths count: 0
Nexthop count: 0
Nexthop:

Destination: 10.27.0.1
Destination type: Linkstate
Installed Time: 00:07:10 ago
  Changed Time: 00:07:10 ago
Lock: 2 Flags: BA--
Memory: prev: 0x23fc980 this: 0x23fd140 next: 0x23de700
Associated Area: 0.0.0.0
Path Type: Intra-Area
LS Origin: Router Id: 0.0.0.0 Adv: 10.27.0.1
Options: --|R|-|--|E|V6
Router Bits: -----
Prefix Options: xxx
Metric Type: 1
Metric: 100 (1)
Paths count: 0
Nexthop count: 1
Nexthop:
  fe80::290:bff:fe7a:87c1 TenGigabitEthern
```

(continues on next page)

(continued from previous page)

```
Destination: 10.27.0.1 Net-ID: 0.0.0.13
Destination type: Linkstate
Installed Time: 00:07:10 ago
  Changed Time: 00:07:10 ago
Lock: 2 Flags: BA--
Memory: prev: 0x23fd140 this: 0x23de700 next: (nil)
Associated Area: 0.0.0.0
Path Type: Intra-Area
LS Origin: Network Id: 0.0.0.13 Adv: 10.27.0.1
Options: --|R|--|E|V6
Router Bits: -----
Prefix Options: xxx
Metric Type: 1
Metric: 100 (0)
Paths count: 0
Nexthop count: 1
Nexthop:
  :: TenGigabitEthernet
```

To view SPF calculation information:

```
tnsr# show route dynamic ospf6 spf
+-10.2.0.1 [0]
  +-10.27.0.1 Net-ID: 0.0.0.13 [100]
    +-10.27.0.1 [100]
```

11.5 Routing Information Protocol (RIP)

RIP is a simple interior routing protocol (IGP), and facilitates routing between private links or segments of local networks. It is a distance vector routing protocol, informing neighbors of known routes, gateways, and hop counts to destinations.

TNSR supports both RIPv1 ([RFC 1058](#)) and RIPv2 ([RFC 1723](#)).

RIP is widely supported and simple, but lacks the speed, efficiency, or capabilities of more powerful routing protocols such as BGP or OSPF.

11.5.1 RIP Required Information

Though RIP is a simple routing protocol, there are a few values that must be determined before a working configuration is possible. More information about these values can be found in [RIP Server Configuration](#).

RIP Version The version of RIP utilized by TNSR must be set to either 1 or 2. This value must match the version used by other connected routers.

Network(s) The subnet(s) for which routes will be advertised by RIP. Note that this value is not used directly, but is used to locate active subnets on interfaces which match.

Active Interface(s) Interfaces participating in RIP, connected to a segment with other routers also running RIP. This is optional if the subnet of the interface is also covered by the **Network(s)** value.

Neighbor(s) Neighboring router(s) running RIP with which TNSR will exchange routes. RIP will find neighbors automatically, but it is helpful to know which neighbors to look for when troubleshooting.

The example in this section uses the following values:

Table 6: Example RIP Configuration

Item	Value
RIP Version	2
Network(s)	10.2.0.0/16
Active Interface(s)	TenGigabitEthernet6/0/0
Neighbor(s)	203.0.113.27

11.5.2 RIP Example

This example configuration implements a RIP setup using the required information from *Example RIP Configuration*.

```
tnsr(config)# route dynamic rip
tnsr(config-frr-rip)# server
tnsr(config-rip)# version 2
tnsr(config-rip)# network prefix 10.2.0.0/16
tnsr(config-rip)# network interface TenGigabitEthernet6/0/0
tnsr(config-rip)# exit
tnsr(config-frr-rip)# enable
tnsr(config-frr-rip)# exit
tnsr(config)#
```

A similar configuration may be applied to neighboring routers also connected to the same network as the TenGigabitEthernet6/0/0 interface. Adjust the networks, neighbors, and interface names as needed.

11.5.3 RIP Configuration

RIP behavior can be customized in several ways, including features such as authentication. The available configuration parameters are covered throughout this section.

Enable RIP

The RIP service has a master enable/disable toggle that must be set before RIP will operate. Enable RIP using the `enable` command in `config-frr-rip` mode:

```
tnsr(config)# route dynamic rip
tnsr(config-frr-rip)# enable
```

To disable the service, use `no enable` or `disable`.

The RIP service is managed as described in *Service Control*.

RIP Server Configuration

To configure the RIP service, start in `config-frr-rip` mode and run the `server` command:

```
tnsr(config-frr-rip)# server
tnsr(config-rip)#
```

This changes into `config-rip` mode, which contains the following commands:

allow-ecmp Allow equal cost multi-path routing, where the same destination is reachable through multiple routers.

default-information originate Transmit default route information to RIP neighbors.

distance default <value> Administratively sets the default distance to the given value (1–255).

distance <prefix> distance <value> [access-list <acl-name>] Sets custom distance values for specific network prefixes.

prefix The prefix for which this distance is set.

distance <value> The distance value to advertise for this prefix.

access-list <acl-name> An optional access list used to filter this distance based on specific subnets or addresses inside the given prefix.

distribution-list interface <interface> (access-list|prefix-list) (in|out) <name> Applies either the given access list or prefix list to routes distributed from networks on the specified interface. This allows control over which routes will be distributed by RIP to neighbors.

interface <interface> The interface which is the source of routes filtered by this directive. May be * or a specific interface name.

access-list (in|out) <name> An access list to filter against in the specified direction.

prefix-list (in|out) <name> A prefix list to filter against in the specified direction.

interface <if-name> Defines an interface for use with RIP, and enters *RIP Interface Configuration* mode.

key-chain <name> Defines a key chain for use with RIP, and enters *RIP Keychain Configuration* mode.

neighbor <ip4-address> Defines the address of a neighboring router with which TNSR will exchange routes using RIP. When a neighbor is defined in this manner, RIP will always transmit to the neighbor even on passive interfaces.

network (interface <if-name>|prefix <prefix>) Defines which networks will have routes distributed by RIP to neighbors. These can be specified by interface or prefix.

Note: These values are not used directly, but are used by RIP to locate active subnets which match the given interface or prefix.

interface <if-name> Advertise routes for networks directly connected to the given interface.

prefix <prefix> Advertise routes for active networks matching the given prefix.

For example, if 10.2.0.0/16 is given and 10.2.0.0/24 and 10.2.1.0/24 are both present on active interfaces, then those two prefixes will be advertised to neighbors, not 10.2.0.0/16.

offset-list <interface> (in|out) <acl-name> Modifies RIP metrics using access-lists.

interface The interface on which metrics will be adjusted. May be * or a specific interface name.

(in|out) The direction in which modifications are made.

in Modify route metrics received from RIP neighbors.

out Modify route metrics advertised to RIP neighbors.

acl-name The name of the access list used to apply metric changes.

passive-interface <interface> [<ip4-address>] Controls whether or not RIP will transmit multicast or unicast packets on interfaces. RIP messages are always accepted in passive mode, and RIP messages are always transmitted to defined neighbors.

Warning: When the default value is set to passive, the meaning of this list is inverted. Instead of specifying passive interfaces, the list defines non-passive interfaces instead.

interface Interface to configure as passive. May be default or a specific interface name.

ip4-address A specific IP address to configure as passive on the given interface.

redistribute <route-source> [(metric <value>|route-map <name>)] Enables redistribution of routes from another source. Available route sources are listed in *Dynamic Routing Protocol Lists*.

metric <val> Advertise the route as having the given metric.

route-map <map> Apply the given route map to the redistributed route advertisements.

route prefix <ip4-prefix> Creates a static route in RIP for the given prefix, which is advertised to neighbors as reachable through this router.

route-map-filter interface <interface> (in|out) route-map <name> Apply a route-map to RIP routes. See *Dynamic Routing Route Maps* for more information on route maps.

interface The interface on which this route-map will be applied. May be default or a specific interface name.

(in|out) The direction in which the route-map will be applied to routes.

route-map <name> The route-map to apply.

timers (garbage-collection|table-update|timeout) <value> Adjust timer values for RIP. Each timer is specified in seconds and can be set to a value from 5–2147483647.

table-update How often RIP will transmit a copy of its route table to neighbors. Default is 30 seconds.

timeout How long RIP will wait before a route is no longer considered valid after receiving an advertisement. Default is 180 seconds.

garbage-collection The time to wait before removing an invalid route from the routing table. Default is 120 seconds.

For example, if a neighbor stops advertising a route or loses connectivity, then advertisements for that route will no longer be received. The route will eventually reach the timeout value since it is no longer seen in advertisements. Once it reaches the timeout value without an advertisement, it is flagged as invalid. Then once it has been invalid for long enough to reach the garbage collection age, it is removed from the routing table.

The lowest amount of time a route can be in the table while invalid is timeout + garbage-collection, which by default is 180+120 or 300 seconds (5 minutes). The longest time would be that value plus the update time, in this case, 330 seconds total.

version (1|2) The RIP version to use when communicating with RIP neighbors.

1 RIP as described in [RFC 1058](#). An older version of the protocol which utilizes class-based routing (e.g. Class A, Class B, etc) and does not support subnetting or authentication. RIP v1 sends updates using broadcast messages which must be processed by every node on connected segments.

2 RIP as described in [RFC 1723](#). An updated version of the protocol which uses classless routing (CIDR), authentication. RIP v2 sends messages using multicast, allowing only

interested routers to receive the messages by joining the appropriate multicast group (224.0.0.9).

RIP Interface Configuration

In basic configurations, RIP will automatically determine which interfaces to use. However, the interface behavior can be tuned when necessary.

To configure settings for RIP interfaces, start in `config-rip` mode and use the `interface <if-name>` command to enter `config-rip-if` mode.

```
tnsr(config-rip)# interface <if-name>
tnsr(config-rip-if)#
```

`config-rip-if` mode contains the following commands:

authentication mode <mode> [auth-length <type>] Configures RIPv2 authentication for this interface. When authentication is enabled, TNSR will ignore updates from unauthenticated peers, including RIPv1 peers.

Note: Updates from unauthenticated peers are ignored, but requests for routes from unauthenticated peers are still honored.

mode <mode> Selects the authentication mode.

md5 MD5-based HMAC authentication, which is more secure than plain text. Keys for MD5 authentication are configured with the `key-chain` command in `config-rip` mode (*RIP Keychain Configuration*).

text Insecure plain text password authentication. The password is set with the `authentication string` command in this mode.

auth-length <type> Configures the expected length of the authentication data.

rfc RFC-compatible data length (16 bytes).

old-ripd Obsolete `ripd` length (20 bytes), compatible only with older `ripd` implementations.

authentication key-chain <name> The name of a key-chain to use with MD5 authentication (*RIP Keychain Configuration*).

authentication string <auth-string> The string used for plain text authentication. Must be less than 16 characters.

receive version (1|2|both) Configures the RIP versions allowed to be received by TNSR on this interface.

send version (1|2|both) Configures the RIP versions TNSR will transmit on this interface.

split-horizon [poisoned-reverse] Prevents a route from being advertised back to the interface through which it was received. This technique helps to prevent routing loops.

poisoned-reverse Instead of preventing such routes from being advertised, this option causes RIP to actively advertise the networks as unreachable by setting the metric to 16. This is more proactive for preventing routing loops, but the primary drawback is that this does not scale well, due to the size increase of advertisements.

v2-broadcast When set, TNSR will transmit RIPv2 updates using broadcast on this interface instead of using multicast.

RIP Keychain Configuration

Key chains are used for MD5-based HMAC authentication, configured in `config-rip-if` mode (*RIP Interface Configuration*).

To configure Keychain settings in RIP, start in `config-rip` mode and use the `key-chain <name>` command to enter `config-rip-key-chain` mode.

```
tnsr(config-rip)# key-chain <name>
tnsr(config-rip-key-chain)#
```

`config-rip-key-chain` mode contains the following commands:

key <key-id> string <key-string>

key-id A numeric identifier for the key within this chain, can be any integer value from 0-2147483647.

key-string A string containing the contents of the key. This string must match between all nodes using this key-chain.

RIP Debugging Information

The following debugging commands are available in `config-rip` mode. Messages will be logged in accordance with the settings in *Logging*.

debug events Enable debugging information for RIP events.

debug zebra Enables RIP-specific debugging for the dynamic routing manager daemon.

debug packet (send|recv) Enables packet-level RIP debugging.

Packet debugging entries can be limited to a single direction:

send Debug packets sent by this router.

recv Debug packets received by this router.

11.5.4 RIP Status

TNSR supports several commands to display information about the RIP daemon configuration and its status.

See also:

For more general dynamic routing status information, see *Dynamic Routing Manager Status*

Configuration Information

To view the RIP configuration:

```
tnsr(config)# show route dynamic rip config
router rip
  version 2
  network TenGigabitEthernet6/0/0
  network 10.2.0.0/16
exit
```


Status Information

To view the RIP routing database:

```
tnsr(config)# show route dynamic rip
Codes: R - RIP, C - connected, S - Static, O - OSPF, B - BGP
Sub-codes:
      (n) - normal, (s) - static, (d) - default, (r) - redistribute,
      (i) - interface
```

	Network	Next Hop	Metric From	Tag Time
C(i)	10.2.0.0/24	0.0.0.0	1 self	0
C(i)	10.2.8.0/24	0.0.0.0	1 self	0
C(i)	10.2.222.0/24	0.0.0.0	1 self	0
R(n)	10.27.0.0/24	203.0.113.27	2 203.0.113.27	0 02:46
R(n)	10.27.8.0/24	203.0.113.27	2 203.0.113.27	0 02:46
C(i)	203.0.113.0/24	0.0.0.0	1 self	0

To view the RIP status:

```
tnsr(config)# show route dynamic rip status
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 12 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing:
  Default version control: send version 2, receive version 2
    Interface      Send  Recv  Key-chain
    GigabitEthernet3/0/02      2
    TenGigabitEthernet6/0/02    2
    ipip1      2      2
  Routing for Networks:
    10.2.0.0/16
    TenGigabitEthernet6/0/0
  Routing Information Sources:
    Gateway      BadPackets  BadRoutes  Distance  Last Update
    203.0.113.27      0          0         120      00:00:12
  Distance: (default is 120)
```

11.5.5 RIP Tuning

11.6 Dynamic Routing Protocol Lists

Throughout dynamic routing, certain commands accept parameters which specify a supported routing protocol or source of routes. Currently, the following values are valid in these parameters:

- connected** Routes for directly connected networks
- kernel** Routes from the kernel
- system** Routes from system configuration
- bgp** Routes obtained dynamically from BGP neighbors
- ospf** IPv4 routes obtained dynamically from OSPF neighbors

ospf6 IPv6 routes obtained dynamically from OSPF6 neighbors

VIRTUAL ROUTER REDUNDANCY PROTOCOL

Virtual Router Redundancy Protocol (VRRP) is a protocol which allows routers to coordinate control of IP addresses between multiple nodes acting as a single “virtual” router cluster. Multiple nodes coordinating control in this way allows for redundancy, where a single node failing does not adversely affect traffic passing through the virtual router.

The specific version of VRRP used by TNSR is VRRPv3 as defined in [RFC 5798](#), but will be referred to as “VRRP” throughout this document.

With VRRP, one router acts as the primary master node and additional routers act as backup nodes. Commonly there are only two routers in a cluster: A primary node and a secondary node. VRRP supports additional nodes if a use case calls for increased redundancy.

Addresses configured on the primary node are defined as virtual router (VR) addresses on all participating nodes, including the primary node which is considered the owner of the VR addresses. The VR addresses are then used as next hop gateways by peers, rather than traditional addresses. This includes delivery of routed subnets from upstream sources as well as acting as a gateway for local clients. Since peers communicate with the shared virtual addresses, when a failure occurs communications will continue through whichever node is elected master of the VR addresses. This allows traffic to flow with little to no interruption when a node fails.

The current master of VR addresses is determined by an election process. The election process considers the priority value for the VR address on each node first, among other factors. The owner of the VR addresses has the highest possible priority, 255, and additional nodes will have a lower priority from 1–254 (e.g. 100).

Participating nodes advertise their state to peers and listen for these advertisements from peers. Typically only the current master will transmit advertisements. If other nodes fail to see advertisements from a higher priority node in a timely manner defined by the settings, control of the virtual address is assumed by the backup node with the next highest priority. This state information is transmitted via multicast on a local segment, to a multicast destination of 224.0.0.18 for IPv4 and ff02::12 for IPv6.

Warning: Switches (physical or virtual) must allow the multicast advertisements to flow freely. Ensure switch features such as storm control or rate limiting are relaxed or disabled on ports participating in VRRP.

At layer 2, VRRP works by enabling the nodes to essentially share a single MAC address. This MAC address is derived from the ID of the associated VR address, with the form of 00:00:5E:00:01:<id> where <id> is the VR ID (1–255) in hexadecimal. The multicast advertisements from the current master allow compatible switches to direct traffic to the correct port, so that the current master receives traffic destined for the MAC address associated with the VR address being advertised.

Warning: In virtual environments, special switch and VM configuration settings may be required to allow VRRP to function. This settings may include, but are not limited to: vSwitch or VM port promiscuous mode, allowing forged transmits, and allowing MAC address changes. These are necessary for TNSR to properly send and receive not only the VRRP advertisements, but also for traffic using the shared VRRP MAC address.

12.1 VRRP Compatibility

Currently VRRP is only compatible with routed deployments.

12.1.1 VRRP Hardware Compatibility

VRRP requires network interface hardware on which DPDK PMDs support programming an additional MAC address. Without this capability, the interface cannot receive traffic addressed to the VRRP MAC address.

The following DPDK PMDs are supported:

- em
- fm10k
- i40e
- iavf
- ice
- igb
- ixgbe
- mlx4
- mlx5
- virtio

12.1.2 VRRP and NAT

VRRP may not be used on interfaces involved in outbound NAT when the VR priority is 255. Currently there is an interaction between NAT and VRRP in this case which leads to both nodes failing to receive and process VRRP advertisements from peers. When NAT is present on outbound NAT interfaces, use a lower priority value. Conflicting configurations will be rejected by input validation.

12.1.3 VRRP and Reflect ACLs

As there is not yet a method for VRRP cluster nodes to share state data, using `reflect` type ACLs may result in active connections being dropped when control is transferred between cluster nodes. New connections may be made immediately.

12.1.4 VRRP and AWS/Azure

Currently VRRP does not support unicast peers for routed environments such as AWS and Azure. This functionality will be added in a future release.

12.2 VRRP Example

This example is a basic two-node VRRP cluster with one node as the owner of an internal and external VR address, and the other as a backup. This is a routed configuration with a statically routed subnet used for the internal LAN.

In this example, the upstream ISP will deliver a routed subnet (198.51.100.0/24) to the WAN-side VR address (203.0.113.2), and internal clients will use the LAN-side VR address (198.51.100.1) as their gateway.

12.2.1 Diagram

12.2.2 Required Information

These tables contain all required information to configure the cluster.

The information in this first table is related to the setup in general, not a specific cluster node.

Table 1: Example Basic VRRP Configuration Related Information

Item	Value
Upstream Gateway	203.0.113.1
Routed Subnet	198.51.100.0/24
LAN Client Gateway	198.51.100.1

This information is for the primary node, which in this example is called R1.

Table 2: Example Basic VRRP Configuration for R1

Item	Value
R1 WAN Interface	TenGigabitEthernet6/0/0
R1 WAN IP Address	203.0.113.2/24
R1 WAN VR ID	220
R1 WAN VR Address	203.0.113.2
R1 WAN VR Priority	255 (Owner)
R1 LAN Interface	TenGigabitEthernet6/0/1
R1 LAN IP Address	198.51.100.1/24
R1 LAN VR ID	210
R1 LAN VR Address	198.51.100.1
R1 LAN VR Priority	255 (Owner)

This information is for the secondary node, which in this example is called R2. Note that the interface addresses are different than R1, but the same VR address is used.

Table 3: Example Basic VRRP Configuration for R2

Item	Value
R2 WAN Interface	TenGigabitEthernet6/0/0
R2 WAN IP Address	203.0.113.3/24
R2 WAN VR ID	220
R2 WAN VR Address	203.0.113.2
R2 WAN VR Priority	100
R2 LAN Interface	TenGigabitEthernet6/0/1
R2 LAN IP Address	198.51.100.2/24
R2 LAN VR ID	210
R2 LAN VR Address	198.51.100.1
R2 LAN VR Priority	100

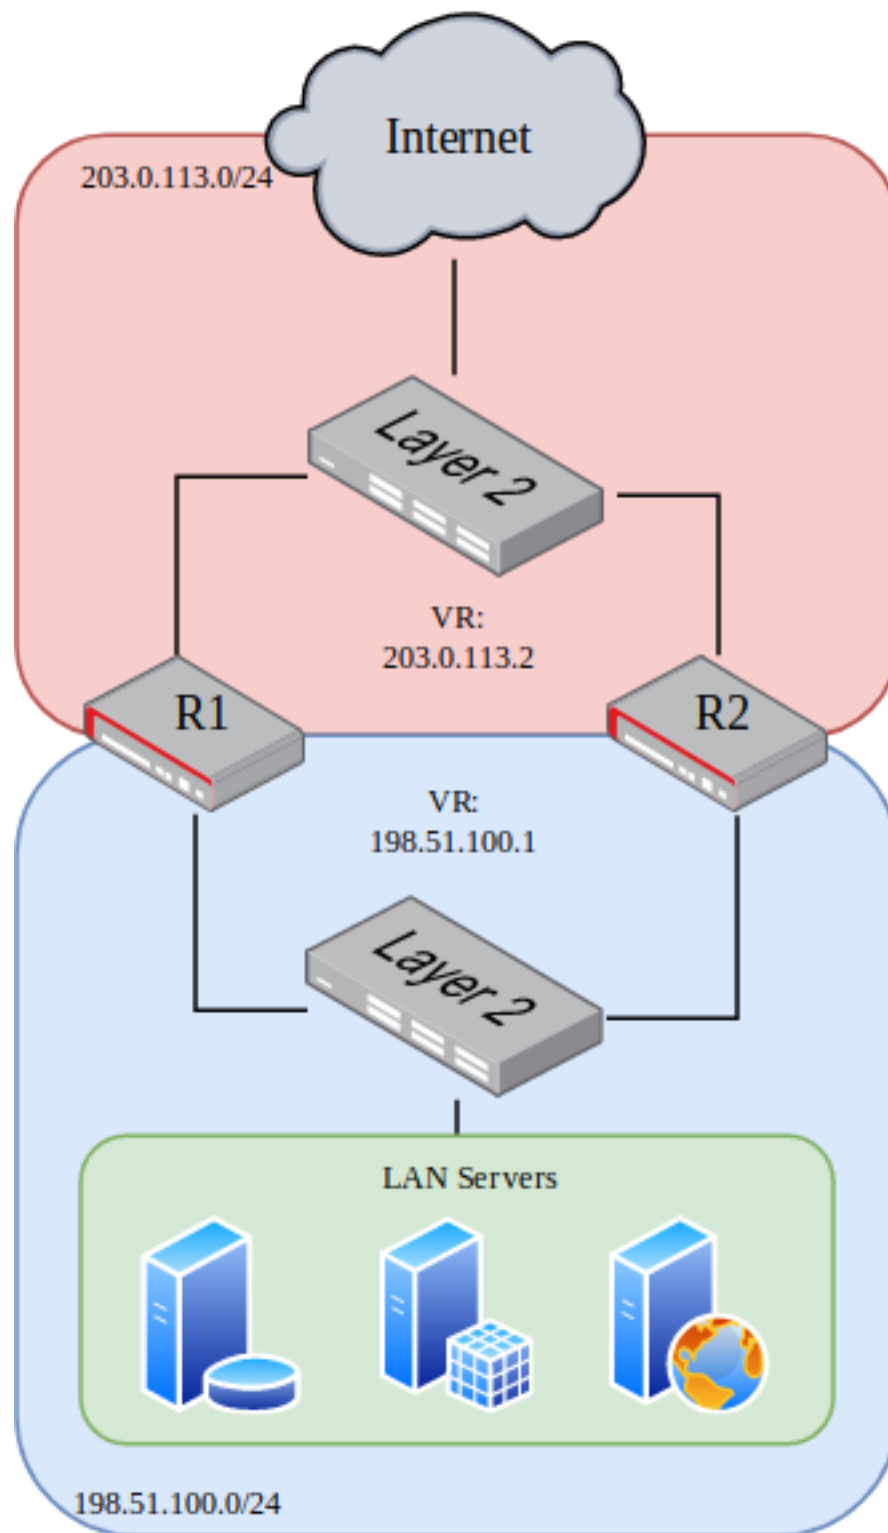


Fig. 1: VRRP Example Diagram

12.2.3 Example Configuration

The configuration commands in this section show how the settings from the table above are applied to each node. Some additional VRRP settings are shown in the commands but not the tables, but they are using the default values, shown for emphasis.

First, configure the R1 WAN interface:

```
r1 tnsr(config)# int TenGigabitEthernet6/0/0
r1 tnsr(config-interface)# description WAN
r1 tnsr(config-interface)# ip address 203.0.113.2/24
r1 tnsr(config-interface)# ip vrrp-virtual-router 220
r1 tnsr(config-vrrp4)# preempt true
r1 tnsr(config-vrrp4)# priority 255
r1 tnsr(config-vrrp4)# v3-advertisement-interval 100
r1 tnsr(config-vrrp4)# virtual-address 203.0.113.2
r1 tnsr(config-vrrp4)# exit
r1 tnsr(config-interface)# exit
r1 tnsr(config)#
```

Next, configure the R1 LAN interface:

```
r1 tnsr(config)# int TenGigabitEthernet6/0/1
r1 tnsr(config-interface)# description LAN
r1 tnsr(config-interface)# ip address 198.51.100.1/24
r1 tnsr(config-interface)# ip vrrp-virtual-router 210
r1 tnsr(config-vrrp4)# preempt true
r1 tnsr(config-vrrp4)# priority 255
r1 tnsr(config-vrrp4)# v3-advertisement-interval 100
r1 tnsr(config-vrrp4)# virtual-address 198.51.100.1
r1 tnsr(config-vrrp4)# exit
r1 tnsr(config-interface)# exit
r1 tnsr(config)#
```

R1 is now complete. Move on to the R2 WAN interface:

```
r2 tnsr(config)# int TenGigabitEthernet6/0/0
r2 tnsr(config-interface)# description WAN
r2 tnsr(config-interface)# ip address 203.0.113.3/24
r2 tnsr(config-interface)# ip vrrp-virtual-router 220
r2 tnsr(config-vrrp4)# preempt true
r2 tnsr(config-vrrp4)# priority 100
r2 tnsr(config-vrrp4)# v3-advertisement-interval 100
r2 tnsr(config-vrrp4)# virtual-address 203.0.113.2
r2 tnsr(config-vrrp4)# exit
r2 tnsr(config-interface)# exit
r2 tnsr(config)#
```

Finally, configure the R2 LAN interface:

```
r2 tnsr(config)# int TenGigabitEthernet6/0/1
r2 tnsr(config-interface)# description LAN
r2 tnsr(config-interface)# ip address 198.51.100.2/24
r2 tnsr(config-interface)# ip vrrp-virtual-router 210
r2 tnsr(config-vrrp4)# preempt true
r2 tnsr(config-vrrp4)# priority 100
r2 tnsr(config-vrrp4)# v3-advertisement-interval 100
r2 tnsr(config-vrrp4)# virtual-address 198.51.100.1
```

(continues on next page)

(continued from previous page)

```
r2 tnsr(config-vrrp4)# exit
r2 tnsr(config-interface)# exit
r2 tnsr(config)#
```

At this point, the interface and VRRP configuration is complete for both nodes.

LAN clients in 198.51.100.0/24 can use the LAN VR address of 198.51.100.1 as their default gateway.

12.3 VRRP Configuration

VRRP is configured on a per-interface basis from within `config-interface` mode. To define a new VR address, use `ip vrrp-virtual-router <vrid>` for IPv4 or `ipv6 vrrp-virtual-router <vrid>` for IPv6 when configuring an interface.

The `<vrid>` must be an integer from 1–255. This identifier must be identical for all nodes in the same cluster using a specific VR address. The VR ID must also be different from VR IDs used for other VR addresses on any other VRRP router on the network segment connected to this interface.

Note: The VR ID must only be unique on a single layer 2 network segment. The same VR ID may be used on different segments.

Note: In situations where it is unclear whether or not there is other VRRP traffic on a segment, run packet captures looking for VRRP to see if any turns up. There would typically be at least one VRRP advertisement per second from other nodes on the network. A packet capture would also show which VR IDs are active on the segment and thus should be avoided.

Tip: Though it is common to use the last octet of the VR address as the VR ID, this is not required.

Example which creates a new virtual router address:

```
tnsr(config)# int TenGigabitEthernet6/0/0
tnsr(config-interface)# ip vrrp-virtual-router 220
tnsr(config-vrrp4)#
```

This command enters `config-vrrp4` (IPv4) or `config-vrrp6` (IPv6) mode to configure the properties of the VR address. This mode includes the following commands:

virtual-address <ip-address> The IPv4 or IPv6 address which will be shared by the virtual router. Also referred to as the “Virtual Router Address” or “VR Address”.

For the primary node, or owner, for this address (priority 255), the same IP address must be configured on an interface.

accept-mode (true|false) Controls whether TNSR will accept packets delivered to this virtual address while in master state if it is not the IP address owner (Priority of 255). The default is `false`.

Deployments that rely on pinging the virtual address or using it for services such as DNS or IPsec should enable this feature.

Note: Accept mode has no effect when the VR address priority is set to 255. In that case, the router with priority 255 is considered the owner of the address and will already receive traffic without accept mode.

Note: IPv6 Neighbor Solicitations and Neighbor Advertisements MUST NOT be dropped when accept-mode is 'false'.

preempt (true|false) Instructs TNSR whether or not to preempt a lower priority peer to become master. The default value is `true`, and the owner of a VR address will always preempt other nodes, no matter how this value is set. When set to `false`, a failed node will not take back over from the current master when it recovers, but would wait until a new election occurs.

priority <priority> The priority for the VR address on this host. Higher values are preferred during the master election process, with the highest priority router currently operating winning the election.

The primary node, which is the owner of the VR address, must use a priority of 255 and no other node should have that priority. Lower priority nodes should use unique priority values, evenly distributed throughout the 1–254 range, depending on the number of nodes. The default value is 100.

Warning: VRRP priority 255 is not compatible with NAT. To use VRRP on an interface configured for outbound NAT, use a lower priority (1–254) instead. Since no router will be the owner of the VR address in this case, enable `accept-mode` to receive traffic for the VR address if communication with the TNSR host is necessary.

v3-advertisement-interval <interval> The interval, specified in centiseconds (hundredths of a second), at which VRRP advertisements will be sent by this node. The default value is 100, or one second. The value may be in the range of 1–4095.

track-interface <interface> priority-decrement <value> This command configures interface tracking, which allows the status of a *different* interface to affect the priority value advertised for this VR address. This allows TNSR to demote itself when other interfaces fail in some way, otherwise known as “preemption”.

The following conditions constitute a failure which results in a priority decrease when tracking an interface:

- The tracked interface is administratively disabled
- The tracked interface suffers a link loss
- The tracked interface no longer has an IP address matching the address family of this VR address

When the priority is decreased by the configured amount, other routers with the same VR address may preempt this router and assume a master role if they now have a higher priority for the VR address. In other words, this allows a VR address to demote itself upon detecting a problem without harshly changing the VR address status directly.

Note: Consider a scenario with all routers sharing the same VR address configured with interface tracking, and all suffer the same failure. For example, a dead switch. In that case, the current master would still be master even with adjusted priority values since all affected routers would have adjusted their priorities by the same amount. If instead the VR address state were adjusted directly to assume

a backup role, then there would be no active master remaining, and connectivity would be lost to this VR address.

interface The interface monitored by TNSR for making VR address priority adjustments. This must be a *different* interface, not the interface holding the VR address.

value The amount by which the priority value will be decreased when the status of <interface> changes to a failed state. May be from 0–255. The chosen value is up to the administrator and varies depending on the importance of the interface being tracked and the desirability of triggering a preemption.

For example, if this router is typically the owner of the VR address with a priority of 255, and the next highest router has a priority of 200, then a value of 60 would ensure that a failure will decrease the priority sufficiently to allow the other router to assume the master role.

Note: The advertised priority value can only decrease to a minimum of 1.

12.4 VRRP Status

The status of VRRP VR addresses is included in the output of `show interface [<if-name>]`. To view only the VRRP status and no other information, use `show interface [<if-name>] ip vrrp-virtual-router` for IPv4 or `show interface [<if-name>] ipv6 vrrp-virtual-router` for IPv6.

```
r1 tnsr# show interface ip vrrp-virtual-router
Interface: TenGigabitEthernet6/0/0
  IPv4 VRRP:
    VR: 220
      State: master, Priority: 255, Flags:          Addresses: 203.0.113.2
      Timers: Adv 100cs, Master down 300cs, Skew 0cs

Interface: TenGigabitEthernet6/0/1
  IPv4 VRRP:
    VR: 210
      State: master, Priority: 255, Flags:          Addresses: 198.51.100.1
      Timers: Adv 100cs, Master down 300cs, Skew 0cs
```

```
r2 tnsr# show interface ip vrrp-virtual-router
Interface: TenGigabitEthernet6/0/0
  IPv4 VRRP:
    VR: 220
      State: backup, Priority: 100, Flags:          Addresses: 203.0.113.2
      Timers: Adv 100cs, Master down 341cs, Skew 41cs

Interface: TenGigabitEthernet6/0/1
  IPv4 VRRP:
    VR: 210
      State: backup, Priority: 100, Flags:          Addresses: 198.51.100.1
      Timers: Adv 100cs, Master down 341cs, Skew 41cs
```

In a properly configured cluster in a normal state, the output should be similar to the above sample. The primary node will show *master* for the state of all configured VR addresses, and the secondary node will show *backup*.

ACCESS LISTS

Access Lists can be used to control ingress or egress traffic or to match hosts, networks and other contexts. An ACL contains a set of rules that defines source and destination hosts or networks to match, along with other aspects of traffic such as protocol and port number. Access Lists have an implicit final deny action. Any traffic not matched with an explicit permit rule will be dropped. Access Lists assume “any” for a value unless otherwise specified.

Access Lists can be stateful (`reflect`), or work without state tracking (`permit`).

Access Lists must be defined first and then applied to an interface along with a specific direction.

Host ACLs operate differently, as they govern traffic for interfaces in the host operating system rather than inside TNSR.

13.1 Standard ACLs

A standard ACL works with IPv4 or IPv6 traffic at layer 3. The name of an ACL is arbitrary so it may be named in a way that makes its purpose obvious.

ACLs consist of one or more rules, defined by a sequence number that determines the order in which the rules are applied. A common practice is to start numbering at a value higher than 0 or 1, and to leave gaps in the sequence so that rules may be added later. For example, the first rule could be 10, followed by 20.

Each rule must have an **action** and a defined `ip-version`. Rules can also define a `source`, `destination`, `protocol`, and other attributes for matching packets.

description `<text>` Text describing the purpose of this ACL.

action (`deny`|`permit`|`reflect`) Determines what happens to packets matched by the rule. This is required.

deny Drop a packet matching this rule.

permit Pass a single packet matching the rule. Since this action is per-packet and stateless, a separate ACL may also be required to pass traffic in the opposite direction.

reflect Permit a packet matching this rule and use a stateful packet processing path. Track the session and automatically permit return traffic in the opposite direction.

ip-version (`ipv4`|`ipv6`) Controls whether IPv4 or IPv6 packets will be matched by the rule. This is required, and also governs validation of the source and destination when applicable.

(source|destination) Define matching criteria for a rule based on where a packet came from or where it is going.

source address `<ip-address>` Match the source address of a packet. The given address must match the type set for `ip-version`.

source port any Match any TCP or UDP source port number (0 through 65535). Only valid when `protocol` is set to TCP or UDP. This is the default behavior when the rule does not contain a source port value.

source port <port-first> [- <port-last>] Match the specified TCP or UDP source port or range of source ports. When supplying a range, the first port must be lower than the last port. Only valid when `protocol` is set to `tcp` or `udp`.

destination address <ip-address> Match the destination address of a packet. The given address must match the type set for `ip-version`.

destination port any Match any TCP or UDP destination port number (0 through 65535). Only valid when `protocol` is set to TCP or UDP. This is the default behavior when the rule does not contain a destination port value.

destination port <port-first> [- <port-last>] Match the specified TCP or UDP destination port or range of destination ports. When supplying a range, the first port must be lower than the last port. Only valid when `protocol` is set to `tcp` or `udp`.

Note: Matching a source or destination port is only possible when the protocol is explicitly set to `tcp` or `udp`.

protocol (any|icmp|icmpv6|tcp|udp|<proto-number>) Sets the protocols which will be matched by this rule. This may be one of: `any`, `icmp`, `icmpv6`, `tcp`, `udp`, or a numeric [protocol number](#) from 0–255. If no protocol is specified, then the rule will match any protocol.

tcp flags value <v> mask <m> For rules matching TCP packets, `tcp flags` further restrict the match. This statement requires both a `value` and `mask`, which may be given in either order. The `value` and `mask` together define the flags matched out of a possible set of flags. These flags are specified numerically using the standard values for the flags: URG=32, ACK=16, PSH=8, RST=4, SYN=2, FIN=1. Add the values together to reach the desired value.

For example, with stateful filtering a common way to detect the start of a TCP session is to look for the TCP SYN flag with a mask of SYN+ACK. That way it will match only when SYN is set and ACK is not set. Using the values from the previous paragraph yields: `tcp flags value 2 mask 18`

icmp (codetype) <first> [- <last>] For rules matching ICMP protocol packets, `icmp type` and `icmp code` restrict matches to a specific value or range. The type and code are entered numerically in the range of 0–255. For a list of possible type and code combinations, see the [IANA ICMP Parameters list](#).

icmp (codetype) any Match any ICMP code or type. This is the default behavior.

13.1.1 Standard ACL Example

The following example ACL will block only SSH (tcp port 22) to 203.0.113.2 and permit all IPv4 other traffic:

```
tnsr(config)# acl blockssh
tnsr(config-acl)# rule 10
tnsr(config-acl-rule)# action deny
tnsr(config-acl-rule)# ip-version ipv4
tnsr(config-acl-rule)# destination address 203.0.113.2/32
tnsr(config-acl-rule)# destination port 22
tnsr(config-acl-rule)# protocol tcp
tnsr(config-acl-rule)# exit
```

(continues on next page)

(continued from previous page)

```
tnsr(config-acl)# rule 20
tnsr(config-acl-rule)# action permit
tnsr(config-acl-rule)# ip-version ipv4
tnsr(config-acl-rule)# exit
tnsr(config-acl)# exit
tnsr(config)# int GigabitEthernet0/14/1
tnsr(config-interface)# access-list input acl blockssh sequence 10
tnsr(config-interface)# exit
tnsr(config)#
```

Deconstructing the above example, the ACL behaves as follows:

- The name of the ACL is `blockssh`
- The first rule is 10. This leaves some room before it in case other rules should be matched before this rule in the future.
- Rule 10 will **deny** traffic matching:
 - A destination of a single IPv4 address, `203.0.113.2`
 - A destination of a single TCP port, 22 (ssh)
 - A source of **any** is implied since it is not specified
- The second rule is 20. The gap between 10 and 20 leaves room for future expansion of rules between the two existing rules.
- Rule 20 will **permit** all other IPv4 traffic, since there is no source or destination given.

The ACL is then applied to `GigabitEthernet0/14/1` in the inbound direction.

13.2 MACIP ACLs

MACIP ACLs and layer 3 ACLs (*Standard ACLs*) work similarly, but MACIP ACLs match traffic at layer 2 using MAC addresses.

Since MACIP ACLs work with layer 2 information, they can only effectively function on interfaces which support operating at layer 2, such as Ethernet. Additionally, MACIP ACLs can only match layer 2 interface packets from neighboring hosts on directly connected networks.

Warning: The MAC address of a remote host that reaches TNSR via routing through another gateway cannot be determined, thus cannot be matched by a MACIP ACL.

For example, traffic arriving at TNSR from the Internet via Ethernet will typically have a source MAC address of the default gateway or routing peer, and *not* the actual source of the traffic.

MACIP ACLs may only be applied in the input direction, and only match source addresses.

description <text> Text describing the purpose of this ACL.

action <name> Determines how the rule governs packets that match.

deny Drops a packet which matches this rule.

permit Passes a single packet matching the rule.

ip-version (ipv4|ipv6) Controls whether IPv4 or IPv6 packets will be matched by the rule. This is required when an address is present for the rule, and governs validation of the address value when applicable.

address <ip-prefix> Match the source IPv4 or IPv6 address of a packet.

mac address <mac-address> Optionally specifies a MAC address to block, in six groups of two colon-separated hexadecimal values, such as 00:11:22:33:44:55. When unset, the default value is 00:00:00:00:00:00 and uses the same value for a mask, which will match any MAC address.

mac mask <mac-mask> Optionally specifies a mask which defines portions of a MAC address to match, similar to an IP Prefix value. Given in six groups of two colon-separated hexadecimal values, such as ff:ff:ff:00:00:00, which matches the first half of a given MAC address. A mask of ff:ff:ff:ff:ff:ff matches an entire MAC address exactly. A mask of 00:00:00:00:00:00 matches any MAC address, and is the default behavior when no mask is set.

13.2.1 MACIP ACL Example

```
tnsr(config)# macip blockamac
tnsr(config-macip)# rule 10
tnsr(config-macip-rule)# action deny
tnsr(config-macip-rule)# mac address 00:11:22:33:44:55
tnsr(config-macip-rule)# mac mask ff:ff:ff:ff:ff:ff
tnsr(config-macip-rule)# exit
tnsr(config-macip)# exit
tnsr(config)# int GigabitEthernet0/14/2
tnsr(config-interface)# access-list macip blockamac
tnsr(config-interface)# exit
tnsr(config)#
```

13.3 Viewing ACL and MACIP Information

The `show acl [<name>]` command prints a list of defined ACLs and their actions. If <name> is given, then output is limited to the specified ACL.

```
tnsr# show acl

Access Control List: blockssh
  IPv Seq Action      Source          Dest Proto      SP/T  DP/C Flag Mask
  --- -- -
ipv4  10  deny  0.0.0.0/0 203.0.113.2/32 tcp   0-65535 22-22 0x00 0x00
ipv4  20  permit 0.0.0.0/0 0.0.0.0/0    0
```

The `show macip [<name>]` command works the same way for MACIP entries:

```
tnsr(config)# show macip

MACIP ACL: blockamac
  AF Seq Action  IP Prefix      MAC Address
  --- -- -
ipv4  10  deny 0.0.0.0/0 00:11:22:33:44:55 ff:ff:ff:ff:ff:ff
```

13.4 ACL and NAT Interaction

When NAT is active, ACL rules are always processed before NAT on interfaces where NAT is applied, in any direction.

The remainder of the section refers to the following example static NAT rule:

```
nat static mapping tcp local 10.2.0.129 22 external 203.0.113.2 222
```

In this example, that rule is applied on the external-facing interface containing 203.0.113.2.

13.4.1 Inbound ACL Rules

ACL Rules set to be processed in the **inbound** direction on an interface (`access-list input acl <name> sequence <seq>`) will match on the `external` address and/or port in a static NAT rule. In the above example, this means an inbound ACL would match on a destination IP address of 203.0.113.2 and/or a destination port of 222.

13.4.2 Outbound ACL Rules

ACL Rules set to be processed in the **outbound** direction on an interface (`access-list output acl <name> sequence <seq>`) will match on the `local` address and/or port in a static NAT rule. In the above example, this means an outbound ACL would match on a source IP address of 10.2.0.129 and/or a source port of 22.

13.5 Host ACLs

TNSR can also create host ACLs to control traffic on host interfaces, such as the management interface. These ACLs are implemented using Netfilter.

As mentioned in *Default Allowed Traffic*, TNSR includes a default set of host ACLs which protect host OS interfaces. Host ACLs created by administrators can override or augment the default blocking behavior.

ACLs are ordered by sequence number, and evaluated from the start to the end, stopping when a match is found. Each ACL contains one or more rules which define matching criteria and actions taken.

To create a new ACL, from `config` mode, use the command `host acl <acl-name>`, with the name to use for the new ACL. This command enters `config-host-acl` mode, where the following commands are available:

description <text> A text description of the host ACL.

sequence <acl-seq> The sequence number of this ACL. This sequence number controls the order of the ACLs when TNSR generates the host OS ruleset.

rule <rule-seq> Creates a new rule in this ACL with the given sequence number and enters `config-host-acl-rule` mode. The sequence number of the rule controls the order of the individual rules inside this ACL.

Inside `config-host-acl-rule` mode, the following commands are available:

action (deny|permit) Controls whether packets matching this rule will be passed (`permit`) or dropped (`deny`).

description <text> A text description of this rule.

match input-interface <host-interface> When set, this rule will only match traffic on the given host interface name. This is an interface name as seen by the host operating system, and not a TNSR interface.

match ip address (source|destination) <ip-prefix> Matches based on a given source or destination network.

match ip icmp type <type> [code <code>] Matches a specific IPv4 ICMP type and optionally matches the ICMP code as well. To match ICMP, the IP protocol must be set to `icmp`. Allowed types include: `address-mask-reply`, `address-mask-request`, `destination-unreachable`, `echo-reply`, `echo-request`, `info-reply`, `info-request`, `parameter-problem`, `redirect`, `router-advertisement`, `router-solicitation`, `source-quench`, `time-exceeded`, `timestamp-reply`, and `timestamp-request`.

match ip icmpv6 type <type> [code <code>] Matches a specific IPv6 ICMP type and optionally matches the ICMP code as well. To match ICMP, the IP protocol must be set to `icmp`. Allowed types include: `destination-unreachable`, `echo-reply`, `echo-request`, `mld-listener-query`, `mld-listener-reduction`, `mld-listener-report`, `nd-neighbor-advert`, `nd-neighbor-solicit`, `nd-redirect`, `nd-router-advert`, `nd-router-solicit`, `packet-too-big`, `parameter-problem`, `router-renumbering`, and `time-exceeded`.

match ip port (source|destination) <port-num> Matches the given source or destination port number. To match a port, the protocol must be `tcp` or `udp`.

match ip port (source|destination) range start <low-port-num> [end <high-port-num>] Matches the given source or destination port range, given as a lower start port number and a higher ending port number. To match a port, the protocol must be `tcp` or `udp`.

match ip protocol (icmp|tcp|udp|<proto-number>) Matches the specified IP protocol. When unset, any protocol will match the rule. However, this option must be set to enable protocol-specific matching such as ports (TCP or UDP) or ICMP types/codes. To match protocols other than TCP, UDP, and ICMP, specify the [protocol number](#) from 0-255.

match ip tcp flag (ack|cwr|ecf|fin|psh|rst|syn|urg) Matches a specific TCP flag. May only be used when protocol is set to `tcp`.

match ip version (4|6) Matches based on whether a packet is IPv4 (4), or IPv6 (6). This is required when matching by source or destination address.

match mac address (source|destination) <mac> Matches based on the source or destination MAC address. This is only valid for neighboring hosts on interfaces which provide layer 2 information, such as Ethernet.

13.5.1 Host ACL Example

This example configures a rule to allow traffic from the remote system 203.0.113.54 to reach a local host OS daemon on port 12345, used by the [TNSR IDS](#) daemon:

```
tnsr(config)# host acl tnsrids
tnsr(config-host-acl)# sequence 10
tnsr(config-host-acl)# description TNSR IDS
tnsr(config-host-acl)# rule 100
tnsr(config-host-acl-rule)# description Pass to tnsrids
tnsr(config-host-acl-rule)# action permit
tnsr(config-host-acl-rule)# match ip address source 203.0.113.54/32
tnsr(config-host-acl-rule)# match ip version 4
tnsr(config-host-acl-rule)# match ip protocol tcp
tnsr(config-host-acl-rule)# match ip port destination 12345
```


13.5.2 Host ACL Status

To see the list of current host ACLs, use the following command:

```
tnsr# show host ruleset
table inet tnsr_filter {
    chain tnsr_input_mgmt_local {
        jump tnsrids
    }

    chain tnsr_input_mgmt_default {
        tcp dport ssh accept
        tcp dport http accept
        tcp dport https accept
        ip protocol icmp accept
        ip6 nexthdr ipv6-icmp accept
        tcp dport bgp accept
        ip protocol ospf accept
        udp dport isakmp accept
        tcp dport ntp accept
        udp dport ntp accept
        tcp dport domain accept
        udp dport domain accept
        udp dport snmp accept
        udp dport bootps accept
        ip ttl 1 udp dport 33434-33524 counter packets 0 bytes 0 accept
    }

    chain tnsr_input {
        type filter hook input priority 0; policy accept;
        iifname "lo" accept
        ct state established,related accept
        jump tnsr_input_mgmt_local
        jump tnsr_input_mgmt_default
        drop
    }

    chain tnsr_forward {
        type filter hook forward priority 0; policy drop;
    }

    chain tnsrids {
        tcp dport 12345 counter packets 0 bytes 0 accept
    }
}
```

IPSEC

IPsec provides a standards-based VPN implementation compatible with other IPsec implementations. The IPsec sub-system in TNSR is handled by [strongSwan](#).

Currently, TNSR supports routed IPsec, allowing BGP or static routes to send traffic through IPsec.

14.1 Required Information

Before attempting to configure an IPsec tunnel, several pieces of information are required in order for both sides to build a tunnel. Typically the administrators of both tunnel endpoints will negotiate and agree upon the values to use for an IPsec tunnel.

At a minimum, these pieces of information should be known to both endpoints before attempting to configure a tunnel:

Local Address The IP address on TNSR which will be used to send and accept IPsec traffic from the peer.

Local IKE Identity The IKE identifier for TNSR, typically an IP address and the same as **Local Address**.

Local Network(s) A list of local networks which will communicate through the IPsec tunnel to hosts on **Remote Network(s)**. This is not entered into the configuration on TNSR for routed IPsec, but will be needed by the peer.

Remote Address The IP address of the IPsec peer.

Remote IKE Identity The identifier for the IPsec peer, typically the same as **Remote Address**.

Remote Network(s) A list of networks at the peer location with which hosts in the **Local Network(s)** will communicate. If using static routing, routes must be manually added for these networks using the **Remote IPsec Address** and `ipipX` interface. If BGP is used with IPsec, this will be handled automatically.

IKE Version Either 1 for IKEv1 or 2 for IKEv2. IKEv2 is stronger and more capable, but not all IPsec equipment can properly handle IKEv2.

IKE Lifetime The maximum amount of time that an IKE session can stay alive until it is renegotiated.

IKE Encryption The encryption algorithm used to encrypt IKE messages.

IKE Integrity The integrity algorithm used to authenticate IKE messages

IKE DH/MODP Group Diffie-Hellman group for key establishment, given in bits.

IKE Authentication The type of authentication used to verify the identity of the peer.

Pre-Shared Key When using Pre-Shared Key for IKE Authentication, this key is used on both sides to authenticate the peer.

SA Lifetime The amount of time that a child security association can be active before it is rekeyed.

SA Encryption The encryption algorithm used to encrypt tunneled traffic.

SA Integrity The integrity algorithm used to authenticate tunneled traffic.

SA DH/MODP Group Diffie-Hellman group for security associations, in bits.

Local IPsec Address The local IP address for the `ipipX` interface, used for routing traffic to/from IPsec peers.

Remote IPsec Address The remote IP address for the peer on `ipipX`, used as a gateway for routing, or a BGP neighbor.

Warning: If NAT is active on the same interface acting as an IPsec endpoint, then NAT forwarding must also be enabled. See [NAT Forwarding](#).

14.2 IPsec Example

14.2.1 Required Information

This table contains the *Required Information* used to form the IPsec tunnel used in this example.

Table 1: Example IPsec Configuration

Item	Value
Local Address	203.0.113.2
Local IKE Identity	203.0.113.2
Local Network(s)	10.2.0.0/16
Remote Address	203.0.113.25
Remote IKE Identity	203.0.113.25
Remote Network(s)	10.25.0.0/16
IKE Version	2
IKE Lifetime	28800
IKE Encryption	AES-128
IKE Integrity	SHA1
IKE DH/MODP Group	2048 (14)
IKE Authentication	Pre-Shared Key
Pre-Shared Key	mysupersecretkey
SA Lifetime	3600
SA Encryption	AES-128
SA Integrity	SHA1
SA DH/MODP Group	2048 (14)
Local IPsec Address	172.32.0.1/30
Remote IPsec Address	172.32.0.2

14.2.2 Example Configuration

This configuration session implements the tunnel described by the settings in *Example IPsec Configuration*:

```
tnsr(config)# ipsec tunnel 0
tnsr(config-ipsec-tun)# local-address 203.0.113.2
tnsr(config-ipsec-tun)# remote-address 203.0.113.25
tnsr(config-ipsec-tun)# crypto config-type ike
tnsr(config-ipsec-tun)# crypto ike
tnsr(config-ipsec-crypto-ike)# version 2
tnsr(config-ipsec-crypto-ike)# lifetime 28800
tnsr(config-ipsec-crypto-ike)# proposal 1
tnsr(config-ike-proposal)# encryption aes128
tnsr(config-ike-proposal)# integrity sha1
tnsr(config-ike-proposal)# group modp2048
tnsr(config-ike-proposal)# exit
tnsr(config-ipsec-crypto-ike)# identity local
tnsr(config-ike-identity)# type address
tnsr(config-ike-identity)# value 203.0.113.2
tnsr(config-ike-identity)# exit
tnsr(config-ipsec-crypto-ike)# identity remote
tnsr(config-ike-identity)# type address
tnsr(config-ike-identity)# value 203.0.113.25
tnsr(config-ike-identity)# exit
tnsr(config-ipsec-crypto-ike)# authentication local
tnsr(config-ike-auth)# round 1
tnsr(config-ike-auth-round)# type psk
tnsr(config-ike-auth-round)# psk mysupersecretkey
tnsr(config-ike-auth-round)# exit
tnsr(config-ike-auth)# exit
tnsr(config-ipsec-crypto-ike)# authentication remote
tnsr(config-ike-auth)# round 1
tnsr(config-ike-auth-round)# type psk
tnsr(config-ike-auth-round)# psk mysupersecretkey
tnsr(config-ike-auth-round)# exit
tnsr(config-ike-auth)# exit
tnsr(config-ipsec-crypto-ike)# child 1
tnsr(config-ike-child)# lifetime 3600
tnsr(config-ike-child)# proposal 1
tnsr(config-ike-child-proposal)# encryption aes128
tnsr(config-ike-child-proposal)# integrity sha1
tnsr(config-ike-child-proposal)# group modp2048
tnsr(config-ike-child-proposal)# exit
tnsr(config-ike-child)# exit
tnsr(config-ipsec-crypto-ike)# exit
tnsr(config-ipsec-tun)# exit
tnsr(config)# interface ipip0
tnsr(config-interface)# ip address 172.32.0.1/30
tnsr(config-interface)# exit
tnsr(config)# route ipv4 table ipv4-VRF:0
tnsr(config-route-table-v4)# route 10.25.0.0/16
tnsr(config-rttbl4-next-hop)# next-hop 0 via 172.32.0.2 ipip0
tnsr(config-rttbl4-next-hop)# exit
tnsr(config-route-table-v4)# exit
tnsr(config)# exit
```

This example is used as a reference through the remainder of the chapter.

Tip: If the TNSR device hardware supports cryptographic acceleration, enable it for optimal performance. See [IPsec Cryptographic Acceleration](#) for details.

14.3 IPsec Configuration

The `ipsec tunnel <n>` command, issued from `config` mode, changes to IPsec tunnel mode. This is denoted by `config-ipsec-tun` in the prompt.

The identifier number for tunnel entries starts at 0 and increments by one. To determine the next tunnel number for a new entry, run `ipsec tunnel ?` and TNSR will print the existing tunnel ID numbers.

This command creates an IPsec tunnel with an identifier of 0:

```
tnsr(config)# ipsec tunnel 0
tnsr(config-ipsec-tun)#
```

The remainder of the configuration is covered in the following sections.

14.3.1 IPsec Endpoints

Next, the IPsec tunnel needs endpoints, defined using the following commands from within `config-ipsec-tun` mode:

local-address Defines the IP address used by TNSR for this IPsec tunnel. This address must exist on a TNSR interface.

remote-address Defines the IP address or hostname of the remote peer.

IPsec Endpoint Example

```
tnsr(config-ipsec-tun)# local-address 203.0.113.2
tnsr(config-ipsec-tun)# remote-address 203.0.113.25
```

14.3.2 IPsec Keys

Inside `config-ipsec-tun` mode, the following commands are available for IPsec key management.

crypto config-type {ike|manual} Configures the type of key management TNSR will use for this tunnel.

ike Internet Key Exchange (IKE). The most common method of key management. IPsec tunnels utilize IKE to dynamically handle key exchange when both parties are negotiating a security association.

manual Static key management.

crypto ike Enters IKE `config-ipsec-crypto-ike` mode to configure IPsec IKE behavior, which is the bulk of the remaining work for most IPsec tunnels.

IKE Configuration

Inside `config-ipsec-crypto-ike` mode, the following commands are available to configure basic IKE behavior:

version <x> Instructs TNSR to use either IKEv1 or IKEv2. Use 2 for IKEv2, which is more secure, or 1 for IKEv1 which is more common and more widely supported.

lifetime <x> Sets the maximum time for this IKE session to be valid, in seconds within the range 120 . .214783647. Default value is 14400 seconds (4 hours). Commonly set to 28800 seconds (8 hours). This value should be longer than the IKE child lifetime, discussed later.

dpd-interval <x> Optional time to wait between sending Dead Peer Detection (DPD) polls, given in seconds within the range 0–65535.

key-renewal (reauth|rekey) Controls the method used to update keys on an established IKE security association (SA) before the lifetime expires.

reauth TNSR performs a full teardown and re-establishment of IKE and child SAs.

rekey Inline rekeying while SAs stay active. Only available in IKEv2.

proposal <name> Configures a new *IKE proposal* and enters `config-ike-proposal` mode.

identity (local|remote) Configures *IKE identity* validation and enters `config-ike-identity` mode.

authentication (local|remote) Configures *IKE authentication* and enters `config-ike-auth` mode.

Additional `config-ipsec-crypto-ike` mode commands are available to configure other aspects of the IPsec tunnel, such as proposals, identity, and authentication. These are covered next.

IKE Example

This example tells TNSR to use IKE for key management, and then sets the tunnel to IKEv2 and a lifetime of 8 hours.

```
tnsr(config-ipsec-tun)# crypto config-type ike
tnsr(config-ipsec-tun)# crypto ike
tnsr(config-ipsec-crypto-ike)# version 2
tnsr(config-ipsec-crypto-ike)# lifetime 28800
```

Additional IKE Configuration

The remainder of the IKE configuration is covered in the following sections.

IKE Proposal

IKE Proposals instruct TNSR how the key exchange will be encrypted and authenticated. TNSR supports a variety of encryption algorithms, integrity/authentication hash algorithms, pseudo-random functions (PRF), and Diffie-Hellman (DH) group specifications. These choices must be coordinated between both endpoints.

Tip: Some vendor IPsec implementations refer to IKE/ISAKMP as “Phase 1”, which may help when attempting to map values supplied by a peer to their corresponding values in TNSR.

From within `config-ipsec-crypto-ike` mode, use the `proposal <name>` command to start a new proposal and enter `config-ike-proposal` mode. In `config-ike-proposal` mode, the following commands are available:

encryption <ea-name> Configures the *encryption algorithm* to use for the proposal.

integrity <ia-name> Configures the *integrity algorithm* to use for the proposal.

prf <prf-name> Configures the *pseudo-random function* (PRF) to use for the proposal.

group <group-name> Configures the *Diffie-Hellman group* (DH Group) to use for the proposal.

Tip: To see a list of supported choices for each option, follow the initial command with a `?`, such as `encryption ?`.

Each of these is described in more detail in the following sections.

Encryption Algorithms

TNSR supports many common, secure encryption algorithms. Some older and insecure algorithms are not supported. Algorithms based on AES are common and secure, and are widely supported by other VPN implementations.

AES-GCM, or [AES Galois/Counter Mode](#) is an efficient and fast authenticated encryption algorithm, which means it provides data privacy as well as integrity validation, without the need for a separate integrity algorithm.

Additionally, AES-based algorithms can often be accelerated by AES-NI.

Warning: TNSR includes the Triple-DES (3DES) algorithm for compatibility with legacy systems, but it is not considered secure. Specifically, 3DES is considered broken by attacks such as [Sweet32](#). Use stronger encryption algorithms where possible.

A full list of encryption algorithms supported by TNSR:

```
tnsr(config-ike-proposal)# encryption ?
<cr>
3des                Triple-DES
aes128              128 bit AES-CBC
aes128ccm12         128 bit AES-CCM with 12 byte ICV
aes128ccm16         128 bit AES-CCM with 16 byte ICV
aes128ccm8          128 bit AES-CCM with 8 byte ICV
aes128ctr           128 bit AES-Counter
aes128gcm12         128 bit AES-GCM with 12 byte ICV
aes128gcm16         128 bit AES-GCM with 16 byte ICV
aes128gcm8          128 bit AES-GCM with 8 byte ICV
aes192              192 bit AES-CBC
aes192ccm12         192 bit AES-CCM with 12 byte ICV
aes192ccm16         192 bit AES-CCM with 16 byte ICV
aes192ccm8          192 bit AES-CCM with 8 byte ICV
aes192ctr           192 bit AES-Counter
aes192gcm12         192 bit AES-GCM with 12 byte ICV
aes192gcm16         192 bit AES-GCM with 16 byte ICV
aes192gcm8          192 bit AES-GCM with 8 byte ICV
aes256              256 bit AES-CBC
aes256ccm12         256 bit AES-CCM with 12 byte ICV
aes256ccm16         256 bit AES-CCM with 16 byte ICV
aes256ccm8          256 bit AES-CCM with 8 byte ICV
aes256ctr           256 bit AES-Counter
aes256gcm12         256 bit AES-GCM with 12 byte ICV
aes256gcm16         256 bit AES-GCM with 16 byte ICV
aes256gcm8          256 bit AES-GCM with 8 byte ICV
camellia128         128 bit Camellia
camellia128ccm12    128 bit Camellia-CCM with 12 byte ICV
camellia128ccm16    128 bit Camellia-CCM with 16 byte ICV
camellia128ccm8     128 bit Camellia-CCM with 8 byte ICV
camellia128ctr      128 bit Camellia-Counter
```

(continues on next page)

(continued from previous page)

camellia192	192 bit Camellia
camellia192ccm12	192 bit Camellia-CCM with 12 byte ICV
camellia192ccm16	192 bit Camellia-CCM with 16 byte ICV
camellia192ccm8	192 bit Camellia-CCM with 8 byte ICV
camellia192ctr	192 bit Camellia-Counter
camellia256	256 bit Camellia
camellia256ccm12	256 bit Camellia-CCM with 12 byte ICV
camellia256ccm16	256 bit Camellia-CCM with 16 byte ICV
camellia256ccm8	256 bit Camellia-CCM with 8 byte ICV
camellia256ctr	256 bit Camellia-Counter
chacha20poly1305	256 bit ChaCha20/Poly1305 with 16 byte ICV

Integrity Algorithms

Integrity algorithms provide authentication of messages and randomness, ensuring that packets are authentic and were not altered by a third party before arriving, and also for constructing keying material for encryption.

Note: When using an authenticated encryption algorithm like AES-GCM with a child Security Association (SA) as opposed to IKE/ISAKMP, an integrity option **should not** be configured, as it is redundant and reduces performance.

When an authenticated encryption algorithm is used with IKE, configure a Pseudo-Random Function (PRF) instead of an Integrity Algorithm. If an integrity algorithm is defined in this case, TNSR will attempt to map the chosen algorithm to an equivalent PRF.

A full list of integrity algorithms supported by TNSR:

```
tnsr(config-ike-proposal)# integrity ?
<cr>
aescmac          AES-CMAC 96
aesxcbc          AES-XCBC 96
md5              MD5 96
sha1             SHA1 96
sha256           SHA2 256 bit blocks, 128 bits output
sha384           SHA2 384 bit blocks, 192 bits output
sha512           SHA2 512 bit blocks, 256 bits output
```

Pseudo-Random Functions

A Pseudo-Random Function (PRF) is similar to an integrity algorithm, but instead of being used to authenticate messages, it is only used to provide randomness for purposes such as keying material. PRFs are primarily used with an authenticated encryption algorithm type such as AES-GCM, but they can be explicitly defined for use with other integrity algorithms.

If a PRF is not explicitly defined, TNSR will attempt to derive the PRF to use based on the integrity algorithm for a given proposal.

Note: In the case of AES-NI, `prfaesxcbc` is likely the most appropriate choice as it can be accelerated by AES-NI, and it is more widely supported than its improved successor `prfaescmac`.

A full list of pseudo-random functions supported by TNSR:


```
tnsr(config-ike-proposal)# prf ?
<cr>
prfaescmac          AES128-CMAC PRF
prfaesxcbc          AES128-XCBC PRF
prfmd5              MD5 PRF
prfsha1             SHA1 PRF
prfsha256           SHA2-256 PRF
prfsha384           SHA2-384 PRF
prfsha512           SHA2-512 PRF
```

Diffie-Hellman Groups

Diffie-Hellman (DH) exchanges allow two parties to establish a shared secret across an untrusted connection. DH choices can be referenced in several different ways depending on vendor implementations. Some reference a DH group by number, others by size. When referencing by group number, generally speaking higher group numbers are more secure.

Tip: In most cases, modp2048 (Group 14) is the lowest choice considered to provide sufficient security in a modern computing environment.

A full list of DH Groups supported by TNSR:

```
tnsr(config-ike-proposal)# group ?
<cr>
curve25519          Group 31 (Elliptic Curve 25519, 256 bit)
ecp256              Group 19 (256 bit ECP)
ecp384              Group 20 (384 bit ECP)
ecp521              Group 21 (521 bit ECP)
modp1024            Group 2 (1024 bit modulus)
modp1024s160        Group 22 (1024 bit modulus, 160 bit POS)
modp1536            Group 5 (1536 bit modulus)
modp2048            Group 14 (2048 bit modulus)
modp2048s224        Group 23 (2048 bit modulus, 224 bit POS)
modp2048s256        Group 24 (2048 bit modulus, 256 bit POS)
modp3072            Group 15 (3072 bit modulus)
modp4096            Group 16 (4096 bit modulus)
modp6144            Group 17 (6144 bit modulus)
modp768             Group 1 (768 bit modulus)
modp8192            Group 18 (8192 bit modulus)
```

Warning: TNSR supports modp768 (Group 1) and modp1024 (Group 2) for compatibility purposes but they are considered broken by the [Logjam Attack](#) and should be avoided.

TNSR also supports modp1024s160 (Group 22), modp2048s224 (Group 23), and modp2048s256 (Group 24) for compatibility but they should also be avoided as they have a [questionable source of primes](#).

IKE Proposal Example

This example configures one proposal. This proposal uses AES-128 encryption, SHA-1 for integrity hashing, and DH group 14 (2048 bit modulus).

```
tnsr(config-ipsec-crypto-ike)# proposal 1
tnsr(config-ike-proposal)# encryption aes128
tnsr(config-ike-proposal)# integrity sha1
tnsr(config-ike-proposal)# group modp2048
tnsr(config-ike-proposal)# exit
```

IKE Identity

In IKE, each party must ensure it is communicating with the correct peer. One aspect of this validation is the identity information included in IKE. Each router tells the other its own local identity and they each validate it against the stored remote identity. If they do not match, the peer is rejected.

From within `config-ipsec-crypto-ike` mode, use the `identity local` and `identity remote` commands to configure local and remote identity information. In either case, the `identity` command enters `config-ike-identity` mode.

IKE requires both local and remote identities. The local identity is sent to the remote peer during the exchange. The remote identity is used to validate the identity received from the peer during the exchange.

In `config-ike-identity`, the following commands are available:

type <name> Sets the type of identity value. The following types are available:

address IPv4 or IPv6 address in the standard notation for either (e.g. `192.0.2.3` or `2001:db8:1:2::3`)

This is the most common type, with the value set to the address on TNSR used as the `local-address` for the IPsec tunnel.

dn An X.509 distinguished name (e.g. `certificate subject`)

email Email address (e.g. `user@example.com`).

fqdn A fully qualified domain name (e.g. `host.example.com`)

key-id An arbitrary string used as an identity

none Automatically interpret the type based on the value

value <text> The identity value, in a format corresponding to the chosen `type`.

Note: The local identity type and value must both be supplied to the administrator of the remote peer so that it can properly identify this endpoint.

Identity Example

First configure the local identity of this firewall. The identity is an IP address, using the same value as the local address of the IPsec tunnel.

```
tnsr(config-ipsec-crypto-ike)# identity local
tnsr(config-ike-identity)# type address
tnsr(config-ike-identity)# value 203.0.113.2
tnsr(config-ike-identity)# exit
```

Next, configure the remote identity. The remote peer has also chosen to use an IP address, the value of which is the remote address used for the IPsec tunnel.

```
tnsr(config-ipsec-crypto-ike)# identity remote
tnsr(config-ike-identity)# type address
tnsr(config-ike-identity)# value 203.0.113.25
tnsr(config-ike-identity)# exit
```

IKE Authentication

After verifying the identity, TNSR will attempt to authenticate the peer using the secret from its configuration in one or two round passes. In most common configurations there is only a single authentication round, however in IKEv2 a tunnel may have two rounds of unique authentication.

From within `config-ipsec-crypto-ike` mode, use the `authentication local` and `authentication remote` commands to configure local and remote authentication information. In either case, the authentication command enters `config-ike-auth` mode.

TNSR will use the parameters under `authentication local` to authenticate outbound traffic and the `authentication remote` parameters are used to authenticate inbound traffic.

Note: With pre-shared key mode, most real-world configurations use identical values for both local and remote authentication.

From `config-ike-auth` mode, the `round <n>` command configures parameters for round 1 or 2. As mentioned previously, most configurations will only use round 1. The `round` command then enters `config-ike-auth-round` mode.

In `config-ike-auth-round` mode, the following commands are available:

type <name> The type of authentication to perform.

Currently the only authentication type supported by TNSR is `psk` (pre-shared key).

psk <text> For `psk` type authentication, this command defines the pre-shared key value.

IKE Authentication Example

This example only has one single round of authentication, a pre-shared key of `mysupersecretkey`. Thus, the `type` is set to `psk` and then the `psk` is set to the secret value.

Warning: Do not transmit the pre-shared key over an insecure channel such as plain text e-mail!

First, add the local authentication parameters:

```
tnsr(config-ipsec-crypto-ike)# authentication local
tnsr(config-ike-auth)# round 1
tnsr(config-ike-auth-round)# type psk
tnsr(config-ike-auth-round)# psk mysupersecretkey
tnsr(config-ike-auth-round)# exit
tnsr(config-ike-auth)# exit
```

Next, configure the remote authentication parameters. As in most practical uses, this is set identically to the local authentication value.

```
tnsr(config-ipsec-crypto-ike)# authentication remote
tnsr(config-ike-auth)# round 1
tnsr(config-ike-auth-round)# type psk
tnsr(config-ike-auth-round)# psk mysupersecretkey
tnsr(config-ike-auth-round)# exit
tnsr(config-ike-auth)# exit
```

14.3.3 Security Associations

After establishing a secure channel, the two endpoints can negotiate an IPsec security association (IPsec SA) as a “child” entry. TNSR supports adding multiple children as needed, though with routed IPsec only one is necessary.

Tip: Some vendor IPsec implementations refer to IPsec security association child entries as “Phase 2”, which may help when attempting to map values supplied by a peer to their corresponding values in TNSR.

From within `config-ipsec-crypto-ike` mode, the `child <n>` command configures the child noted by the given number. The `child` command enters `ike-child` mode.

Within `ike-child` mode, the following commands are available:

lifetime <x> Sets the maximum time for this child IPsec SA to be valid before it must be rekeyed. The value is given in seconds within the range 120 . . 214783647. Default value is 3600 seconds (one hour). This value must be shorter than the IKE lifetime, discussed earlier.

replay-window (0|64) Number of packets in replay window. The replay window is used to protect the tunnel against attacks where the sequence number is re-used or has been processed recently. Some allowance is helpful in dealing with network link issues that cause packets to arrive late or out-of-order. A value of 0 disables the replay window. A value of 64 enables a 64 packet replay window.

proposal <name> Each child may have one or more `proposal` entries which define acceptable encryption, integrity, and DH Group (Perfect Forward Security, PFS) parameters to encrypt and validate the IPsec SA traffic.

Child SA proposals work similarly to IKE/ISAKMP proposals as described in *IKE Proposal*.

This command enters `config-ike-child-proposal` mode to configure these proposals. in `config-ike-child-proposal` mode, the following commands are available:

encryption <ea-name> Configures the *encryption algorithm* to use for the proposal.

integrity <ia-name> Configures the *integrity algorithm* to use for the proposal.

group <group-name> Configures the *Diffie-Hellman group* (DH Group) to use for the proposal.

sequence-number (esn|noesn) Controls whether or not TNSR will attempt to negotiate extended sequence number (ESN) support with the peer. ESN uses 64-bit sequence numbers instead of the 32-bit sequence numbers. The default is `noesn` which disables ESN negotiation.

Child SA Example

This example only has a single child, thus `child 1`. The child has a lifetime of 3600.

```
tnsr(config-ipsec-crypto-ike)# child 1
tnsr(config-ike-child)# lifetime 3600
```

Next, create a child SA proposal. This example uses AES-128 for encryption, SHA-1 for an authentication hash, and PFS group 14 (2048 bit modulus).

```
tnsr(config-ike-child)# proposal 1
tnsr(config-ike-child-proposal)# encryption aes128
tnsr(config-ike-child-proposal)# integrity sha1
tnsr(config-ike-child-proposal)# group modp2048
```

This completes the configuration for the IPsec tunnel, at this point after exiting back to basic mode the tunnel will attempt to establish a connection to the peer.

```
tnsr(config-ike-child-proposal)# exit
tnsr(config-ike-child)# exit
tnsr(config-ipsec-crypto-ike)# exit
tnsr(config-ipsec-tun)# exit
```

14.3.4 Configuring the IPsec Interface

TNSR supports routed IPsec via the `ipipX` interface. The number of the ipsec interface corresponds to the index number of the tunnel set previously. For example `ipsec tunnel 0` is `ipip0`, and `ipsec tunnel 2` is `ipip2`.

These IPsec interfaces are used to configure routed IPsec connectivity and they behave like most other interfaces. For example, they can have access lists defined to filter traffic.

The `ipipX` interface should be configured with an IP address and the peer will have its own IP address in the same subnet. This allows the two endpoints to communicate directly over the IPsec interface and also gives the peer an address through which traffic for other subnets may be routed. When configured in this way, it acts like a directly connected point-to-point link to the peer.

IPsec Interface Example

In this example, the `ipip0` interface is given an address of `172.32.0.1/30`. The remote peer will be `172.32.0.2/30`.

```
tnsr(config)# interface ipip0
tnsr(config-interface)# ip address 172.32.0.1/30
tnsr(config-interface)# exit
```

14.3.5 IPsec Routes

The IPsec interface allows the peers to talk directly, but in most cases with IPsec there is more interesting traffic to handle. For example, a larger subnet on the LAN side of each peer that must communicate securely.

To allow these networks to reach one another, routes are required. These may be managed manually using static routes, or a dynamic routing protocol such as BGP can manage the routes automatically.

IPsec Static Route Example

This example adds a static route to the main IPv4 routing table for a subnet located behind the peer. Any traffic trying to reach a host inside the `10.25.0.0/16` subnet will be routed through the `ipip0` interface using the peer address in that subnet (`172.32.0.2`) as the next hop.

```
tnsr(config)# route ipv4 table ipv4-VRF:0
tnsr(config-route-table-v4)# route 10.25.0.0/16
tnsr(config-rttbl4-next-hop)# next-hop 0 via 172.32.0.2 ipip0
tnsr(config-rttbl4-next-hop)# exit
tnsr(config-route-table-v4)# exit
tnsr(config)# exit
```

See also:

For a larger example involving BGP for dynamic route management, see *TNSR IPsec Hub for pfSense*.

14.4 IPsec Status Information

To view status information about active IPsec tunnels, use the `show ipsec tunnel` command. This command prints status output for all IPsec tunnels, and it also supports printing tunnel information individually by providing the tunnel ID. This command supports several additional parameters to increase or decrease the amount of information it displays.

The following forms of `show ipsec tunnel` are available:

show ipsec tunnel Display a short summary of all IPsec tunnels.

show ipsec tunnel n Display a short summary of a specific IPsec tunnel n.

show ipsec tunnel [n] verbose Display a verbose list of all IPsec tunnels, optionally limited to a single tunnel n. The output shows detailed information such as active encryption, hashing, DH groups, identifiers, and more.

show ipsec tunnel [n] ike [verbose] Display only IKE parameters of all tunnels. Optionally limited to a single tunnel n and/or expanded details with `verbose`.

show ipsec tunnel [n] child [verbose] Display only IPsec child Security Association parameters of all tunnels. Optionally limited to a single tunnel n and/or expanded details with `verbose`.

14.4.1 IPsec Status Examples

Show the status of tunnel 0:

```
tnsr# show ipsec tunnel 0
IPsec Tunnel: 0
  IKE SA: ipip0    ID: 13    Version: IKEv2
    Local: 203.0.113.2    Remote: 203.0.113.25
    Status: ESTABLISHED    Up: 372s    Reauth: 25275s
  Child SA: child0    ID: 7
    Status: INSTALLED    Up: 372s    Rekey: 2523s    Expire: 3228s
    Received: 0 bytes, 0 packets
    Transmitted: 0 bytes, 0 packets
  Child SA: child0    ID: 8
    Status: INSTALLED    Up: 372s    Rekey: 2813s    Expire: 3228s
    Received: 0 bytes, 0 packets
    Transmitted: 0 bytes, 0 packets
  Child SA: child0    ID: 9
    Status: INSTALLED    Up: 372s    Rekey: 2583s    Expire: 3228s
    Received: 0 bytes, 0 packets
    Transmitted: 0 bytes, 0 packets
```

Adding the verbose keyword also shows detailed information about the encryption parameters:

```
tnsr# show ipsec tunnel 0 verbose
IPsec Tunnel: 0
  IKE SA: ipip0      ID: 13      Version: IKEv2
    Local: 203.0.113.2      Remote: 203.0.113.25
    Status: ESTABLISHED    Up: 479s    Reauth: 25168s
    Local ID: 203.0.113.2    Remote ID: 203.0.113.25
    Cipher: AES_CBC 128    MAC: HMAC_SHA1_96
    PRF: PRF_HMAC_SHA1    DH: MODP_2048
    SPI Init: 1880997989256787091    Resp: 1437908875259838715
    Initiator: yes
  Child SA: child0    ID: 7
    Status: INSTALLED    Up: 479s    Rekey: 2416s    Expire: 3121s
    Received: 0 bytes, 0 packets
    Transmitted: 0 bytes, 0 packets
    Cipher: AES_CBC 128    MAC: HMAC_SHA1_96    PFS: MODP_2048
    SPI in: 3540263882    out: 974161796
  Child SA: child0    ID: 8
    Status: INSTALLED    Up: 479s    Rekey: 2706s    Expire: 3121s
    Received: 0 bytes, 0 packets
    Transmitted: 0 bytes, 0 packets
    Cipher: AES_CBC 128    MAC: HMAC_SHA1_96    PFS: MODP_2048
    SPI in: 2432966668    out: 1361993947
  Child SA: child0    ID: 9
    Status: INSTALLED    Up: 479s    Rekey: 2476s    Expire: 3121s
    Received: 0 bytes, 0 packets
    Transmitted: 0 bytes, 0 packets
    Cipher: AES_CBC 128    MAC: HMAC_SHA1_96    PFS: MODP_2048
    SPI in: 2318058408    out: 1979056986
```

Specifying the `ike` or `child` parameter filters the output, and these also support `verbose` output.

Note: The first Child SA entry uses DH information from the parent IKE SA, and not its own PFS setting. As such, Child SA entries in this situation will display %IKE at the end of their PFS value to indicate the source. The PFS value configured on the Child SA is used when a Child SA is rekeyed.

```
tnsr# show ipsec tunnel 0 ike
IPsec Tunnel: 0
  IKE SA: ipip0      ID: 13      Version: IKEv2
    Local: 203.0.113.2      Remote: 203.0.113.25
    Status: ESTABLISHED    Up: 372s    Reauth: 25275s
```

```
tnsr# show ipsec tunnel 0 ike verbose
IPsec Tunnel: 0
  IKE SA: ipip0      ID: 13      Version: IKEv2
    Local: 203.0.113.2      Remote: 203.0.113.25
    Status: ESTABLISHED    Up: 479s    Reauth: 25168s
    Local ID: 203.0.113.2    Remote ID: 203.0.113.25
    Cipher: AES_CBC 128    MAC: HMAC_SHA1_96
    PRF: PRF_HMAC_SHA1    DH: MODP_2048
    SPI Init: 1880997989256787091    Resp: 1437908875259838715
    Initiator: yes
```

14.5 IPsec Cryptographic Acceleration

There are three types of cryptographic acceleration available for use on TNSR:

- Software cryptographic acceleration
- CPU-based AES-NI cryptographic acceleration
- Hardware-based Intel® QuickAssist Technology (QAT) cryptographic acceleration

The list above is in order of likely performance boost, from least to most. In other words, software acceleration is slowest, QAT is fastest. The availability of AES-NI CPU instructions and QAT vary by platform and installed hardware.

14.5.1 Software Cryptographic Acceleration

TNSR will automatically configure software cryptographic acceleration for VPP if an IPsec tunnel is defined in the configuration. To enable this configuration, the VPP service must be restarted manually so it can enable the feature and allocate additional memory.

Note: The cryptographic accelerator setting applies to all tunnels, so the restart is only required after the first IPsec tunnel configured by TNSR. The restart is not required for additional tunnels or when changing IPsec settings.

Restart the VPP dataplane from the TNSR basic mode CLI using the following command:

```
tnsr# config
tnsr(config)# service dataplane restart
```

If the TNSR configuration contains no IPsec tunnels, TNSR will not require the memory resources associated with cryptographic acceleration and TNSR will not require a restart of the VPP dataplane service.

14.5.2 AES-NI cryptographic acceleration

AES-NI cryptographic acceleration takes advantage of AES acceleration instructions available in most modern CPUs. Since this feature relies on CPU support, it is not available on all hardware and, depending on the hypervisor and its configuration, may not be passed through from a host to a VM.

AES-NI offers a significant performance boost with AES-based ciphers, especially with AEAD ciphers such as AES-GCM.

AES-NI is configured in the dataplane using the `dataplane dpdk vdev` command, which is covered in [DPDK Configuration](#).

14.5.3 QAT cryptographic acceleration

TNSR Supports hardware compatible with Intel® QuickAssist Technology (QAT) for accelerating cryptographic operations, such as IPsec. This requires the presence of a compatible QAT device, which may be a component of the hardware platform or an add-in card such as the CPIC devices sold by Netgate.

Note: This hardware can be found in CPIC cards as well as many C3000 and Skylake Xeon systems. Netgate [XG-1541](#) and [XG-1537](#) hardware has an add-on option for a CPIC card.

To configure a QAT device, follow the procedures described in *Setup QAT Compatible Hardware* to enable the device in the dataplane configuration.

Warning: There is a known incompatibility between QAT and VT-d on some platforms which can prevent IPsec traffic from passing when QAT acceleration is enabled. See *Disable VT-d in the BIOS* for details.

NETWORK ADDRESS TRANSLATION

Network Address Translation, or NAT, involves changing properties of a packet as it passes through a router. Typically this is done to mask or alter the source or destination to manipulate how such packets are processed by other hosts.

The most common examples are:

- Source NAT, also known as Outbound NAT, which translates the source address and port of a packet to mask its origin.
- Destination NAT, commonly referred to as Static NAT or Port Forwards which translate the destination address and port of a packet to redirect the packet to a different target host behind the router.

TNSR applies NAT based on the configured mode and the presence of directives that set `inside` (internal/local) and `outside` (external/remote) interfaces.

An `inside` interface is a local interface where traffic enters and it will have its source hidden by NAT. An `outside` interface is an interface where that translation will occur as a packet exits TNSR. An example of this is shown in *Outbound NAT*.

Note: NAT is processed *after* ACL rules. For more information, see *ACL and NAT Interaction*.

Note: NAT-specific virtual reassembly parameters have been deprecated in favor of shallow virtual reassembly. See *IP Reassembly*.

15.1 Dataplane NAT Modes

The dataplane has several NAT modes that may be used. This mode is configured via the `dataplane nat mode <mode>` command from `config mode`.

The following modes are available:

- simple** Simple NAT mode. Holds less information for each session, but only works with outbound NAT and static mappings.
- endpoint-dependent** Endpoint-dependent NAT mode. The default mode. Uses more information to track each session, which also enables additional features such as `out-to-in-only` and `twice-nat`.
- deterministic** Deterministic NAT (CGN) mode. Used for large-scale deployments with a focus on performance at a cost of using much more memory.

After changing the NAT mode, the dataplane must be restarted with `service dataplane restart`.

Note: There must be at least one `inside` and `outside` interface for NAT to function, see [Network Address Translation](#) and [Outbound NAT](#) for more details.

15.1.1 Simple NAT

Simple NAT is the most basic NAT mode. It tracks sessions in a hash table using four items:

- Source IP address
- Source port
- Protocol
- FIB table index

Simple NAT has a couple basic options that may be adjusted using the `dataplane nat mode-options simple <option>` command:

out2in-dpo Enables out-to-in DPO

static-mapping-only Static mapping only, disables dynamic translation of connections.

15.1.2 Endpoint-dependent NAT

Endpoint-dependent NAT mode is the default NAT mode on TNSR. Endpoint-dependent NAT mode tracks more information about each connection. As suggested by the name, the key difference is in tracking the destination of the connection:

- Source IP address
- Source port
- Target IP address
- Target port
- Protocol
- FIB table index

Some NAT features require this extra information, notably `out-to-in-only` and `twice-nat`.

15.1.3 Deterministic NAT

Deterministic NAT mode, also known as Carrier-Grade NAT (CGN) mode, is geared for maximum performance at a large scale. This performance comes at a price, however, in that it consumes greater amounts of memory to achieve its goals.

For more information on Deterministic NAT, see [Deterministic NAT](#).

15.2 NAT Options

The NAT options described here control TNSR NAT behavior independent of the chosen mode.

15.2.1 NAT Forwarding

When NAT is active, it will affect traffic to and from services on TNSR, such as IPsec and BGP. When NAT is enabled, by default TNSR will drop traffic that doesn't match an existing NAT session or static NAT rule. To change this behavior, enable NAT forwarding mode:

```
tnsr(config)# nat global-options nat44 forwarding true
```

If NAT is active and there are **no** services present on TNSR which need to communicate using an interface involved with NAT, then it is more secure and efficient to disable forwarding:

```
tnsr(config)# nat global-options nat44 forwarding false
```

15.2.2 NAT Session Timeout Duration

The `tnsr nat global-options timeouts (icmp|tcp_established|tcp_transitory|udp) <seconds>` command controls how long NAT sessions in various states will be retained while idle (no packets passing which match the session entry).

Longer session idle timeouts are friendlier to user connections, at the expense of resource consumption required to retain the NAT sessions for long periods.

The following timeout values can be changed:

icmp Idle timeout for ICMP sessions (e.g. Echo/ping). The default value is 60 seconds.

tcp_established Idle timeout for established TCP connections. Established connections should rarely be forced down in most use cases, so a long timeout is best for this value. The default value is 7440 seconds (2 hours, 4 minutes). It is common to see this set as high as 86400 (24 hours) in deployments with long-lived idle connections.

tcp_transitory Idle timeout for TCP connections which are not fully established (being setup or torn down). The default value is 240 seconds (4 minutes) which is typically sufficient.

udp Idle timeout for UDP sessions. Since UDP is technically stateless and has no formal setup/tear-down for sessions, there is no way for TNSR to determine if a UDP "connection" is established or finished. The default value is 300 seconds (5 minutes) which, combined with client and server keep-alives, is typically sufficient.

A longer idle timeout may be required in certain cases, such as for VoIP connections passing through which expect to reuse specific source ports.

In deployments with many short-lived UDP connections, such as DNS queries, lowering the timeout will help manage session usage/turnover more efficiently.

The `show nat config` command output includes the current timeout values.

15.3 NAT Pool Addresses

Before TNSR can perform any type of NAT, an `inside` and `outside` interface must be set and the outside/external addresses (e.g. WAN-side) must be listed in a NAT pool. These pools are added from configure mode (*Configuration Mode*) in the TNSR CLI (*Entering the TNSR CLI*).

For a single external address, define a NAT pool like so:

```
tnsr(config)# nat pool addresses 203.0.113.2
```

For multiple addresses, use a range:

```
tnsr(config)# nat pool addresses 203.0.113.2 - 203.0.113.5
```

TNSR also supports using an interface to automatically determine the pool addresses:

```
tnsr(config)# nat pool interface GigabitEthernet0/14/1
```

For *Outbound NAT* this is typically the interface set as `ip nat outside`.

15.4 Outbound NAT

Outbound NAT, sometimes referred to as Source NAT, Overload NAT or Port Address Translation (PAT), changes the source address and port of packets exiting a given interface. This is most commonly performed in order to hide the origin of a packet, allowing multiple IPv4 hosts inside a network to share one, or a limited number of, external or outside addresses on a router.

In TNSR, this type of NAT is configured by marking the LAN or internal interface as `inside` and the WAN or external interface as `outside`, for example:

```
tnsr(config)# nat pool addresses 203.0.113.2
tnsr(config)# interface GigabitEthernet0/14/1
tnsr(config-interface)# ip nat outside
tnsr(config-interface)# exit
tnsr(config)# interface GigabitEthernet0/14/2
tnsr(config-interface)# ip nat inside
tnsr(config-interface)# exit
tnsr(config)# nat global-options nat44 forwarding true
tnsr(config)#
```

Traffic originating on the inside interface and exiting the outside interface will have its source address changed to match that of the outside interface.

Warning: The address of the outside interface **must** exist as a part of a NAT pool (*NAT Pool Addresses*) or connectivity from the inside interface will not function with NAT configured. Use either an address pool as shown above, or `nat pool interface <name>` where <name> is the same interface that contains `ip nat outside`.

Warning: When activating `ip nat outside`, services on TNSR may fail to accept or initiate traffic on that interface depending on the NAT mode. For services on TNSR to function in combination with `ip nat outside`, endpoint-dependent NAT mode must be enabled. In TNSR 18.11 and later, this is the default mode.

The following commands set TNSR to endpoint-dependent NAT mode:

```
tnsr(config)# dataplane nat mode endpoint-dependent
tnsr(config)# service dataplane restart
```

Additionally, NAT forwarding must be enabled for this traffic to be accepted by TNSR. See *NAT Forwarding* for details.

15.5 Static NAT

Static NAT entries alter traffic, redirecting it to a static host on an internal network, or mapping it to a static address on the way out:

```
tnsr(config)# nat pool addresses <external address>
tnsr(config)# nat static mapping (icmp|tcp|udp) local <local address> [local port]
                        external (external address|external interface) [external port]
                        [twice-nat] [out-to-in-only] [route-table <rt-tbl-name>]
```

There are two common use cases for static NAT in practice: Port Forwarding and 1:1 NAT.

Warning: Remember to add the address of the outside interface as a part of a NAT pool (*NAT Pool Addresses*) or the static NAT entry will fail to commit.

Warning: The out-to-in-only and twice-nat features require endpoint-dependent NAT mode. In TNSR 18.11 and later, this is the default mode.

The following commands set TNSR to endpoint-dependent NAT mode:

```
tnsr(config)# dataplane nat mode endpoint-dependent
tnsr(config)# service dataplane restart
```

15.5.1 Port Forwards

Port forwards redirect a port on an external NAT pool address to a port on a local host. A port forward is accomplished by specifying ports in the static NAT command:

```
tnsr(config)# nat pool addresses 203.0.113.2
tnsr(config)# nat static mapping tcp local 10.2.0.5 22 external 203.0.113.2 222
```

In the above example, a TCP connection to port 222 on 203.0.113.2 will be forwarded to port 22 on 10.2.0.5. The source address remains the same.

15.5.2 1:1 NAT

1:1 NAT, also called One-to-One NAT or in some cases “Network Address Translation”, maps all ports of an external address for a given protocol to an internal address. This mapping works for inbound and outbound packets. To create a 1:1 mapping, make a static NAT entry which does not specify any ports:

```
tnsr(config)# nat pool addresses 203.0.113.3
tnsr(config)# nat static mapping tcp local 10.2.0.5 external 203.0.113.3
```

15.5.3 Twice NAT

Twice NAT changes both the source and destination address of inbound connection packets. This works similar to a static NAT port forward, but requires an additional NAT address specification.

First, add the internal address for source translation:

```
tnsr(config)# nat pool addresses 10.2.0.2 twice-nat
```

Next, add the external address to which the client originally connects:

```
tnsr(config)# nat pool addresses 203.0.113.2
```

Finally, add the static mapping which sets up the destination translation:

```
tnsr(config)# nat static mapping tcp local 10.2.0.5 22 external 203.0.113.2 222 twice-  
→nat
```

In the above example, a TCP connection to port 222 on 203.0.113.2 will be forwarded to port 22 on 10.2.0.5. When the packet leaves TNSR, the source is translated so the connection appears to originate from 10.2.0.2 using a random source port.

Warning: This feature requires endpoint-dependent NAT mode. In TNSR 18.11 and later, this is the default mode.

The following commands set TNSR to endpoint-dependent NAT mode:

```
tnsr(config)# dataplane nat mode endpoint-dependent  
tnsr(config)# service dataplane restart
```

15.6 Dual-Stack Lite

Dual-Stack Lite, also known as DS-Lite, is a mechanism which facilitates large scale IPv4 NAT by encapsulating IPv4 packets inside IPv6 packets for delivery to a Carrier-Grade NAT (CGN) endpoint. This allows providers to provision end users with only a routed IPv6 address, and any IPv4 traffic is carried through IPv6 to a CGN device. Once the IPv6 packet reaches the CGN device, the IPv4 packet is extracted, has NAT applied, and is forwarded. The CGN device will apply NAT using one of its routable IPv4 addresses, shared between DS-Lite users.

By using encapsulation, DS-Lite avoids multiple layers of NAT between the customer and the Internet. An end-user network which connects to a DS-Lite provider should not perform any IPv4-IPv4 NAT on the traffic before it reaches a router configured for DS-Lite.

DS-Lite is considered an IPv6 transition mechanism as it allows providers to reduce their dependence on scarce IPv4 routable addresses, while still giving clients full access to IPv4 and IPv6 resources. It also removes the need to use potentially conflicting IPv4 private address space for IPv4 routing inside a provider network.

There are two endpoints to DS-Lite connections:

- DS-Lite Basic Bridging BroadBand (B4) element on the customer end
- DS-Lite Address Family Transition Router (AFTR) element at the provider end

From a customer perspective, their side is before (B4) DS-Lite and the ISP side is after (AFTR) DS-Lite.

TNSR can operate in either capacity: As a CPE DS-Lite B4 client endpoint, or as an AFTR endpoint providing DS-Lite connectivity and IPv4 NAT to clients.

15.6.1 Acting as a B4 Endpoint

For a customer premise equipment (CPE) role which connects to an ISP offering DS-Lite service, the following steps are required:

First, configure IPv6 connectivity to the ISP and local/"LAN" IPv4 with a private address.

Next, enable DS-Lite NAT and restart the dataplane:

```
tnsr(config)# dataplane nat dslite-ce
tnsr(config)# service dataplane restart
```

Next, configure the local IPv6 address TNSR will use for its DS-Lite B4 endpoint. This tunnel address will be allocated by the ISP:

```
tnsr(config)# dslite b4 endpoint <customer-b4-ip6-address>
```

Next, configure the remote IPv6 DS-Lite tunnel AFTR endpoint address given by the ISP:

```
tnsr(config)# dslite aftr endpoint <isp-aftr-ip6-address>
```

Finally, configure a route to the AFTR endpoint address through the interface connected to the ISP.

```
tnsr(config)# route ipv6 table ipv6-VRF:0
tnsr(config-route-table-v6)# route <isp-aftr-ip6-address>
tnsr(config-rttbl6-next-hop)# next-hop 0 via <isp-dslite-gateway> <isp-interface>
```

15.6.2 Acting as an AFTR Endpoint

For a provider role as a DS-Lite AFTR endpoint serving customers, the following steps are required:

First, configure IPv6 and IPv4 connectivity such that this TNSR instance has both IPv6 and IPv4 connectivity to the Internet.

Next, configure the local AFTR IPv6 tunnel address TNSR will use to receive DS-Lite encapsulated packets from customer equipment:

```
tnsr(config)# dslite aftr endpoint <isp-b4-ip6-address>
```

Next, configure one or more routable ("public") IPv4 addresses for the DS-Lite NAT pool. These addresses are used by TNSR to apply NAT to outgoing IPv4 traffic which arrived via DS-Lite:

```
tnsr(config)# dslite pool address <ipv4-addr-first> [- <ipv4-addr-last>]
```

Finally, add a route back to the customer B4 tunnel endpoint via the router address for that customer.

```
tnsr(config)# route ipv6 table ipv6-VRF:0
tnsr(config-route-table-v6)# route <customer-b4-ip6-address>
tnsr(config-rttbl6-next-hop)# next-hop 0 via <customer-ip6-address> <customer-
↪interface>
```

IPv4 packets arriving through DS-Lite from a customer will be removed from the encapsulation, have NAT applied, and then be forwarded upstream (e.g. to the Internet). Reply packets will come back, and then go back through NAT and DS-Lite to reach customers.

15.6.3 DS-Lite Status

To view active DS-Lite sessions, use the following command:

```
tnsr# show dslite
```


15.7 Deterministic NAT

Deterministic NAT mode, also known as Carrier-Grade NAT (CG-NAT), is geared for maximum performance at a large scale. This performance comes at a price, however, in that it consumes greater amounts of memory than other, simpler modes.

To switch the NAT mode used by TNSR, see [Dataplane NAT Modes](#).

Deterministic NAT pre-allocates 1000 external ports per inside address, which can increase memory requirements significantly. Each single session requires approximately 15 Bytes of memory.

Deterministic NAT enforces maximum numbers of NAT sessions per user, and only works for TCP, UDP, and ICMP protocols.

Deterministic NAT requires a mapping, configured as follows:

```
tnsr(config)# nat deterministic mapping inside <inside-prefix> outside <outside-
->prefix>
```

In this command, the parameters to replace are:

inside <inside-prefix> The internal subnet containing local users, for example, 198.18.0.0/15.

outside <outside-prefix> The external subnet to which these users will be mapped using deterministic NAT. For example, 203.0.113.128/25.

Configured mappings may be viewed as follows:

```
tnsr(config)# show nat deterministic-mappings
Deterministic Mappings
-----

```

Inside	Outside	Ratio	Ports	Sessions
198.14.0.0/15	203.0.113.128/25	1024	63	0

```

NAT Reassembly Parameters
-----

```

15.8 NAT Status

TNSR offers several ways to view the active NAT configuration, rules, and sessions. These start with `nat show`, and are all available in `config` and `master` mode.

15.8.1 View NAT Configuration

To view the current NAT configuration parameters (not rules), use `show nat config`:

```
tnsr# show nat config

NAT Configuration Parameters
-----
translation hash buckets 1024
translation hash memory 134217728
deterministic false
```

(continues on next page)

(continued from previous page)

```

user hash buckets 128
user hash memory 67108864
max translations per user 100
outside Route Table ipv4-VRF:0
inside Route Table ipv4-VRF:0
dynamic mapping enabled
forwarding is enabled
UDP timeout 300s
TCP established connections timeout 7440s
TCP transitory connections timeout 240s
ICMP timeout 60s

```

15.8.2 View Static Mappings

To view currently configured static NAT mappings, use `show nat static-mappings`:

```

tnsr# show nat static-mappings

Static Mappings

Proto Local IP      Port External IP Port Interface Twice NAT Out to In Route Table
-----
tcp 10.2.0.5         22 203.0.113.2  222                               ipv4-VRF:0

```

15.8.3 View Deterministic Mappings

To view currently configured deterministic NAT mappings, use `show nat deterministic-mappings`:

```

tnsr# show nat deterministic-mappings
Deterministic Mappings
-----

Inside      Outside          Ratio    Ports    Sessions
-----
198.14.0.0/15 203.0.113.128/25 1024      63       0

```

15.8.4 View Dynamic Configuration

To view the IP addresses or interfaces currently assigned for use by NAT, use `show nat dynamic addresses` or `show nat dynamic interfaces`, depending on the TNSR NAT configuration:

```

tnsr# show nat dynamic addresses

Pool Addresses  Route Table    Twice NAT
-----
203.0.113.2

```

15.8.5 View Interfaces

To view the interfaces which are currently marked as inside and outside for NAT purposes, use `show nat interface-sides`:

```
tnsr# show nat interface-sides
```

Interfaces	Side
GigabitEthernet0/14/0	outside
GigabitEthernet3/0/0	inside

15.8.6 View NAT Sessions

To view a summary of outgoing NAT sessions by source address, use `show nat sessions`:

```
tnsr# show nat sessions
```

```
NAT sessions
```

IP address	Static	Dynamic	Route Table
10.2.0.1	0	4	ipv4-VRF:0
203.0.113.2	0	1	ipv4-VRF:0

To see more detail for each specific session, add `verbose` to the previous command, which becomes `show nat sessions verbose`:

```
tnsr# show nat sessions verbose
```

```
NAT sessions detail
```

Proto	Inside/Outside/Ext	Type	Route Table	Last used	Bytes/pkts
udp	10.2.0.1:123	dynamic	ipv4-VRF:0	143	498
	203.0.113.2:16253				6
	52.6.160.3:123				
udp	10.2.0.1:123	dynamic	ipv4-VRF:0	143	498
	203.0.113.2:18995				6
	184.105.182.7:123				
udp	10.2.0.1:123	dynamic	ipv4-VRF:0	145	498
	203.0.113.2:53893				6
	69.36.182.57:123				
udp	10.2.0.1:123	dynamic	ipv4-VRF:0	207	498
	203.0.113.2:44109				6
	198.50.238.163:123				

15.9 NAT Examples

The examples in this section describe and demonstrate use cases and packet flows for typical scenarios involving NAT.

15.9.1 AWS NAT Examples

When using TNSR with AWS, it is relatively easy to unintentionally create an asymmetric routing situation. AWS knows about your local networks and will happily egress traffic with NAT for them, when other networking setups

would otherwise drop or fail to hand off the traffic.

The examples in this section covers what would happen with a TNSR setup in AWS with two instances: An internal LAN instance with a local “client” system making an outbound request, and an external WAN instance that is intended to handle public-facing traffic. TNSR sits between the WAN and LAN instance to route traffic. In AWS, the VPC routing table is configured such that the LAN instance uses TNSR for its default gateway. The expected flow is that traffic flows from clients, through TNSR, to the Internet and back the same path.

This table lists the networks and addresses used by these examples.

Item	Value
AWS Networks	192.0.2.0/24 (LAN), 198.18.5.0/24 (WAN), 203.0.113.0/24 (External)
AWS Gateways	192.0.2.1 (LAN), 198.18.5.1 (WAN), 203.0.113.1 (External)
TNSR LAN	192.0.2.2/24
TNSR WAN	198.18.5.2
TNSR GW	198.18.5.1 (AWS Gateway)
LAN Client	192.0.2.5/24
LAN Client GW	192.0.2.2 (TNSR LAN)
Server	198.51.100.19/24
Server GW	198.51.100.1

AWS Example without NAT

In this example, TNSR is not configured to perform NAT. This example steps through each portion of a packet and its reply, and then discusses the problems at the end.

First, the client initiates a connection using a packet which arrives on the TNSR LAN interface

Proto	Source	Destination	Via
TCP	192.0.2.5:1025	198.51.100.19:443	192.0.2.2

TNSR performs a FIB lookup. The destination IP address is not within the the subnets configured on the TNSR instance interfaces, so it matches the default route

Proto	Source	Destination	Via
TCP	192.0.2.5:1025	198.51.100.19:443	Default

TNSR forwards the packet out its WAN interface to its default gateway on the WAN. TNSR is not configured for NAT, thus it does not perform any translation.

Proto	Source	Destination	Via
TCP	192.0.2.5:1025	198.51.100.19:443	198.18.5.1

The packet reaches the AWS internet gateway connected to the VPC. Its source IP address is still the private IP address of the LAN instance.

Proto	Source	Destination	Via
TCP	192.0.2.5:1025	198.51.100.19:443	198.18.5.1

The AWS internet gateway performs NAT. It recognizes the source IP address as belonging to the LAN instance and rewrites it to the public IP address of the LAN instance.

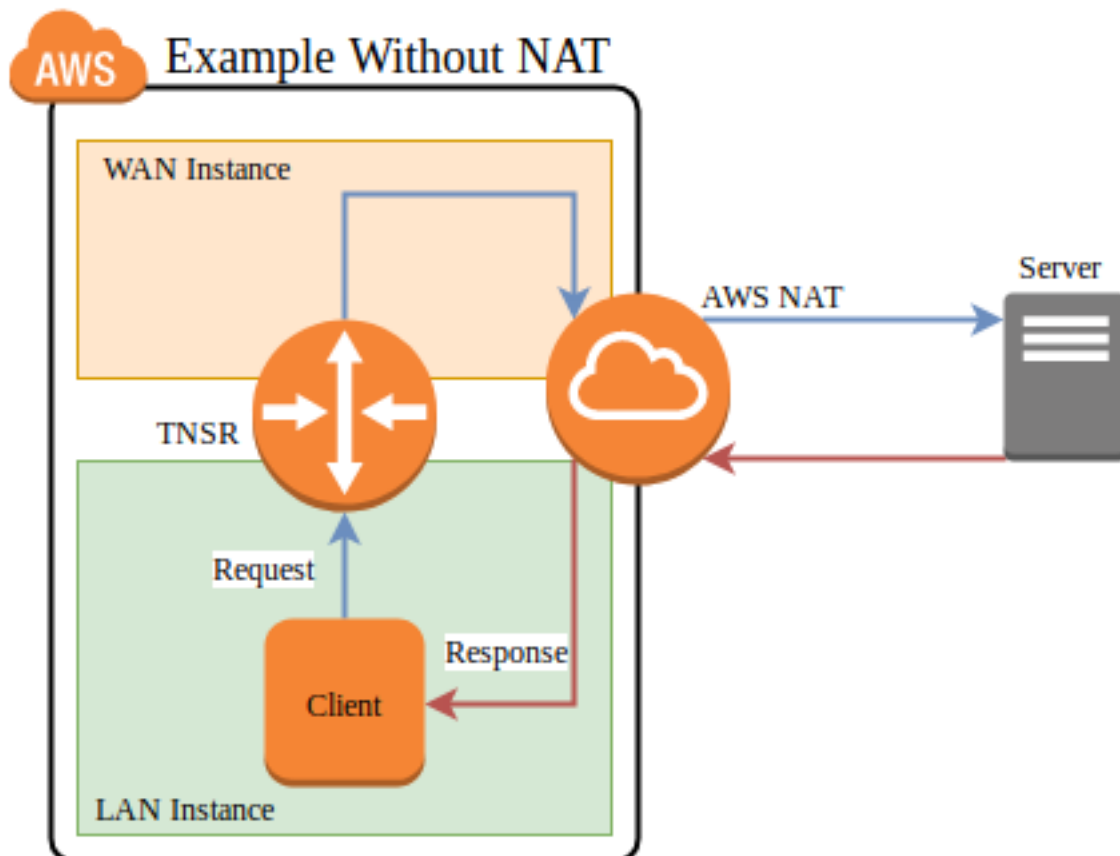


Fig. 1: AWS example packet flow without NAT

Proto	Source	Destination	Via
TCP	203.0.113.50:40250	198.51.100.19:443	Default

The AWS internet gateway forwards the packet to the internet.

Proto	Source	Destination	Via
TCP	203.0.113.50:40250	198.51.100.19:443	203.0.113.1

The destination host sends a reply to the public IP address of the LAN instance. It arrives at the AWS internet gateway.

Proto	Source	Destination	Via
TCP	198.51.100.19:443	203.0.113.50:40250	198.51.100.1

The AWS internet gateway performs NAT. It recognizes the destination IP address as belonging to LAN instance and rewrites it to the private IP address of the LAN instance.

Proto	Source	Destination	Via
TCP	198.51.100.19:443	192.0.2.5:1025	Direct L2 LAN

The AWS internet gateway knows how to reach the private IP address of the LAN instance directly, so it forwards the reply packet directly to the LAN instance, skipping the TNSR instance.

Proto	Source	Destination	Via
TCP	198.51.100.19:443	192.0.2.5:1025	Direct L2 LAN

The packet arrives at the client.

The return path skipped TNSR, so TNSR is only seeing half the packets for the connection. At best this means the asymmetric routing will bypass any filtering or inspection of the replies (IDS/IPS), and at worst it could mean subsequent packets would be dropped instead of passing through TNSR.

AWS Example with NAT

In this example, TNSR has NAT configured such that its LAN is defined as an `inside` interface and its WAN is an `outside` interface. See [Outbound NAT](#) for details. Packets leaving the WAN will be translated such that they leave with a source address set to the TNSR WAN interface IP address.

First, the client initiates a connection using a packet which arrives on the TNSR LAN interface

Proto	Source	Destination	Via
TCP	192.0.2.5:1025	198.51.100.19:443	192.0.2.2

TNSR performs a FIB lookup. The destination IP address is not within the the subnets configured on the TNSR instance interfaces, so it matches the default route

Proto	Source	Destination	Via
TCP	192.0.2.5:1025	198.51.100.19:443	Default

TNSR applies NAT and forwards the packet out its WAN interface to its default gateway on the WAN subnet.

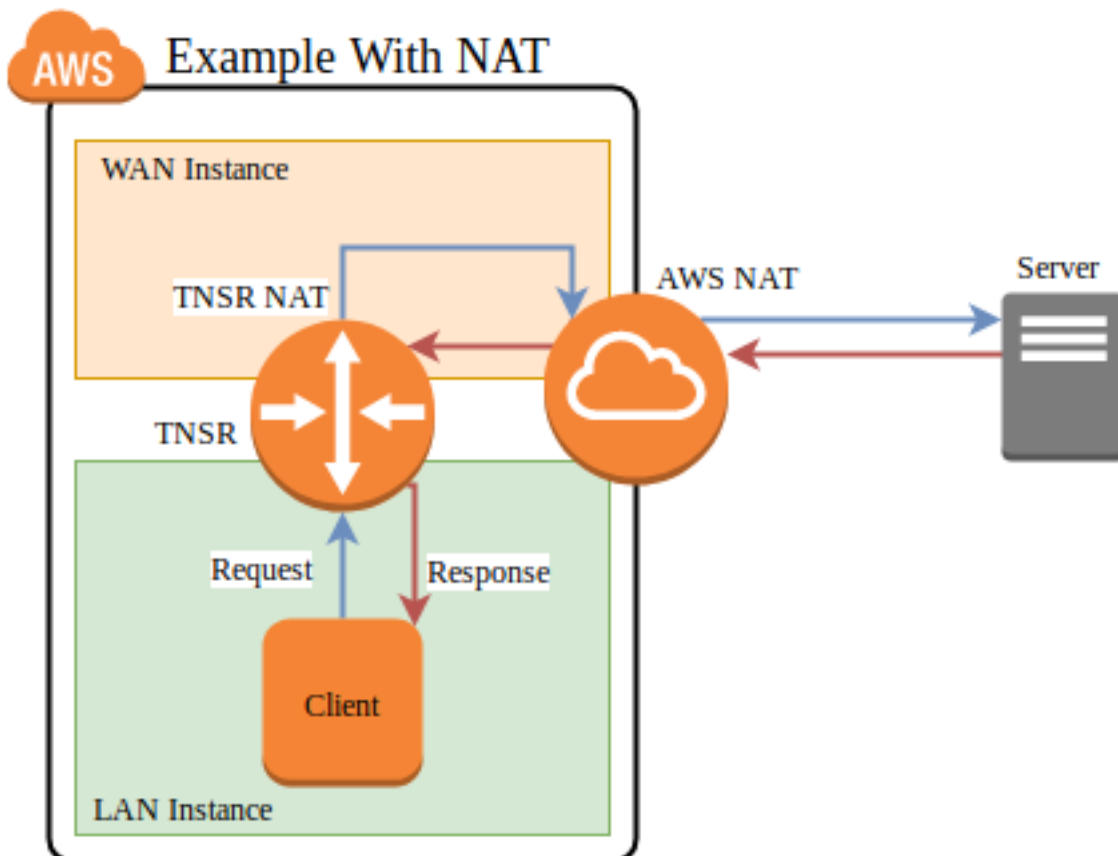


Fig. 2: AWS example packet flow with NAT

Proto	Source	Destination	Via
TCP	198.18.5.2:34567	198.51.100.19:443	198.18.5.1

The packet reaches the AWS internet gateway connected to the VPC. Its source IP address is the private IP address of the TNSR WAN instance.

Proto	Source	Destination	Via
TCP	198.18.5.2:34567	198.51.100.19:443	198.18.5.1

The AWS internet gateway performs NAT. It recognizes the source IP address as belonging to the WAN instance and rewrites it to the public IP address of the WAN instance.

Proto	Source	Destination	Via
TCP	203.0.113.50:40250	198.51.100.19:443	Default

The AWS internet gateway forwards the packet to the internet.

Proto	Source	Destination	Via
TCP	203.0.113.50:40250	198.51.100.19:443	203.0.113.1

The destination host sends a reply to the public IP address of the WAN instance. It arrives at the AWS internet gateway.

Proto	Source	Destination	Via
TCP	198.51.100.19:443	203.0.113.50:40250	198.51.100.1

The AWS internet gateway performs NAT. It recognizes the destination IP address as belonging to WAN instance and rewrites it to the private IP address of the WAN instance. The AWS internet gateway knows how to reach the private IP address of the WAN instance directly, so it forwards the reply packet directly to the WAN instance.

Proto	Source	Destination	Via
TCP	198.51.100.19:443	198.18.5.2:34567	Direct L2 WAN

The packet arrives at the TNSR WAN, which performs NAT. It recognizes the source and destination as matching an existing NAT state belonging to the LAN client and rewrites the destination address to the LAN client. TNSR knows how to reach the client LAN IP address directly, so it forwards the reply packet.

Proto	Source	Destination	Via
TCP	198.51.100.19:443	192.0.2.5:1025	Direct L2 LAN

The packet arrives back at the client.

In this case, the NAT performed on TNSR ensured that the AWS gateway delivered the reply back to TNSR instead of handing it off directly. This allowed the packet and its reply to use the same path outbound and inbound.

MAP (MAPPING OF ADDRESS AND PORT)

MAP is short for Mapping of Address and Port. It is a carrier-grade IPv6 transition mechanism capable of efficiently transporting high volumes of IPv4 traffic across IPv6 networks.

There are two MAP implementations in TNSR Enterprise: MAP-T which uses translation and MAP-E which uses encapsulation.

With MAP, IPv4 requests are forwarded from an end user Customer Edge (CE) device through an IPv6 Border Relay (BR) router which processes and forwards the requests to IPv4 destinations. Customer IPv6 requests can proceed directly to IPv6 destinations without going through the BR, which lowers the burden on the BR.

MAP is stateless, thus capable of handling large scale traffic volume without additional overhead for tracking individual connections. Each CE device receives a public IPv4 address but may only use a specific port range on that address. In this way, multiple users may share a public address without an additional layer of NAT. Since this relationship is predetermined, the ports are also available bidirectionally, which is not possible with other solutions such as Carrier-Grade NAT/NAT444.

MAP-T and MAP-E require port information to operate, thus fragments must be reassembled at the BR before forwarding. This is due to the fact that protocol and port information are only present in the first packet. Intelligent caching & forwarding may be employed for handling fragments.

TNSR can currently act as a BR, providing service to CE clients.

16.1 MAP Configuration

MAP configurations consist of MAP domains, MAP rules, and interface configuration.

16.1.1 MAP Domains

A MAP domain encompasses a set of addresses, translation parameters, and MAP rules. Groups of CE devices belong to specific MAP domains.

A MAP domain is created in config mode using the `nat nat64 map <domain name>` command from within config mode. That command enters `config-map` mode.

This mode, `config-map`, contains a number of MAP options specific to a MAP domain:

description A short text description noting the name or purpose of this MAP domain.

port-set <length><offset> A port set is, as the name implies, a set of ports. This is typically divided up into multiple sets of ports, the exact size and ranges of which are calculated using the port set length and offset, discussed next. With MAP, users are overloaded onto a single IP address, with different port sets on a single IP address being allocated to multiple users. In this way, users can share individual IP addresses but only have access to specific ranges of ports.

port-set length <psid-length> Determines the number of port sets to allocate inside the available 16-bit port range (1-65536). A larger port set length allows for more users to share an address, but allocates them each a smaller number of ports. For example, a port set length of 8 uses 8 bits to define the port set, leaving the remaining 8 bits for use by each customer, or 256 ports each.

port-set offset <psid-offset> Determines the position of the port set identifier inside the available bits which represent the port. An offset of 0 means the identifier is first, and the ports per user will be contiguous. Placing the offset in the middle of the available space will allow users to utilize multiple ranges that are not contiguous, but each user will have slightly less ports available. For example, with a port set length of 8, but an offset of 2, each user can utilize only 192 ports instead of 256, since it is split into three ranges of 64 ports each. The offset cannot be larger than the port set length subtracted from the total available bits (16).

There are minor security benefits when using multiple non-contiguous port ranges since it is more difficult for an attacker to guess which ports belong to a given customer, but the loss of port capacity may outweigh this benefit in most environments.

embedded-address bit-length <ea-width> The Embedded Address Bits value is the sum of the bits needed for the IPv4 prefix and the port set length. For example, if the IPv4 prefix is a /24, that requires 8 bits to embed and allows 256 addresses for users. A port set length of 8 allows for 256 port sets. With a port set offset of 0, this yields a maximum of 65,536 users sharing 256 IPv4 addresses, each of which can use 256 ports.

Note: To utilize MAP rules, this value must be 0.

ipv4 prefix <ip4-prefix> The IPv4 Prefix is available pool of IPv4 addresses which can be utilized by MAP clients. The size of this prefix must be represented in the Embedded Address Bits. For example, a /24 prefix network requires 8 bits to uniquely identify an address.

ipv6 prefix <ip6-prefix> The IPv6 prefix contains the range of possible addresses assigned to clients. The end-user network must be at least a 64 prefix, leaving 64 bits to represent both this prefix and the embedded address bits. The smallest possible IPv6 prefix will be 128 bits less the sum of the end user network and embedded address bits. For example, with an embedded address length of 16, 48 bits remain for the IPv6 prefix. Shorter prefixes (e.g. 44) allow for additional IPv6 subnets to be assigned to clients.

ipv6 source <ip6-src> The IPv6 source address on the router used as the MAP domain BR address and Tunnel source. This address should exist on the interface used for mapping. For MAP-T, this must have a prefix length of either /64 or /96. For MAP-E, this is a single address (/128) and not a prefix.

mtu <mtu-val> The Maximum Transmission Unit (MTU) is the largest packet which can traverse the link without fragmentation. This must be set appropriately due to the importance of MAP fragment handling, as required information to calculate targets is only in the first packet and not additional fragments.

16.1.2 MAP Rules

MAP rules exist inside a MAP domain and are configured from within `config-map` mode. MAP rules map specific port sets to specific MAP CE end user addresses. These are 1:1 manual mappings and take the place of automatic calculation, and as such to use MAP rules, the `embedded-address bit-length` must be 0.

A map rule takes the following form:

```
rule port-set <psid> ipv6-destination <ip6-destination>
```

The components of a rule are:

port-set <psid> The port set ID (PSID) to match for this rule.

ipv6-destination <ip6-destination> The MAP CE IPv6 address to associate with this specific port set ID.

16.1.3 MAP Interface Configuration

TNSR must be told which interface is used with MAP, and how that interface will operate.

Within config-interface mode (*Configure Interfaces*), there are two possible settings for MAP:

map <enable|disable> Enables or disables MAP for this interface.

map translate When present and MAP is enabled, the interface operates in translate mode (MAP-T). When not set, encapsulation is used instead (MAP-E).

16.1.4 View MAP Configuration

The MAP configuration can be viewed with the `show map [<map-domain-name>]` command. Without a given domain name, information is printed for all MAP domains, plus the MAP parameters.

```
tnsr# show map cpoc
```

Name	IP4 Prefix	IP6 Prefix	IP6 Src Pref	EA Bits	PSID Off	PSID Len
↪MTU						
↪---						
cpoc	192.168.1.0/24	2001:db8::/32	1234:5678:90ab:cdef::/64	16	6	4
↪1280						

```
tnsr# show map
MAP Parameters
```

```
-----
Fragment: outer
Fragment ignore-df: false
ICMP source address: 0.0.0.0
ICMP6 unreachable msgs: disabled
Pre-resolve IPv4 next hop: 0.0.0.0
Pre-resolve IPv6 next hop: ::
Security check enabled: true
Security check fragments enabled: false
Traffic-class copy: enabled
Traffic-class value: 0
```

Name	IP4 Prefix	IP6 Prefix	IP6 Src Pref	EA Bits	PSID Off	PSID Len
↪MTU						
↪---						
cpoc	192.168.1.0/24	2001:db8::/32	1234:5678:90ab:cdef::/64	16	6	4
↪1280						

16.2 MAP Parameters

MAP Parameters control the behavior of MAP-T and MAP-E. These parameters are configured by the `nat nat64 map parameters` command from within `config` mode, which enters `config-map-param` mode where the individual values are set.

From within `config-map-param` mode, the following commands are available:

fragment ignore-df Allows TNSR to perform IPv4 fragmentation even when packets contain the do-not-fragment (DF) bit. This improves performance by moving the burden of fragmentation to the endpoint rather than the MAP relay.

fragment (inner|outer) Controls whether TNSR will fragment the inner (encapsulated or translated) packets or the outer (tunnel) packets.

icmp source-address <ipv4-address> Sets the IPv4 address used by TNSR to send relayed ICMP error messages.

icmp6 unreachable-msgs (enable|disable) When enabled, TNSR will generate ICMPv6 unreachable messages when a packet fails to match a MAP domain or fails a security check.

pre-resolve (ipv4|ipv6) next-hop <ip46-address> Manually configures the next hop for IPv4 or IPv6 routing of MAP traffic, which bypasses a routing table lookup. This increases performance, but means that the next hop cannot be determined dynamically or by routing protocol.

security-check (enable|disable) Enables or disables validation of decapsulated IPv4 addresses against the external IPv6 address on single packets or the first fragment of a packet. Disabling the check increases performance but potentially allows IPv4 address spoofing.

security-check fragments (enable|disable) Extends the previous security check to all fragments instead of only inspecting the first packet.

tcp mss <mss-value> Sets the MSS value for MAP traffic, typically the MTU less 40 bytes.

traffic-class tc <tc-val> Sets the Class/TOS field of outer IPv6 packets to the specified value.

traffic-class copy (enable|disable) When enabled, copies the class/TOS field from the inner IPv4 packet header to the outer IPv6 header. This is enabled by default, but disabling can slightly improve performance.

Note: MAP-specific virtual reassembly parameters have been deprecated in favor of shallow virtual reassembly. See [IP Reassembly](#).

16.2.1 View MAP Parameters

The current value of MAP parameters can be displayed by the `show map` command:

```
tnsr# show map
MAP Parameters
-----
Fragment: outer
Fragment ignore-df: false
ICMP source address: 0.0.0.0
ICMP6 unreachable msgs: disabled
Pre-resolve IPv4 next hop: 0.0.0.0
Pre-resolve IPv6 next hop: ::
Security check enabled: true
```

(continues on next page)

(continued from previous page)

```
Security check fragments enabled: false
Traffic-class copy: enabled
Traffic-class value: 0
```

```
Name IP4 Prefix      IP6 Prefix      IP6 Src Pref      EA Bits PSID Off PSID Len
--MTU
-----
cpoc 192.168.1.0/24 2001:db8::/32 1234:5678:90ab:cdef::/64 16      6      4
--1280
```

16.3 MAP Example

16.3.1 Environment

MAP Border Relay	
Item	Value
MAP Domain Name	cpoc
IPv6 Prefix	2001:db8::/32
IPv6 Source Prefix	1234:5678:90ab:cdef::/64
IPv4 Prefix	192.168.1.0/24
Port Set Length	8
Port Set Offset	0
Embedded Address Bits	16
MTU	1300
Interface	GigabitEthernet0/14/0
IPv6 Address	fd01:2::1/64
IPv4 Address	203.0.113.2/24

16.3.2 TNSR Border Relay Configuration

This shows an example Border Relay (BR) configuration in TNSR to provide service to MAP-T Customer Edge (CE) clients. This example assumes some configuration details are already in place, such as the IPv4 prefix already being routed to the BR from upstream, and default routes configured in TNSR for upstream gateways.

First, configure the interface connected to the upstream network. There could be separate interfaces for reaching the Internet and for reaching the CE network, but this example uses a single interface.

```
tnsr(config)# interface GigabitEthernet0/14/0
tnsr(config-interface)# ip address 203.0.113.2/24
tnsr(config-interface)# ipv6 address fd01:2::1/64
tnsr(config-interface)# exit
```

Next, configure the MAP domain:

```
tnsr(config)# nat nat64 map cpoc
tnsr(config-map)# ipv4 prefix 192.168.1.0/24
tnsr(config-map)# ipv6 prefix 2001:db8::/32
tnsr(config-map)# ipv6 source 1234:5678:90ab:cdef::/64
```

(continues on next page)

(continued from previous page)

```
tnsr(config-map)# embedded-address bit-length 16
tnsr(config-map)# port-set length 4
tnsr(config-map)# port-set offset 6
tnsr(config-map)# mtu 1280
tnsr(config-map)# exit
```

Then add a static route:

```
tnsr(config)# route ipv6 table ipv6-VRF:0
tnsr(config-route-table-v6)# route 2001:db8::/32
tnsr(config-rttbl6-next-hop)# next-hop 0 via fd01:2::2 GigabitEthernet0/14/0
tnsr(config-rttbl6-next-hop)# exit
tnsr(config-route-table-v6)# exit
```

Lastly, enable MAP and MAP-T translation for the interface:

```
tnsr(config)# interface GigabitEthernet0/14/0
tnsr(config-interface)# map translate
tnsr(config-interface)# map enable
tnsr(config-interface)# exit
```

See also:

For information on configuring other operating systems to act as a CE, consult their documentation or check the links in *Additional MAP Reading and Tools* for additional information.

16.4 MAP Types

16.4.1 MAP-T (Translation)

With MAP-T, translations are made using mapping rules that can calculate addresses and ports based on information embedded in an IPv6 address, along with several known parameters.

MAP-T clients determine where to send translated IPv4 traffic using the Default Mapping Rule (DMR) IPv6 /64 prefix.

16.4.2 MAP-E (Encapsulation)

MAP-E is similar to MAP-T, but instead of translating IPv4 traffic and encoding information in the address, the IPv4 requests are encapsulated in IPv6 between the CE and BR as described in [RFC 2473](#).

MAP-E clients send all IPv4 encapsulated traffic to the BR IPv6 address.

16.4.3 Additional MAP Reading and Tools

MAP is a complex topic and much of it is outside the scope of TNSR documentation. There are a number of additional resources that have information on MAP along with examples for other operating systems and example environments.

We recommend the following links as starting points for MAP information.

- CableLabs MAP Technical Report [CL-TR-MAP-V01-160630](#)
- Charter MAP-T deployment presentation [MAP-T NANOG Video](#) / [MAP-T NANOG Slides](#)
- Cisco [MAP Simulation Tool](#)

- [MAP-E RFC 7597](#)
- [MAP-T RFC 7599](#)

DYNAMIC HOST CONFIGURATION PROTOCOL

The Dynamic Host Configuration Protocol (DHCP) service on TNSR provides automatic addressing to clients on an interface. Typically, this service uses a local, internal interface such as one connected to a LAN or DMZ.

17.1 DHCP Configuration

The main IPv4 DHCP configuration mode, entered with `dhcp4 server`, defines global options for IPv4 DHCP that affect the general behavior of DHCP as well as options that cover all subnets and pools.

To enter IPv4 DHCP configuration mode, enter:

```
tnsr# configure
tnsr(config)# dhcp4 server
tnsr(config-kea-dhcp4)#
```

From this mode, there are a variety of possibilities, including:

subnet Subnet configuration, see *Subnet Configuration*.

description Description of the DHCP server

option A DHCP Option declaration, see *DHCP Options*.

decline-probation-period <n> Decline lease probation period, in seconds.

echo-client-id <boolean> Controls whether or not the DHCP server sends the client-id back to the client in its responses.

interface listen <if-name> The interface upon which the DHCP daemon will listen. **This is required.**

interface socket (rawludp) Controls whether the DHCP daemon uses raw or UDP sockets.

lease filename <path> Lease database file

lease lfc-interval <n> Lease file cleanup frequency, in seconds.

lease persist <boolean> Whether or not the lease database will persist.

logging <logger-name> Controls which events are logged by the DHCP daemon. Enters `config-kea-dhcp4-log` mode. See *DHCP Logging* for more information.

match-client-id <boolean> When true, DHCP will attempt to match clients first based on client ID and then by MAC address if the client ID doesn't produce a match. When false, it prefers the MAC address.

next-server <IP Address> Specifies a TFTP server to be used by a client.

rebind-timer <n> Sets the period, in seconds, at which a client must rebind its address.

renew-timer <n> Sets the period, in seconds, at which a client must renew its lease.

valid-lifetime <n> The period of time, in seconds, for which a lease will be valid.

Some of these values may be set here globally, and again inside subnets or pools. In each case, the more specific value will be used. For example, if an option is defined in a pool, that would be used in place of a global or subnet definition; A subnet option will be favored over a global option. In this way, the global space may define defaults and then these defaults can be changed if needed for certain areas.

17.1.1 DHCP Options

DHCP Options provide information to clients beyond the basic address assignment. These options give clients other aspects of the network configuration, tell clients how they should behave on the network, and give them information about services available on the network. Common examples are a default gateway, DNS Servers, Network Time Protocol servers, network booting behavior, and dozens of other possibilities.

See also:

For a list of Standard IPv4 DHCP options, see [Standard IPv4 DHCP Options](#). This list also includes the type of data expected and whether or not they take multiple values.

The general form of an option is:

```
tnsr(config-kea-dhcp4)# option <name>
tnsr(config-kea-dhcp4-opt)# data <comma-separated values>
tnsr(config-kea-dhcp4-opt)# exit
```

This example defines a global domain name for all clients in all subnets:

```
tnsr(config-kea-dhcp4)# option domain-name
tnsr(config-kea-dhcp4-opt)# data example.com
tnsr(config-kea-dhcp4-opt)# exit
```

This example defines a default gateway for a specific subnet:

```
tnsr(config-kea-subnet4)# option routers
tnsr(config-kea-subnet4-opt)# data 10.2.0.1
tnsr(config-kea-subnet4-opt)# exit
```

To see a list of option names, enter:

```
tnsr(config-kea-dhcp4)# option ?
```

When defining options the data can take different forms. The DHCP daemon uses comma-separated value (CSV) format by default and it will automatically convert the text representation of a value to the expected data in the daemon.

Inside the `option` configuration mode, the following choices are available:

always-send <boolean> Controls whether the DHCP server will always send this option in a response, or only when requested by a client. The default behavior varies by option and is documented in [Standard IPv4 DHCP Options](#)

csv-format <boolean> Toggles between either CSV formatted data or raw binary data. This defaults to `true` unless an option does not have a default definition. In nearly all cases this option should be left at the default.

data <data> Arbitrary option data. Do not enclose in quotes. To see option data types and expected formats, see [Standard IPv4 DHCP Options](#)

space <name> Option space in which this entry exists, defaults to `dhcp4`.

Standard IPv4 DHCP Options

This list contains information about the standard IPv4 DHCP options, sourced from the [Kea Administrator Manual](#) section on [DHCP Options](#).

For a list of the Types and their possible values, see [DHCP Option Types](#).

Name	Code	Type	Array	Always Return
time-offset	2	int32	false	false
routers	3	ipv4-address	true	true
time-servers	4	ipv4-address	true	false
name-servers	5	ipv4-address	true	false
domain-name-servers	6	ipv4-address	true	true
log-servers	7	ipv4-address	true	false
cookie-servers	8	ipv4-address	true	false
lpr-servers	9	ipv4-address	true	false
impress-servers	10	ipv4-address	true	false
resource-location-servers	11	ipv4-address	true	false
boot-size	13	uint16	false	false
merit-dump	14	string	false	false
domain-name	15	fqdn	false	true
swap-server	16	ipv4-address	false	false
root-path	17	string	false	false
extensions-path	18	string	false	false
ip-forwarding	19	boolean	false	false
non-local-source-routing	20	boolean	false	false
policy-filter	21	ipv4-address	true	false
max-dgram-reassembly	22	uint16	false	false
default-ip-ttl	23	uint8	false	false
path-mtu-aging-timeout	24	uint32	false	false
path-mtu-plateau-table	25	uint16	true	false
interface-mtu	26	uint16	false	false
all-subnets-local	27	boolean	false	false
broadcast-address	28	ipv4-address	false	false
perform-mask-discovery	29	boolean	false	false
mask-supplier	30	boolean	false	false
router-discovery	31	boolean	false	false
router-solicitation-address	32	ipv4-address	false	false
static-routes	33	ipv4-address	true	false
trailer-encapsulation	34	boolean	false	false
arp-cache-timeout	35	uint32	false	false
ieee802-3-encapsulation	36	boolean	false	false
default-tcp-ttl	37	uint8	false	false
tcp-keepalive-interval	38	uint32	false	false
tcp-keepalive-garbage	39	boolean	false	false
nis-domain	40	string	false	false
nis-servers	41	ipv4-address	true	false
ntp-servers	42	ipv4-address	true	false
vendor-encapsulated-options	43	empty	false	false
netbios-name-servers	44	ipv4-address	true	false
netbios-dd-server	45	ipv4-address	true	false
netbios-node-type	46	uint8	false	false

Continued on next page

Table 1 – continued from previous page

Name	Code	Type	Array	Always Return
netbios-scope	47	string	false	false
font-servers	48	ipv4-address	true	false
x-display-manager	49	ipv4-address	true	false
dhcp-option-overload	52	uint8	false	false
dhcp-message	56	string	false	false
dhcp-max-message-size	57	uint16	false	false
vendor-class-identifier	60	binary	false	false
nwip-domain-name	62	string	false	false
nwip-suboptions	63	binary	false	false
tftp-server-name	66	string	false	false
boot-file-name	67	string	false	false
user-class	77	binary	false	false
client-system	93	uint16	true	false
client-ndi	94	record (uint8, uint8, uint8)	false	false
uuid-guid	97	record (uint8, binary)	false	false
subnet-selection	118	ipv4-address	false	false
domain-search	119	binary	false	false
vivco-suboptions	124	binary	false	false
vivso-suboptions	125	binary	false	false

DHCP Option Types

binary An arbitrary string of bytes, specified as a set of hexadecimal digits.

boolean Boolean value with allowed values `true` or `false`.

empty No value, data is carried in suboptions.

fqdn Fully qualified domain name (e.g. `www.example.com`).

ipv4-address IPv4 address in dotted-decimal notation (e.g. `192.0.2.1`).

ipv6-address IPv6 address in compressed colon notation (e.g. `2001:db8::1`).

ipv6-prefix An IPv6 address and prefix length

record Structured data of other types (except `record` and `empty`).

string Any arbitrary text.

int8 8-bit signed integer with values between `-128` to `127`.

int16 16-bit signed integer with values between `-32768` to `32767`.

int32 32 bit signed integer with values between `-2147483648` and `2147483647`.

psid Port Set ID and length with values in the format `<id>/<length>`. Defines a port set for use by the client. The value of `<id>` is `0-65535`, value of `<length>` is `0-16`.

tuple A length encoded as an 8-bit unsigned integer followed by a string of this length.

uint8 8 bit unsigned integer with values between `0` and `255`.

uint16 16 bit unsigned integer with values between `0` and `65535`.

uint32 32 bit unsigned integer with values between `0` and `4294967295`.

IPv4 DHCP Option Definitions

TNSR also supports custom DHCP option definitions. These allow new options not listed in *Standard IPv4 DHCP Options* to be defined and they may also be used to change the standard definitions. Once created, these entries may be used as any other option, in the manner covered earlier in this document.

To create a new option definition, start in `config-kea-dhcp4` mode and use the `option-def <name>` command:

```
tnsr(config-kea-dhcp4) # option-def <name>
tnsr(config-kea-dhcp4-optdef) #
```

Note: To create a custom option, use a unique custom name; to redefine a standard option, enter its name from *Standard IPv4 DHCP Options*.

This command enters `config-kea-dhcp4-optdef` mode where the following additional commands are available:

- array <true|false>** When set to `true`, the data for this option is an array of values.
- code <code-val>** The code number for the DHCP option.
- encapsulate <encap>** Encapsulated option space name.
- record-types <types>** Record field type list.
- space <space-name>** Option space name.
- type <type>** Type of the option data, see *DHCP Option Types* for a list of types and allowed values.

17.1.2 Subnet Configuration

A `subnet` defines a network in which the DHCP server will provide addresses to clients, for example:

```
tnsr(config-kea-dhcp4) # subnet 10.2.0.0/24
tnsr(config-kea-subnet4) # interface GigabitEthernet0/14/2
```

From within the `subnet4` configuration mode, the following commands can be used:

- id <id>** Sets an optional unique identifier for this subnet.
- interface <name>** **Required.** The interface on which the subnet is located.
- option** Defines an option specific to this subnet (*DHCP Options*).
- pool** Defines a pool of addresses to serve inside this subnet. (*Address Pool Configuration*).
- reservation <ipv4-address>** Defines a host reservation to tie a client MAC address to a static IP address assignment.

At a minimum, the subnet itself must contain an `interface` definition and a `pool`.

17.1.3 Address Pool Configuration

A `pool` controls which addresses inside the `subnet` can be used by clients, for example:

```
tnsr(config-kea-subnet4) # pool 10.2.0.128-10.2.0.191
tnsr(config-kea-subnet4-pool) #
```

A pool may be defined as an address range (inclusive) as shown above in `<ipv4-addr>-<ipv4-addr>` format, or as a prefix, such as `10.2.0.128/26`.

Options can be defined inside a pool that only apply to clients receiving addresses from that pool.

17.1.4 Host Reservations

A reservation sets up a static IP address reservation for a client inside a subnet. For example:

```
tnsr(config-kea-subnet4)# reservation 10.2.0.20
tnsr(config-kea-subnet4-reservation)#
```

This reservation ensures that a client always obtains the same IP address, and can also provide the client with DHCP options that differ from the main subnet configuration.

Reservations are defined from within `config-kea-subnet4` mode, and take the form of `reservation <ipv4-address>`. That command then enters `config-kea-subnet4-reservation` mode, which contains the following options:

hostname <hostname> The hostname for this client.

mac-address <mac-address> **Mandatory.** The MAC address of the client, used to uniquely identify the client and assign this reserved IP address. The same MAC address cannot be used in more than one reservation on a single subnet.

option <dhcp4-option> DHCP options specific to this client. See *DHCP Options* for details on configuring DHCP options.

At a minimum, a reservation entry requires the `ipv4-address` which defines the reservation itself, and a `mac-address` to identify the client.

Warning: While it is possible to define a reservation inside a pool, this can lead to address conflicts in certain cases, such as when a different client already holds a lease for the new reservation.

The best practice is to keep reservations outside of the dynamic assignment pool.

Host reservation example:

```
tnsr(config-kea-subnet4)# reservation 10.2.0.20
tnsr(config-kea-subnet4-reservation)# mac-address 00:0c:29:4c:b3:9b
tnsr(config-kea-subnet4-reservation)# hostname mint-desktop
tnsr(config-kea-subnet4-reservation)# exit
tnsr(config-kea-subnet4)#
```

17.2 DHCP Logging

DHCP logging is configured in `config-kea-dhcp4-log` mode. To enter this mode, start in `config-kea-dhcp4` mode and issue the `logging <logger-name>` command.

The `<logger-name>` parameter must be one of these names:

kea-dhcp4 Default DHCPv4 logging behavior, used when no settings exist for a specific logger.

kea-dhcp4.alloc-engine DHCPv4 lease allocation events.

kea-dhcp4.bad-packets DHCPv4 packets that are dropped or rejected.

kea-dhcp4.callouts DHCPv4 hook point callout registration/execution.

kea-dhcp4.commands DHCPv4 commands received over the command channel.

kea-dhcp4.ddns DHCPv4 client FQDN and Hostname option processing.

kea-dhcp4.dhcp4 DHCPv4 server basic operations.

kea-dhcp4.dhcpsrv DHCPv4 libkea-dhcpsrv library default logging.

kea-dhcp4.eval DHCPv4 client classification expression evaluation.

kea-dhcp4.hooks DHCPv4 hook registration/deregistration.

kea-dhcp4.hosts DHCPv4 host reservation management.

kea-dhcp4.leases DHCPv4 lease allocation.

kea-dhcp4.options DHCPv4 option processing/parsing/encoding.

kea-dhcp4.packets DHCPv4 packet reception/transmission.

kea-dhcp4.stat-cmds-hooks DHCPv4 libdhcp_stat_cmds library default logging.

For example:

```
tnsr(config-kea-dhcp4) # logging kea-dhcp4
tnsr(config-kea-dhcp4-log) #
```

See also:

See the [Kea documentation for Logging](#), for a list of values and their meanings.

17.2.1 DHCP Logging Options

`config-kea-dhcp4-log` mode contains the following commands:

debug-level <level> The amount of debug information to log, when a log message is classified using the debug severity level. From 0 (lowest detail) to 99 (most detail).

output <location> Sets the log output location and enters `config-kea-dhcp4-log-out` mode to further configure behavior for the given location. The `<location>` must be one of:

- stdout** Log messages to standard output.
- stderr** Log messages to standard error.
- syslog** Log messages to syslog using the daemon name.
- syslog:<name>** Log messages to syslog using the given `<name>`.
- <filename>** The full path and filename in which to log messages.

severity (debug|error|fatal|info|warn) The severity level of messages to write in the log:

- fatal** Errors severe enough to cause the daemon to exit.
- error** Errors which are notable but otherwise do not prevent the daemon from functioning.
- warn** Unusual conditions which are not normal but also not causing problems.
- info** General noteworthy events and other similar messages.
- debug** Debugging messages, which are highly informative but also verbose. The amount of debugging information can be further controlled by the `debug-level` command.

17.2.2 DHCP Log Output

The `output <location>` command from `config-kea-dhcp4-log` mode enters `config-kea-dhcp4-log-out` mode, which contains the following commands:

flush (false|true) Flush the log buffer after writing each line. Lowers performance but ensures that every log message is written completely. This can help in cases where the daemon dies before the log buffer is emptied.

maxsize <size> When logging to a file, this option controls the maximum size of the log file before it is rotated. Default value is 10240000 (10MB).

maxver <rotate> When logging to a file and rotation is enabled, this command controls how many previous log files are retained when performing log rotation. The default value is 1, which is also the minimum allowed value.

If the value of `maxsize` is greater than 204800, the daemon will perform log rotation when the size of the log file exceeds `maxsize`. When rotating logs, the current log will be renamed to `<filename>.1`. If `<filename>.1` already exists, it will be renamed to `<filename>.2`, and so on until the number of log files reaches the value of `maxver`. Older log files beyond the limit set by `maxver` are removed.

17.3 DHCP Service Control and Status

17.3.1 Enable the DHCP Service

Enable the DHCP4 server:

```
tnsr(config)# dhcp4 enable
tnsr(config)#
```

17.3.2 Disable the DHCP Service

Similar to the DHCP `enable` command, disable the DHCP4 service from configuration mode:

```
tnsr(config)# dhcp4 disable
tnsr(config)#
```

17.3.3 Check the DHCP Service Status

Check the status of the DHCP services from configuration mode:

```
tnsr(config)# service dhcp status
DHCPv4 server: active
DHCPv6 server: inactive
DHCP DDNS: inactive
Control Agent: inactive
Kea DHCPv4 configuration file: /etc/kea/kea-dhcp4.conf
Kea DHCPv6 configuration file: /etc/kea/kea-dhcp6.conf
Kea DHCP DDNS configuration file: /etc/kea/kea-dhcp-ddns.conf
Kea Control Agent configuration file: /etc/kea/kea-ctrl-agent.conf
keactrl configuration file: /etc/kea/keactrl.conf
```

17.3.4 View the DHCP Configuration

View the current Kea DHCP Daemon and Control TNSR Configuration:

```
tnsr# show kea
```

View the current Kea DHCP Daemon TNSR Configuration:

```
tnsr# show kea dhcp4
```

View the current Kea DHCP daemon configuration file:

```
tnsr# show kea dhcp4 config-file
```

View the current Kea Control TNSR Configuration:

```
tnsr# show kea keactrl
```

View the current Kea Control Configuration file:

```
tnsr# show kea keactrl config-file
```

17.3.5 View the DHCP Lease Database

View the database of active DHCP leases:

```
tnsr# show kea dhcp4 leases
```

IP address	HW address	Hostname	CLTT	Valid	Subnet	ID	State
10.2.0.129	00:0c:29:4c:b3:9b	mintclient	1581082364	7200		1	default
10.2.0.130	00:0c:29:41:dc:ac	graybeard	1581081332	7200		1	default
10.2.0.131	00:0c:29:13:54:93	doctor	1581081363	7200		1	default

17.4 DHCP Service Example

Configure the DHCP IPv4 Service from configuration mode (*Configuration Mode*). This example uses the interface and subnet from *Example Configuration*:

```
tnsr(config)# dhcp4 server
tnsr(config-kea-dhcp4)# description LAN DHCP Server
tnsr(config-kea-dhcp4)# interface listen GigabitEthernet0/14/2
tnsr(config-kea-dhcp4)# option domain-name
tnsr(config-kea-dhcp4-opt)# data example.com
tnsr(config-kea-dhcp4-opt)# exit
tnsr(config-kea-dhcp4)# subnet 10.2.0.0/24
tnsr(config-kea-subnet4)# pool 10.2.0.128-10.2.0.191
tnsr(config-kea-subnet4-pool)# exit
tnsr(config-kea-subnet4)# interface GigabitEthernet0/14/2
tnsr(config-kea-subnet4)# option domain-name-servers
tnsr(config-kea-subnet4-opt)# data 8.8.8.8, 8.8.4.4
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-subnet4)# option routers
tnsr(config-kea-subnet4-opt)# data 10.2.0.1
```

(continues on next page)

(continued from previous page)

```
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-subnet4)# exit
tnsr(config-kea-dhcp4)# exit
tnsr(config)# dhcp4 enable
tnsr(config)#
```

The above example configures `example.com` as the domain name supplied to all clients. For the specific subnet in the example, the TNSR IP address inside the subnet is supplied by DHCP as the default gateway for clients, and DHCP will instruct clients to use `8.8.8.8` and `8.8.4.4` for DNS servers.

Note: The subnet definition requires an interface.

DNS RESOLVER

TNSR uses the [Unbound](#) Domain Name System Resolver to handle DNS resolution and client queries.

Unbound is a recursive caching DNS resolver. Unbound can validate DNS data integrity with DNSSEC, and supports query privacy using DNS over TLS.

By default Unbound will act as a DNS resolver, directly contacting root DNS servers and other authoritative DNS servers in search of answers to queries. Unbound can also act as a DNS Forwarder, sending all DNS queries to specific upstream servers.

18.1 DNS Resolver Configuration

Unbound can be configured with a wide array of optional parameters to fine-tune its behavior. Due to the large number of options, this documentation is split into several parts, with related options listed together.

These options are all found in `config-unbound` mode, which is entered by the command `unbound server` from configuration mode ([Configuration Mode](#)).

enable/disable These commands enable or disable options that do not require additional parameters, they can only be turned on or off. The specific options are discussed in other areas of this chapter such as [Security Tuning](#) and [Cache & Performance Tuning](#).

verbosity <n> Sets the verbosity of the logs, from 0 (no logs) through 5 (high). Default value is 1. Each level provides the information from the lower levels plus additional data.

- Level 1: Operational Information
- Level 2: Additional details
- Level 3: Per-query logs with query level information
- Level 4: Algorithm level information
- Level 5: Client identification for cache misses

interface <x.x.x.x> [port <n>] Configures an interface IP address that Unbound will use for binding as a server, and an optional port specification. In most cases there should be an interface definition for a TNSR IP address in each local network, plus a definition for localhost (127.0.0.1 as shown in [Resolver Mode Example](#)). The port number defaults to 53 and should not be changed in most use cases.

outgoing-interface <ip-address> Configures an interface IP address that Unbound will use when making outbound DNS queries to upstream servers (roots or forwarders).

Note: If this is not configured, Unbound will make queries using the host OS default route, and not TNSR interfaces or routes.

port <n> Sets the default port which Unbound will use to listen for client queries. Defaults to 53.

enable/disable ip4 Tells Unbound to use, or not use, IPv4 for answering or performing queries. Default is enabled. Unless TNSR has no IPv4 connectivity, this should be left enabled.

enable/disable ip6 Tells Unbound to use, or not use, IPv6 for answering or performing queries. Default is enabled. Unless there is a situation where TNSR is configured with IPv6 addresses but lacks working connectivity to upstream networks via IPv6, this should remain enabled.

enable/disable udp Tells Unbound to use, or not use, UDP for answering or performing queries. Default is enabled. In nearly all cases, DNS requires UDP to function, except special cases such as a pure DNS over TLS environment. Thus, this should nearly always be left enabled.

enable/disable tcp Tells Unbound to use, or not use, TCP for answering or performing queries. Default is enabled. TCP is generally required for functional DNS, especially for queries with large answers. DNS over TLS also requires TCP. Unless a use case specifically calls for UDP DNS only, this should remain enabled.

access-control Configures access control list entries for Unbound. See [Access Control Lists](#).

forward-zone Enters `config-unbound-fwd-zone` mode. See [Forward Zones](#).

18.1.1 Access Control Lists

Access Control Lists in Unbound determine which clients can and cannot perform queries against the DNS Resolver as well as aspects of client behavior.

The default behavior is to allow access from TNSR itself (localhost), but refuse queries from other clients.

Example:

```
tnsr(config)# unbound server
tnsr(config-unbound)# access-control 10.2.0.0/24 allow
```

The general form of the command is:

```
tnsr(config-unbound)# access-control <IPv4 or IPv6 Network Prefix> <action>
```

The **IPv4 or IPv6 Network Prefix** is a network specification, such as `10.2.0.0/24` or `2001:db8::/64`. For a single address, use `/32` for IPv4 or `/128` for IPv6.

The **Action** types are:

allow Allow access to recursive and local data queries for clients in the specified network.

allow_snoop Allow access to recursive and local data queries for clients in the specified network, additionally this allows access to cache snooping. Cache snooping is a technique to use nonrecursive queries to examine the contents of the cache for debugging or identifying malicious data.

refuse Stops queries from clients in the specified network, but sends a DNS response code REFUSED error. This is the default behavior for networks other than localhost, since it is friendly and protocol-safe response behavior.

refuse_non_local Similar to `refuse` but allows queries for authoritative local data. Recursive queries are refused.

deny Drops and does not respond to queries from clients in the specified network. In most cases a `refuse` action is preferable since DNS is not designed to handle a non-response. A lack of response may cause clients to send additional unwanted queries.

deny_non_local Allows queries for authoritative local-data only, all other queries are dropped without a response.

18.1.2 Forward Zones

In Unbound, a Forward Zone controls how queries are handled on a per-zone basis. This can be used to send queries for a specific domain or zone to a specific DNS server, or it can be used to setup forwarding mode sending all queries to one or more upstream recursive DNS servers.

Forward Zone Examples

Example to override the default resolver behavior and forward all queries to an upstream DNS server:

```
tnsr(config)# unbound server
tnsr(config-unbound)# forward-zone .
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.8.8
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.4.4
```

This forwards the root zone (`.`) and all zones underneath to the specified servers, in this case, `8.8.8.8` and `8.8.4.4`.

Example to send queries for one specific domain to an alternate server:

```
tnsr(config)# unbound server
tnsr(config-unbound)# forward-zone example.com
tnsr(config-unbound-fwd-zone)# nameserver address 192.0.2.5
```

This example sends all queries for `example.com` and subdomains underneath `example.com` to the server at `192.0.2.5`. This is useful for sending queries for internal domains to a local authoritative DNS server, or an internal DNS server reachable through a VPN.

Forward Zone Configuration

To enter `config-unbound-fwd-zone` mode, start from `config-unbound` mode and use the `forward-zone <zone-name>` command. The `<zone-name>` takes the form of the domain part of a fully qualified domain name (FQDN), but may also be `.` to denote the root zone.

nameserver address <ip-address> [port <port>] [auth-name <name>] Specifies a DNS server for this zone by IP address. Optionally, a port number may be given (default 53). `auth-name` sets the FQDN of the DNS server for use in validating certificates with DNS over TLS.

nameserver host <host-name> Specifies a DNS server for this zone by FQDN. This hostname will be resolved before use.

enable/disable forward-first When enabled, if a query fails to the forwarding DNS servers it will be retried using resolver mode through the root DNS servers. By default this behavior is disabled.

enable/disable forward-tls-upstream When enabled, queries to the DNS servers in this zone are sent using DNS over TLS, typically on port 853. This mode provides query privacy by encrypting communication between Unbound and upstream DNS servers in the zone. Default is disabled as this feature is not yet widely supported by other platforms.

Multiple DNS server address or host entries may be given for a forward zone. These servers are not queried sequentially and are not necessarily queried simultaneously. Unbound tracks the availability and performance of each DNS server in the zone and will attempt to use the most optimal server for a query.

18.1.3 Local Zones

Unbound can host local zone data to complement, control, or replace upstream DNS data. This feature is commonly used to supply local clients with host record responses that do not exist in upstream DNS servers, or to supply local clients with a different response, akin to a DNS view.

Local Zone Example

This basic example configures a local zone for `example.com` and two hostnames inside. If a client queries TNSR for these host records, it will respond with the answers configured in the local zone. If a client requests records for a host under `example.com` not listed in this local zone, then the query is resolved as usual through the usual resolver or forwarding server mechanisms.

```
tnsr(config)# unbound server
tnsr(config-unbound)# local-zone example.com
tnsr(config-unbound-local-zone)# type transparent
tnsr(config-unbound-local-zone)# hostname server.example.com
tnsr(config-unbound-local-host)# address 192.0.2.5
tnsr(config-unbound-local-host)# exit
tnsr(config-unbound-local-zone)# hostname db.example.com
tnsr(config-unbound-local-host)# address 192.0.2.6
tnsr(config-unbound-local-host)# exit
```

Local Zone Configuration

Local zones are configured in `config-unbound` mode (*DNS Resolver Configuration*) using the `local-zone <zone-name>` command. This defines a new local zone and enters `config-unbound-local-zone` mode.

Within `config-unbound-local-zone` mode, the following commands are available:

description <descr> A short text description of the zone

type <type> The type for this local zone, which can be one of:

transparent Gives local data, and resolves normally for other names. If the query matches a defined host but not the record type, the client is sent a NOERROR, NODATA response. This is the most common type and most likely the best choice for most scenarios.

typetransparent Similar to transparent, but will forward requests for records that match by name but not by type.

deny Serve local data, drop queries otherwise.

inform Like transparent, but logs the client IP address.

inform_deny Drops queries and logs the client IP address.

no_default Normally resolve AS112 zones.

redirect Serves zone data for any subdomain in the zone.

refuse Serve local data, else reply with REFUSED error.

static Serve local data, else NXDOMAIN or NODATA answer.

hostname <fqdn> Defines a new hostname within the zone, and enters `config-unbound-local-host` mode. A local zone may contain multiple hostname entries.

Note: Include the domain name when creating a hostname entry.

Inside `config-unbound-local-host` mode, the following commands are available:

description <descr> A short text description of this host

address <ip-address> The IPv4 or IPv6 address to associate with this hostname for forward and reverse (PTR) lookups.

18.1.4 Security Tuning

Unbound can be tuned to provide stronger (or weaker) security and privacy, depending on the needs of the network and features supported by clients and upstream servers.

enable caps-for-id Experimental support for draft [dns-0x20](#). This feature combats potentially spoofed replies by randomly flipping the 0x20 bit of ASCII letters, which switches characters between upper and lower case. The answer is checked to ensure the case in the response matches the request exactly. This is disabled by default since it is experimental, but is safe to enable unless the upstream server does not copy the query question to the response identically. Most if not all servers follow this convention, but it is unknown if this behavior is truly universal.

enable harden dnssec-stripped Require DNSSEC for trust-anchored zones. If the DNSSEC data is absent, the zone is marked as bogus. If disabled and no DNSSEC data is received in the response, the zone is marked insecure. Default behavior is enabled. If disabled, there is a risk of a forced downgrade attack on the response that disables security on the zone.

enable harden glue Trust glue only if the server is authorized. Default is enabled.

enable hide identity When enabled, queries are refused for `id.server` and `hostname.bind`, which prevents clients from obtaining the server identity. Default behavior is disabled.

enable hide version When enabled, queries are refused for `version.server` and `version.bind`, preventing clients from determining the version of Unbound. Default behavior is disabled.

thread unwanted-reply-threshold <threshold> When set, Unbound tracks the total number of unwanted replies in each thread. If the threshold is reached, Unbound will take defensive action and logs a warning. This helps prevent cache poisoning by clearing the RRSet and message caches when triggered. By default this behavior is disabled. If this behavior is desired, a starting value of 10000000 (10 million) is best. Change the value in steps of 5-10 million as needed.

jostle timeout <t> Timeout in milliseconds, used when the server is very busy. This timeout should be approximately the same as the time it takes for a query to reach an upstream server and receive a response (round trip time). If a large number of queries are received by Unbound, than half the active queries are allowed to complete and the other half are replaced by new queries. This helps reduce the effectiveness of a denial of service attack by allowing the server to ignore slow queries when under load. The default value is 200 msec.

18.1.5 Cache & Performance Tuning

port outgoing range <n> Sets the number of source ports Unbound may use per thread to connect when making outbound queries to upstream servers. A larger number of ports provides protection against

spoofing. Default value varies by platform. A large number of ports yields better performance but it also consumes more host resources.

edns reassembly size <s> Number to advertise as the EDNS reassembly buffer size, in bytes. This value is sent in queries and must not be set larger than the default message buffer size, 65552. The default value is 4096, which is recommended by RFC. May be set lower to alleviate problems with fragmentation resulting in timeouts. If the default value is too large, try 1472, or 512 in extreme cases. Avoid setting that low as it will cause many queries to fall back to TCP which can negatively impact performance.

host cache num-hosts <num> Number of hosts to hold in the cache, defaults to 10000. Larger caches can result in increased performance but consume more host resources.

host cache slabs <s> Number of slabs in the host cache. Larger numbers help prevent lock contention by threads when performing cache operations. The value is a power of 2, between 0 . . 10

host cache ttl <t> The amount of time, in seconds, that entries in the host cache are kept. Default value is 900 seconds.

enable key prefetch When enabled, Unbound will start fetching DNSKEYS when it sees a DS record instead of waiting until later in the process. Prefetching keys will consume more CPU, but reduces latency. The default is disabled.

key cache slabs <s> Number of slabs in the key cache. Larger numbers help prevent lock contention by threads when performing key cache operations. The value is a power of 2, between 0 . . 10. Setting to a number close to the number of CPUs/cores in the host is best.

enable message prefetch Prefetch message cache items before they expire to keep entries in the cache updated. When enabled, Unbound will consume approximately 10% more throughput and CPU time but it will keep popular items primed in the cache for better client performance. Disabled by default.

message cache size <s> Size of the message cache, in bytes. The message cache stores DNS meta-information such as message formats. Default value is 4 MB.

message cache slabs <s> Number of slabs in the message cache. Larger numbers help prevent lock contention by threads when performing message cache operations. The value is a power of 2, between 0 . . 10. Setting to a number close to the number of CPUs/cores in the host is best.

rrset cache size <s> Size of the RRset cache, in bytes. The RRset cache stores resource records. Default value is 4 MB.

rrset cache slabs <s> Number of slabs in the RRset cache. Larger numbers help prevent lock contention by threads when performing RRset cache operations. The value is a power of 2, between 0 . . 10. Setting to a number close to the number of CPUs/cores in the host is best.

rrset-message cache ttl maximum <max> Maximum time that values in the RRset and message caches are kept in the cache, specified in seconds. The default value is 86400 (1 day). When set lower, Unbound will be forced to query for data more often, but it will also ignore very large TTLs in DNS responses.

rrset-message cache ttl minimum <max> Minimum time that values in the RRset and message caches are kept in the cache, specified in seconds. The default value is 0, which honors the TTL specified in the DNS response. Higher values may ignore the TTL set by the response, which means a record may be out of sync with the source, but it also prevents queries from being repeated frequently when a very low TTL is set by the domain.

socket receive-buffer size <s> SO_RCVBUF socket receive buffer size for incoming queries on the listening port(s). Larger values result in less drops during spikes in activity. The default is 0 which uses the system default value. Cannot be set higher than the maximum value for the operating system, such as the one shown in the `net.core.rmem_max` sysctl OID.

tcp buffers incoming <n> Number of incoming TCP buffers that Unbound will allocate per thread. Larger values can handle higher loads, but will consume more resources. The default value is 10. A value of 0 will disable acceptance of TCP queries.

tcp buffers outgoing <n> Number of outgoing TCP buffers that Unbound will allocate per thread. Larger values can handle higher loads, but will consume more resources. The default value is 10. A value of 0 will disable TCP queries to authoritative DNS servers.

thread num-queries <n> Number of queries serviced by each thread simultaneously. If more queries arrive and there is no room to answer them, the new queries will be dropped, unless older/slower queries can be dropped by using the `jostle timeout`. Default varies by platform but is typically 512 or 1024.

thread num-threads <n> Number of threads created by Unbound for serving clients. Defaults to one thread per CPU/core. To disable threading, set to 1.

enable serve-expired When enabled, Unbound will immediately serve answers to clients using expired cache entries if they exist. Unbound still performs the query and will update the cache with the result. This can result in faster, but potentially incorrect, answers for client queries. Default is disabled.

18.2 DNS Resolver Service Control and Status

18.2.1 Enable the DNS Resolver

Enable the DNS Resolver:

```
tnsr(config)# unbound enable
tnsr(config)#
```

18.2.2 Disable the DNS Resolver

Similar to the `enable` command, disable the DNS Resolver from configuration mode:

```
tnsr(config)# unbound disable
tnsr(config)#
```

18.2.3 Check the DNS Resolver Status

Check the status of the DNS Resolver from configuration mode:

```
tnsr(config)# service unbound status
* unbound.service - Unbound recursive Domain Name Server
  Loaded: loaded (/usr/lib/systemd/system/unbound.service; disabled; vendor preset:
  ↳ disabled)
  Active: active (running) since Wed 2018-08-22 15:26:05 EDT; 55min ago
  Process: 26675 ExecStartPre=/usr/sbin/unbound-anchor -a /var/lib/unbound/root.key -
  ↳ c /etc/unbound/icannbundle.pem (code=exited, status=0/SUCCESS)
  Process: 26673 ExecStartPre=/usr/sbin/unbound-checkconf (code=exited, status=0/
  ↳ SUCCESS)
  Main PID: 26679 (unbound)
  CGroup: /system.slice/unbound.service
```

(continues on next page)

(continued from previous page)

```
└─26679 /usr/sbin/unbound -d

Aug 22 15:26:05 tnsr.example.com systemd[1]: Starting Unbound recursive Domain Name_
↳Server...
Aug 22 15:26:05 tnsr.example.com unbound-checkconf[26673]: unbound-checkconf: no_
↳errors in /etc/unbound/unbound.conf
Aug 22 15:26:05 tnsr.example.com systemd[1]: Started Unbound recursive Domain Name_
↳Server.
Aug 22 15:26:05 tnsr.example.com unbound[26679]: [26679:0] notice: init module 0:_
↳subnet
Aug 22 15:26:05 tnsr.example.com unbound[26679]: [26679:0] notice: init module 1:_
↳validator
Aug 22 15:26:05 tnsr.example.com unbound[26679]: [26679:0] notice: init module 2:_
↳iterator
Aug 22 15:26:05 tnsr.example.com unbound[26679]: [26679:0] info: start of service_
↳(unbound 1.6.6).
```

18.2.4 View the DNS Resolver Configuration

View the current Unbound DNS Resolver daemon configuration file:

```
tnsr# show unbound config-file
```

18.3 DNS Resolver Examples

Configure the DNS Resolver Service from configuration mode (*Configuration Mode*). These examples use the interface and subnet from *Example Configuration*.

18.3.1 Resolver Mode Example

For Resolver mode, the configuration requires only a few basic options:

```
tnsr# configure
tnsr(config)# unbound server
tnsr(config-unbound)# interface 127.0.0.1
tnsr(config-unbound)# interface 10.2.0.1
tnsr(config-unbound)# outgoing-interface 203.0.113.2
tnsr(config-unbound)# access-control 10.2.0.0/24 allow
tnsr(config-unbound)# exit
tnsr(config)# unbound enable
```

This example enables the Unbound DNS Resolver and configures it to listen on localhost as well as 10.2.0.1 (GigabitEthernet0/14/2, labeled LAN in the example). It uses 203.0.113.2, which is the example WAN interface address, for outgoing queries. The example also allows clients inside the local subnet, 10.2.0.0/24, to perform DNS queries and receive responses.

18.3.2 Forwarding Mode Example

For Forwarding mode, use the configuration above plus these additional commands:

```
tnsr# configure
tnsr(config)# unbound server
tnsr(config-unbound)# forward-zone .
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.8.8
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.4.4
tnsr(config-unbound-fwd-zone)# exit
tnsr(config-unbound)# exit
```

This example builds on the previous example but instead of working in resolver mode, it will send all DNS queries to the upstream DNS servers 8.8.8.8 and 8.8.4.4.

NETWORK TIME PROTOCOL

The Network Time Protocol (NTP) service on TNSR synchronizes the host clock with reference sources, typically remote servers. It also acts as an NTP server for clients.

19.1 NTP Configuration

The NTP daemon has a variety of options to fine-tune its timekeeping behavior.

interface sequence <seq> <action> <address> Interface binding options. The default behavior when no `interface` configuration entries are present is to bind to all available addresses on the host.

seq The sequence number controls the order of the interface definitions in the NTP daemon configuration.

action The action taken for NTP traffic on this interface, it can be one of:

drop Bind the daemon to this interface, but drop NTP traffic.

ignore Do not bind the daemon to this interface.

listen Bind the daemon to this interface and use it for NTP traffic.

address The address or interface to bind. This may be:

prefix An IPv4/IPv6 prefix, which will bind to only that specific address.

interface An interface name, which will bind to every address on that interface.

all Bind to all interfaces and addresses on TNSR.

server <address|host> <server> Defines an NTP peer with which the daemon will attempt to synchronize the clock. This command enters `config-ntp-server` mode. The server may be specified as:

address <IPv4/IPv6 Address> An IPv4 or IPv6 address specifying a single NTP server.

host <fqdn> A fully qualified domain name, which will be resolved using DNS.

Within `config-ntp-server` mode, additional commands are available that control how NTP interacts with the specified server:

iburst Use 8 packets on unreachable servers, which results in faster synchronization at startup and when a peer is recovering.

maxpoll <poll> Maximum polling interval for NTP messages. This is specified as a power of 2, in seconds. May be between 7 and 17, defaults to 10 (1024 seconds).

noselect Instructs NTP to not use the server for synchronization, but it will still connect and display statistics from the server.

prefer When set, NTP will prefer this server if it and multiple other servers are all viable candidates of equal quality.

operational-mode server This entry is a single server. When the server is specified as an FQDN, if the DNS response contains multiple entries then only one is selected. Can also be used with IPv4/IPv6 addresses directly, rather than FQDN entries.

operational-mode pool This entry is a pool of servers. Only compatible with FQDN hosts. NTP will expect multiple records in the DNS response and will use all of these entries as distinct servers. This is a reliable way to configure multiple NTP peers with minimal configuration.

Warning: An `operational-mode` is required.

tinker panic <n> Sets the NTP panic threshold, in seconds. This is a sanity check which will cause NTP to fail if the difference between the local and remote clocks is too great. Commonly set to 0 to disable this check so that NTP will still synchronize when its clock is off by a large factor. The default value is 1000.

tos orphan <n> Configures the stratum of orphan mode servers from 1 to 16. When all UTC reference peers below this stratum are unreachable, clients in the same subnet may use each other as references as a last resort.

driftfile <file> Full path to the filename used by the NTP daemon to store clock drift information to improve accuracy over time. This file and its directory must be writable by the `ntp` user or group.

statsdir <file> Full path to statistics directory used by the NTP daemon. This directory must be writable by the `ntp` user or group.

<enable|disable> monitor Enables or disables the monitoring facility used to poll the NTP daemon for information about peers and other statistics. This is enabled by default, and is also enabled if `limited` is present in any `restrict` entries. This is required for `show ntp <x>` commands which display peer information to function.

19.1.1 NTP Restrictions

NTP restrictions control how NTP treats traffic from peers. The *NTP Service Example* at the start of this section contains a good set of restrictions to use as a starting point.

These restrictions are configured using the `restrict` command from within `config-ntp` mode.

restrict <default|source|host|prefix> This command enters `config-ntp-restrict` mode.

The restriction is placed upon an address specified as:

default The default restriction for any host.

source Default restrictions for associated hosts.

host An address specified as an FQDN to be resolved using DNS.

prefix An IPv4 or IPv6 network specification.

In `config-ntp-restrict` mode, the following settings control what hosts matching this restriction can do:

kod Sends a Kiss of Death packet to misbehaving clients. Only works when paired with the `limited` option.

- limited** Enforce rate limits on clients. This does not apply to queries from `ntpq/ntpd` or the `show ntp <x>` commands.
- nomodify** Allows clients to query read only server state information, but does not allow them to make changes.
- nopeer** Deny unauthorized associations. When using a server entry in `pool` mode, this should be present in the `default` restriction but not in the `source` restriction.
- noquery** Deny `ntpq/ntpd/show ntp <x>` queries for NTP daemon information. Does not affect NTP acting as a time server.
- noserve** Disables time service. Still allows `ntpq/ntpd/show ntp <x>` queries
- notrap** Decline mode 6 trap service to clients.

19.1.2 NTP Logging

The NTP Logging configuration controls which type of events are logged by the NTP daemon using `syslog`, and the verbosity of the logs. By default, the NTP daemon will log all synchronization messages.

The logging configuration is set using the `logconfig` command from within `config-ntp` mode.

logconfig sequence <seq> <action> <class> <type>

- seq** Specifies the sequence for log entries so that the order of parameters may be controlled by the configuration.
- action** Specifies the action for this log entry, as one of:
- set** Set the mask for the log entry. Typically this would be used for the first entry to control which message class+type is logged as the base set of log entries.
 - add** Add log entries matching this specification to the specified total set of logs.
 - delete** Do not log entries matching this specification in the total set of logs.
- class** Specifies the message class, which can be one of:
- all** All message classes
 - clock** Messages about local clock events and information.
 - peer** Messages about peers.
 - sync** Messages about the synchronization state.
 - sys** Messages about system events and status.
- type** Specifies the type of messages to log for each class:
- all** All types of messages.
 - events** Event messages.
 - info** Informational messages.
 - statistics** Statistical information.
 - status** Status changes.

19.2 NTP Service Control and Status

19.2.1 Enable the NTP Service

Enable the NTP server:

```
tnsr(config)# ntp enable
tnsr(config)#
```

19.2.2 Disable the NTP Service

Similar to the NTP enable command, disable the NTP service from configuration mode:

```
tnsr(config)# ntp disable
tnsr(config)#
```

19.2.3 Check the NTP Service Status

Check the status of the NTP services from configuration mode:

```
tnsr(config)# service ntp status
* ntpd.service - Network Time Service
  Loaded: loaded (/usr/lib/systemd/system/ntpd.service; disabled; vendor preset:
  ↳ disabled)
  Active: active (running) since Thu 2018-11-15 07:05:57 EST; 2 weeks 5 days ago
  Main PID: 1744 (ntpd)
  CGroup: /system.slice/ntpd.service
          └─1744 /usr/sbin/ntpd -u ntp:ntp -g

Dec 04 11:38:44 ntpd[1744]: Listen normally on 21 mytap 10.2.99.1 UDP 123
Dec 04 11:38:44 ntpd[1744]: Listen normally on 22 vpp5 fe80::208:a2ff:fe09:95b5 UDP
  ↳ 123
Dec 04 11:38:44 ntpd[1744]: Listen normally on 23 vpp1 fe80::208:a2ff:fe09:95b1 UDP
  ↳ 123
Dec 04 11:38:44 ntpd[1744]: Listen normally on 24 vpp1 fe80::5 UDP 123
Dec 04 11:38:44 ntpd[1744]: Listen normally on 25 vpp5 fe80::15 UDP 123
Dec 04 11:38:44 ntpd[1744]: Listen normally on 26 mytap fe80::c41e:7bff:fea5:462a UDP
  ↳ 123
Dec 04 11:38:44 ntpd[1744]: new interface(s) found: waking up resolver
```

19.2.4 View NTP Peers

The NTP peer list shows the peers known to the NTP daemon, along with information about their network availability and quality. For more information on peer associations, see [View NTP Associations](#).

```
tnsr(config)# show ntp peers
```

Id	Host	Ref ID	Stratum	Reach	Poll	Delay	Offset	Jitter
17417	5.9.80.113	192.53.103.103	2	0xff	512	134.456	-1.936	3.904
17418	95.216.39.155	131.188.3.223	2	0xff	512	151.370	-1.582	4.883
17419	145.239.118.233	85.199.214.98	2	0xec	512	126.181	4.112	21.541
17420	178.128.4.44	204.123.2.5	2	0xff	512	80.998	2.906	4.140

19.2.5 View NTP Associations

The NTP peer associations list shows how the NTP daemon is using each peer, along with its status. These peers are listed by ID. For more information on each peer, see [View NTP Peers](#).

```
tnsr(config)# show ntp associations
Id      Status Persistent Auth En Authentic Reachable Broadcast Selection Event
↪Count
-----
↪--
17417 0x931a true      false false true      false outlier sys_peer 1
17418 0x941a true      false false true      false candidate sys_peer 1
17419 0x941a true      false false true      false candidate sys_peer 1
17420 0x961a true      false false true      false sys.peer sys_peer 1
```

19.2.6 View NTP Daemon Configuration File

View the current NTP Daemon configuration file, generated by the settings in TNSR:

```
tnsr# show ntp config-file
#
# NTP config autogenerated
#

tinker panic 0

tos orphan 12

logconfig =syncall +clockall

restrict ::/0 kod limited nomodify nopeer notrap
restrict default kod limited nomodify nopeer notrap
restrict source kod limited nomodify notrap

pool pool.ntp.org maxpoll 9
```

19.3 NTP Service Example

Configure the NTP Service from configuration mode ([Configuration Mode](#)). This example uses `pool.ntp.org` in pool mode so that multiple DNS results are used as reference servers.

```
tnsr(config)# ntp server
tnsr(config-ntp)# tos orphan 12
tnsr(config-ntp)# tinker panic 0
tnsr(config-ntp)# logconfig sequence 1 set sync all
tnsr(config-ntp)# logconfig sequence 2 add clock all
tnsr(config-ntp)# restrict default
tnsr(config-ntp-restrict)# kod
tnsr(config-ntp-restrict)# limited
tnsr(config-ntp-restrict)# nomodify
tnsr(config-ntp-restrict)# nopeer
tnsr(config-ntp-restrict)# notrap
tnsr(config-ntp-restrict)# exit
```

(continues on next page)

(continued from previous page)

```
tnsr(config-ntp)# restrict source
tnsr(config-ntp-restrict)# kod
tnsr(config-ntp-restrict)# limited
tnsr(config-ntp-restrict)# nomodify
tnsr(config-ntp-restrict)# notrap
tnsr(config-ntp-restrict)# exit
tnsr(config-ntp)# server host pool.ntp.org
tnsr(config-ntp-server)# operational-mode pool
tnsr(config-ntp-server)# maxpoll 9
tnsr(config-ntp-server)# exit
tnsr(config-ntp)# exit
tnsr(config)# ntp enable
tnsr(config)#
```

19.4 NTP Best Practices

Use a minimum of three servers, either as three separate server entries or a pool containing three or more servers. This is to ensure that if the clock on any one server becomes skewed, the remaining two sources can be used to determine that the skewed server is no longer viable. Otherwise NTP would have to guess which one is accurate and which is skewed.

There are a large number of public NTP servers available under `pool.ntp.org`. The `pool.ntp.org` DNS entry will return a number of randomized servers in each DNS query response. These can be used individually or as pools. The easiest way is to use the `pool` operational mode, which uses all returned servers as if they were specified individually.

When using entries as individual `server` entries, these responses can be subdivided into mutually exclusive pools of peers to avoid overlap. For example, if a configuration specifies `pool.ntp.org` multiple times for `server` entries, the same IP address could accidentally be selected twice. In this case, use `0.pool.ntp.org`, `1.pool.ntp.org`, `2.pool.ntp.org`, and so on. When queried in this way, the responses will be unique for each number.

Furthermore, there are also pools available for regional and other divisions. For example, to only receive responses for servers in the United States, use `us.pool.ntp.org` as a pool or `<n>.us.pool.ntp.org` as servers. For more information, see <https://www.ntppool.org/en/>

LINK LAYER DISCOVERY PROTOCOL

The Link Layer Discovery Protocol (LLDP) service provides a method for discovering which routers are connected to a LAN segment, and offers a way to discover the topology of a network.

20.1 Configuring the LLDP Service

LLDP is configured in two places: Global router parameters and per-interface parameters.

To enable LLDP, TNSR requires global settings for LLDP and at least one interface must participate in LLDP.

20.1.1 LLDP Router Configuration

Three LLDP commands are available in configuration mode (*Configuration Mode*) to configure global LLDP parameters for this router. All of these commands are required to activate LLDP.

lldp system-name <name> The router hostname advertised by LLDP.

lldp tx-interval <seconds> Transmit interval, which controls the time between LLDP messages in seconds.

lldp tx-hold <value> Transmit hold multiplier, which is multiplied by the transmit interval to calculate the total time used for the Time-To-Live (TTL) of the LLDP message.

Tip: If the transmit interval is 30 and the transmit hold multiplier is 4, then the advertised TTL of the LLDP message is 120 ($4 \times 30 = 120$).

Warning: The LLDP timer options do not have default values and must be manually set.

Example:

```
tnsr(config)# lldp system-name MyRouter
tnsr(config)# lldp tx-interval 30
tnsr(config)# lldp tx-hold 4
```

These parameters can be changed at any time.

20.1.2 LLDP Interface Configuration

Additional LLDP commands are available in `config-interface` mode (*Interface Command*) to configure per-interface LLDP identification:

lldp port-name <name> The name of the interface as advertised by LLDP. This must be present on an interface for it to participate in LLDP.

lldp management (ipv4|ipv6) <ip-address> The IPv4 and/or IPv6 address advertised by LLDP as a means to manage this router on this interface.

lldp management oid <oid> An object identifier associated with the management IP address on this interface.

Warning: LLDP requires `lldp port-name` to be present on at least one interface to function.

Example:

```
tnsr(config)# interface TenGigabitEthernet3/0/0
tnsr(config-interface)# lldp port-name MyPort
tnsr(config-interface)# lldp management ipv4 192.0.2.123
tnsr(config-interface)# lldp management ipv6 2001:db8::1:2:3:4
tnsr(config-interface)# exit
tnsr(config)#
```

Warning: Due to a limitation of the underlying API, all LLDP interface parameters must be configured at the same time. When LLDP parameters change, TNSR requires a *dataplane restart* for the new settings to take effect. See *Known Issues*.

20.2 LLDP Status

LLDP status cannot be viewed natively in TNSR at this time, but the status may be obtained directly from the dataplane at the CLI using the command `shell sudo vppctl show lldp`:

```
tnsr# shell sudo vppctl show lldp
Local interface      Peer chassis ID    Remote port ID    Last heard    Last sent
↪Status
TenGigabitEthernet3/0/0  54:78:1a:c0:ab:80  Fa0/1            17.1s ago    25.3s ago
↪active
```

Alternately, run the command directly from a host shell prompt and not through the TNSR CLI using `sudo vppctl show lldp`.

Warning: In either of these examples, the user account must have access to `sudo` to run the command.

PUBLIC KEY INFRASTRUCTURE

TNSR supports Public Key Infrastructure (PKI) X.509 certificates for various uses by the router and supporting software. PKI uses a pair of keys to encrypt and authenticate data, one public and one private. The private key is known only to its owner, and the public key can be known by anyone.

PKI works in an asymmetric fashion. A message is encrypted using the public key, and can only be decrypted by the private key. The private key can also be used to digitally sign a message to prove it originated from the key holder, and this signature can be validated using the public key. Combined with certificates, this provides a means to identify an entity and encrypt communications.

A Certificate Authority (CA) independently verifies the identity of the entity making a request for a certificate, and then signs a request, yielding a certificate. This certificate can then be validated against the certificate of the CA itself by anyone who has access to that CA certificate. In some cases, this CA may be an intermediate, meaning it is also signed by another CA above it. All together, this creates a chain of trust starting with the root CA all the way down to individual certificates. So as long as the CA is trustworthy, any certificate it has signed can be considered trustworthy.

Due to their size and private nature, certificates and keys are stored on the filesystem and not in the XML configuration. PKI files are stored under the following locations:

- Certificate Authorities: `/etc/pki/tls/tnsr/CA/`
- Certificates and Signing Requests: `/etc/pki/tls/tnsr/certs/`
- Private Keys: `/etc/pki/tls/tnsr/private/`

A key pair, CSR, and certificate associated with each other must all have the same name.

The process for creating a certificate is as follows:

- Create keys for `name`.
- Create a certificate signing request for `name` with the attributes to use for the certificate.
- Submit the CSR to a CA, which will sign the CSR and return a certificate.
- Enter or import the certificate contents for `name` into TNSR.

21.1 Key Management

Warning: Private keys are secret. These keys should never need to leave the firewall, with the exception of backups. The CA does not need the private key to sign a request.

TNSR can generate RSA key pairs with sizes of 2048, 3072, or 4096 bits. Larger keys are more secure than shorter keys. RSA Keys smaller than 2048 bits are no longer considered secure in practice, and are thus not allowed.

21.1.1 Generate a Key Pair

To generate a new key pair named `mycert` with a length of 4096 bits:

```
tnsr# pki private-key mycert generate key-length 4096
-----BEGIN PRIVATE KEY-----
[...]
-----END PRIVATE KEY-----
```

The key pair is stored in a file at `/etc/pki/tls/tnsr/private/<name>.key`.

Note: Remember that the private key, CSR, and certificate must all use identical names!

21.1.2 Importing a Key Pair

In addition to generating a key pair on TNSR, a private key may also be imported from an outside source. The key data can be imported in one of two ways:

- Use `pki private-key <name> enter` then copy and paste the PEM data
- Copy the PEM format key file to the TNSR host, then use `pki private-key <name> import <file>` to import from a file from the current working directory.

Copy and Paste

First, use the `enter` command:

```
tnsr# pki private-key mycert enter
Type or paste a PEM-encoded private key.
Include the lines containing 'BEGIN PRIVATE KEY' and 'END PRIVATE KEY'
```

Next, paste the key data:

```
-----BEGIN PRIVATE KEY-----
<key data>
-----END PRIVATE KEY-----
```

Import from File

First, make sure that the copy of the key file is in PEM format.

Next, copy the key file to TNSR and start the CLI from the directory containing this file. The filename extension is not significant, and may be `key`, `pem`, `txt`, or anything else depending on how the file was originally created.

Next, use the `import` command:

```
tnsr# pki private-key mycert import mycert.key
```

21.1.3 Other Key Operations

To view a list of all current keys known to TNSR:

```
tnsr# pki private-key list
mycert
```

To view the contents of the private key named `mycert` in PEM format:

```
tnsr# pki private-key mycert get
-----BEGIN PRIVATE KEY-----
<key data>
-----END PRIVATE KEY-----
```

Warning: When making a backup copy of this key, store the backup in a protected, secure location. Include the armor lines (BEGIN, END) when making a backup copy of the key.

To delete a key pair which is no longer necessary:

```
tnsr# pki private-key <name> delete
```

Warning: Do not delete a private key associated with a CSR or Certificate which is still in use!

21.2 Certificate Signing Request Management

A certificate signing request, or CSR, combines the public key along with a list of attributes that uniquely identify an entity such as a TNSR router. Once created, the CSR is exported and sent to the Certificate Authority (CA). The CA will sign the request and return a certificate.

21.2.1 Set Certificate Signing Request Attributes

The first step in creating a CSR is to set the attributes which identify this firewall. These attributes will be combined to form the certificate Subject:

```
tnsr# pki signing-request set common-name tnsr.example.com
tnsr# pki signing-request set country US
tnsr# pki signing-request set state Texas
tnsr# pki signing-request set city Austin
tnsr# pki signing-request set org Example Co
tnsr# pki signing-request set org-unit IT
```

The attributes include:

common-name The common name of the entity the certificate will identify, typically the fully qualified domain name of this host, or a username.

country The country in which the entity is located.

state The state or province in which the entity is located.

city The city in which the entity is located.

org The company name associated with the entity.

org-unit The department or division name inside the company.

Note: At a minimum, a common-name must be set to generate a CSR.

Next, set the required digest algorithm which will be used to create a hash of the certificate data:

```
tnsr# pki signing-request set digest sha256
```

This algorithm can be any of the following choices, from weakest to strongest: md5, sha1, sha224, sha256, sha384, or sha512.

Note: SHA-256 is the recommended minimum strength digest algorithm.

Before generating the CSR, review the configured attributes for the CSR:

```
tnsr# pki signing-request settings show
Certificate signing request fields:
  common-name: tnsr.example.com
  country: US
  state: Texas
  city: Austin
  org: Example Co
  org-unit: IT
  digest: sha256
```

If any attributes are incorrect, change them using the commands shown previously.

21.2.2 Generate a Certificate Signing Request

If the attributes are all correct, generate the CSR using the same name as the private key created previously. TNSR will output CSR data to the terminal in PEM format:

```
tnsr# pki signing-request mycert generate

-----BEGIN CERTIFICATE REQUEST-----
<csr data>
-----END CERTIFICATE REQUEST-----
```

The CSR data is stored in a file at `/etc/pki/tls/tnsr/certs/<name>.csr`

Note: Remember that the private key, CSR, and certificate must all use identical names!

The CSR data for existing entries can be displayed in PEM format:

```
tnsr# pki signing-request mycert get
-----BEGIN CERTIFICATE REQUEST-----
<csr data>
-----END CERTIFICATE REQUEST-----
```

Copy and paste the CSR data, including the armor lines (BEGIN, END), from the terminal into a local file, and submit that copy of the CSR to the CA for signing.

Warning: Remember, the private key for the CSR is not required for signing. Do not send the private key to the CA.

21.2.3 Other CSR Operations

A CSR entry may be deleted once the certificate has been imported to TNSR:

```
tnsr# pki signing-request <name> delete
```

To view a list of all CSR entries known to TNSR:

```
tnsr# pki signing-request list
```

To reset the CSR attribute contents:

```
tnsr# pki signing-request settings clear
```

21.3 Certificate Management

After submitting the certificate signing request to the CA, the CA will sign the request and return a signed copy of the certificate. Typically this will be sent in PEM format, the same format used for the CSR and private key.

The certificate data can be imported in one of two ways:

- Use `pki certificate <name> enter` then copy and paste the PEM data
- Copy the PEM format certificate file to the TNSR host, then use `pki certificate <name> import <file>` to import from a file from the current working directory.

The certificate data is stored in a file at `/etc/pki/tls/tnsr/certs/<name>.cert` after entering or importing the contents.

Warning: When importing a certificate created outside of TNSR, The private key must be imported and present before TNSR can import the certificate.

21.3.1 Copy and Paste

First, use the `enter` command:

```
tnsr# pki certificate mycert enter
Type or paste a PEM-encoded certificate.
Include the lines containing 'BEGIN CERTIFICATE' and 'END CERTIFICATE'
```

Note: Remember that the private key, CSR, and certificate must all use identical names!

Next, paste the certificate data:

```
-----BEGIN CERTIFICATE-----  
<cert data>  
-----END CERTIFICATE-----
```

21.3.2 Import from File

First, make sure that the copy of the certificate file is in PEM format. The CA may have delivered the certificate in PEM format, or another format. Convert the certificate to PEM format if it did not come that way.

Next, copy the certificate file to TNSR and start the CLI from the directory containing the certificate file. The filename extension is not significant, and may be `pem`, `crt`, `txt`, or anything else depending on how the file was delivered from the CA.

Next, use the `import` command:

```
tnsr# pki certificate mycert import mycert.pem
```

21.3.3 Other Certificate Operations

To view a list of all certificates known to TNSR:

```
tnsr# pki certificate list
```

To view the PEM data for a specific certificate known to TNSR:

```
tnsr# pki certificate <name> get
```

To delete a certificate:

```
tnsr# pki certificate <name> delete
```

21.4 Certificate Authority Management

As mentioned in *Public Key Infrastructure*, a Certificate Authority (CA) provides a starting point for a chain of trust between entities using certificates. A CA will sign a certificate showing that it is valid, and as long as an entity trusts the CA, it knows it can trust certificates signed by that CA.

By creating or importing a CA into TNSR, TNSR can use that CA to validate other certificates or sign new certificate requests. These certificates can then be used to identify clients connecting to the RESTconf service or other similar purposes.

A CA can be managed in several ways in TNSR. For example:

- Import a CA generated by another device by copy/paste in the CLI
- Import a CA generated by another device from a file
- Generate a new private key and CSR, then self-sign the CSR and set the CA property. The resulting CA is automatically available as a TNSR CA.

21.4.1 Import a CA

TNSR can import a CA from the terminal with copy/paste, or from a file. When importing a CA, the key is optional for validation but required for signing. To import the key, see [Key Management](#). Import the key with the same name as the CA.

To import a CA from the terminal, use the `enter` command. In this example, a CA named `tnsrca` will be imported from the terminal by TNSR:

```
# pki ca tnsrca enter
Type or paste a PEM-encoded certificate.
Include the lines containing 'BEGIN CERTIFICATE' and 'END CERTIFICATE'
-----BEGIN CERTIFICATE-----
<cert data>
-----END CERTIFICATE-----
tnsr(config)#
```

Next, import the private key using the same name:

```
tnsr(config)# pki private-key tnsrca enter
Type or paste a PEM-encoded private key.
Include the lines containing 'BEGIN PRIVATE KEY' and 'END PRIVATE KEY'
-----BEGIN PRIVATE KEY-----
<key data>
-----END PRIVATE KEY-----
```

Alternately, import the CA and key from the filesystem:

```
tnsr(config)# pki ca otherca import otherca.crt
tnsr(config)# pki private-key otherca import otherca.key
```

21.4.2 Creating a Self-Signed CA

TNSR can also create a self-signed CA instead of importing an external CA. For internal uses, this is generally a good practice since TNSR does not need to rely on public CA entries to determine trust for its own clients.

First, generate a new private key for the CA:

```
tnsr(config)# pki private-key selfca generate
-----BEGIN PRIVATE KEY-----
<key data>
-----END PRIVATE KEY-----
```

Next, create a new CSR for the CA:

```
tnsr(config)# pki signing-request set common-name selfca
tnsr(config)# pki signing-request set digest sha256
tnsr(config)# pki signing-request selfca generate
-----BEGIN CERTIFICATE REQUEST-----
<csr data>
-----END CERTIFICATE REQUEST-----
```

Finally, have TNSR self-sign the CSR while setting the CA flag on the resulting certificate:

```
tnsr(config)# pki signing-request selfca sign self enable-ca true
-----BEGIN CERTIFICATE-----
```

(continues on next page)

(continued from previous page)

```
<cert data>
-----END CERTIFICATE-----
```

After signing, the newly created CA is ready for immediate use:

```
tnsr(config)# pki ca list
      tnsrca
      selfca
```

21.4.3 Intermediate CAs

In some cases a CA may rely on another CA. For example, if a root CA signs an intermediate CA and the intermediate CA signs a certificate, then both the root CA and intermediate CA are required by the validation process.

To show this relationship in TNSR, a CA may be appended to another CA:

```
tnsr(config)# pki ca <root ca name> append <intermediate ca name>
```

In the above command, both CA entries must be present in TNSR before using the append command.

21.4.4 Using a CA to sign a CSR

A CA in TNSR with a private key present can also sign a client certificate. The typical use case for this is for RESTconf clients which must have a certificate recognized by a known CA associated with the RESTconf service.

First, generate a client private key and CSR:

```
tnsr(config)# pki private-key tnsrclient generate
-----BEGIN PRIVATE KEY-----
<key data>
-----END PRIVATE KEY-----
tnsr(config)# pki signing-request set common-name tnsrclient.example.com
tnsr(config)# pki signing-request set digest sha256
tnsr(config)# pki signing-request tnsrclient generate
-----BEGIN CERTIFICATE REQUEST-----
<csr data>
-----END CERTIFICATE REQUEST-----
```

Then, sign the certificate:

```
tnsr(config)# pki signing-request tnsrclient sign ca-name tnsrca days-valid 365
↳ digest sha512 enable-ca false
-----BEGIN CERTIFICATE-----
<cert data>
-----END CERTIFICATE-----
```

The sign command takes several parameters, each of which has a default safe for use with client certificates in this context. The above example uses these defaults, but specifies them manually to show how the parameters function. The available parameters are:

days-valid The number of days the resulting certificate will be valid. The default is 365 days (one year). When the certificate expires, it must be signed again for a new term. Certificates with a shorter lifetime are more secure, but longer lifetimes are more convenient.

digest The hash algorithm used to sign the certificate. The default value is sha512.

enable-ca A boolean value which sets the CA flag in the resulting certificate. If a CSR is signed as a CA, the resulting certificate can then be used to sign other certificates. For end user certificates this is not necessary or desired, so the default is `false`.

21.4.5 Other CA Operations

The remaining basic CA operations allow management of CA entries.

To view a list of all CA entries:

```
tnsr(config)# pki ca list
    tnsrca
    selfca
```

To view the contents of a CA certificate:

```
tnsr(config)# pki ca tnsrca get
-----BEGIN CERTIFICATE-----
<cert data>
-----END CERTIFICATE-----
```

To delete a CA entry:

```
tnsr(config)# pki ca tnsrca delete
```

BIDIRECTIONAL FORWARDING DETECTION

Bidirectional Forwarding Detection (BFD) is used to detect faults between two routers across a link, even if the physical link does not support failure detection. Even in cases where physical link issues occur and are detected, BFD can coordinate reaction to these failures rather than each component relying on its own failure detection methods.

TNSR uses UDP as a transport for BFD between directly connected routers (single hop/next hop) as described in [RFC 5880](#) and [RFC 5881](#).

Each BFD session monitors one link. Multiple BFD sessions are necessary to detect faults on multiple links. BFD sessions must be manually configured between endpoints as there is no method for automated discovery.

Note: The BFD implementation on TNSR only supports single hop BFD session in the dataplane. As such, BFD can only be configured on directly connected interfaces, between directly connected peers.

BFD supports session authentication using SHA1 for security, and the best practice is to use authentication when possible.

When using BFD, both endpoints transmit “Hello” packets back and forth between each other. If these packets are not received within the expected time frame, with the expected authentication information, the link is considered down. Links may also be administratively configured as down, and will not recover until manually changed.

TNSR currently supports BFD integration with BGP, OSPF, and OSPF6.

22.1 BFD Sessions

A BFD session defines a relationship between TNSR and a peer so they can exchange BFD information and detect link faults. These sessions are configured by using the `bfd session <name>` command, which enters `config-bfd` mode, and defines a BFD session using the given word for a name.

Example:

```
tnsr# conf
tnsr(config)# bfd session otherrouter
tnsr(config-bfd)# interface GigabitEthernet0/14/0
tnsr(config-bfd)# local address 203.0.113.2
tnsr(config-bfd)# peer address 203.0.113.25
tnsr(config-bfd)# desired-min-tx 100000
tnsr(config-bfd)# required-min-rx 100000
tnsr(config-bfd)# detect-multiplier 3
tnsr(config-bfd)# exit
tnsr(config)# exit
tnsr#
```

22.1.1 Session Parameters

interface <if-name> The Ethernet interface on which to enable BFD.

Warning: This interface must be directly connected to the peer (single hop), as the dataplane does not support BFD over multiple hops. It cannot be used with routing protocols running on loopback interfaces, for example.

local address <ip-address> The local address used as a source for BFD packets. This address must be present on <if-name>.

peer address <ip-address> The remote BFD peer address. The local and remote peer IP addresses must use the same address family (either IPv4 or IPv6)

desired-min-tx <microseconds> The desired minimum transmit interval, in microseconds

required-min-rx <microseconds> The required minimum transmit interval, in microseconds

detect-multiplier <n-packets> A non-zero value that is, roughly speaking, due to jitter, the number of packets that have to be missed in a row to declare the session to be down. Must be between 1 and 255.

Additional parameters for authentication are covered in [BFD Session Authentication](#).

22.1.2 Changing the BFD Administrative State

Under normal conditions the state of a link monitored by BFD is handled automatically. The link state can also be set manually when necessary.

To disable a link and mark it administratively down:

```
tnsr# bfd session <name>
tnsr(config-bfd)# enable false
```

To remove the administrative down and return the link to BFD management:

```
tnsr# bfd session <name>
tnsr(config-bfd)# enable true
```

22.1.3 Viewing BFD Session Status

To see the configuration and status of a BFD session, use the `show bfd session` command:

```
tnsr# show bfd session
Session Number: 0
  Local IP Addr: 203.0.113.2
  Peer  IP Addr: 203.0.113.25
  State: down
  Required Min Rx Interval: 100000 usec
  Desired Min Tx Interval: 100000 usec
  Detect Multiplier: 3
  BFD Key Id: 123
  Configuration Key Id: 14
  Authenticated: true
```

22.1.4 Using BFD Sessions

For BFD to function fully, the BFD session status must be consumed by other interested parties. Currently on TNSR this can be BGP, OSPF, or OSPF6 dynamic routing.

BGP BFD can be enabled for specific BGP neighbors with the `bfd enabled true` command from within `config-bgp-neighbor` mode.

OSPF/OSPF6 BFD can be enabled on specific OSPF interfaces with the `bfd enabled true` command from within `config-ospf[6]-if` mode.

In each case, the BGP neighbor or OSPF/OSPF6 interface must coincide exactly with the settings on a BFD session.

22.2 BFD Session Authentication

TNSR supports SHA1 and meticulous SHA1 authentication. In either mode, a secret key is used to create a hash of the outgoing packets. The key itself is not sent in the packets, only the hash and the ID of the key.

A sequence number is used to help avoid replay attacks. With SHA1, this sequence number is incremented occasionally. With meticulous SHA1, the sequence number is incremented on every packet.

The receiving peer will check for a key matching the given ID and then compare a hash of the BFD payload against the hash sent by the peer. If it matches and the sequence number is valid, the packet is accepted.

22.2.1 Define BFD Keys

There are two keys defined for each BFD session:

conf-key-id The Configuration Key ID. An unsigned 32-bit integer which identifies an internal unique key in TNSR. Neither the key itself nor this ID are **ever** communicated to peers. The secret component of this key must be generated outside of TNSR. It is a group of 1 to 20 hex pair values, such as 4a40369b4df32ed0652b548400.

bfd-key-id The BFD key ID. An unsigned 8-bit integer (0–255) which is the key ID carried in BFD packets, used for verifying authentication.

Warning: Both `conf-key-id` and `bfd-key-id` must be specified, or neither can be present.

To define a new configuration key ID:

```
tnsr(config)# bfd conf-key-id <conf-key-id>
tnsr(config-bfd-key)# authentication type (keyed-shal|meticulous-keyed-shal)
tnsr(config-bfd-key)# secret <(<hex-pair>)[1-20] >
```

For example:

```
tnsr(config)# bfd conf-key-id 123456789
tnsr(config-bfd-key)# authentication type meticulous-keyed-shal
tnsr(config-bfd-key)# secret 4a40369b4df32ed0652b548400
```

22.2.2 Setup BFD Authentication

Authentication will only be active if both the `bfd-key-id` and `conf-key-id` are defined for a BFD session.

An additional `delayed` keyword is also supported for BFD session which tells BFD to hold off any authentication action when receiving BFD messages until a peer attempts to authenticate or uses new credentials.

Warning: Only **one** host can have the `delayed` option enabled, otherwise credentials will never update as both peers will be waiting on the other one to act first.

Warning: BFD implementations vary, so authentication changes may disrupt live BFD sessions. The best practice to avoiding disruption when operating with third party BFD implementations is to set `delayed` on the TNSR side.

When adding authentication to an existing BFD session, or changing active authentication settings, make the changes first on the node with `delayed` set then configure the peer to match.

To activate authentication, add the chosen identifiers to a BFD session:

```
tnsr(config)# bfd session <bfd-session>
tnsr(config-bfd)# bfd-key-id <bfd-key-id>
tnsr(config-bfd)# conf-key-id <conf-key-id>
tnsr(config-bfd)# delayed (true|false)
tnsr(config-bfd)# exit
```

For example:

```
tnsr(config)# bfd session otherrouter
tnsr(config-bfd)# bfd-key-id 123
tnsr(config-bfd)# conf-key-id 123456789
tnsr(config-bfd)# delayed false
tnsr(config-bfd)# exit
```

22.2.3 View BFD Keys

To view a list of keys and their types, use the `show bfd keys` command:

```
tnsr# show bfd keys
Conf Key  Type                               Use Count
-----
123456789 meticulous-keyed-sha1         1
234567890 keyed-sha1                   0
```

To view only one specific key, pass its ID to the same command:

```
tnsr# show bfd keys conf-key-id 123456789
Conf Key  Type                               Use Count
-----
123456789 meticulous-keyed-sha1         1
```

22.3 BFD Example

This example establishes authenticated BFD between two routers which use OSPF to exchange routing information.

22.3.1 Configure BFD Authentication Keys

First, configure and check the authentication keys on both routers.

```
r1 tnsr(config)# bfd conf-key-id 123456789
r1 tnsr(config-bfd-key)# authentication type meticulous-keyed-sha1
r1 tnsr(config-bfd-key)# secret 4a40369b4df32ed0652b548400
r1 tnsr(config-bfd-key)# exit
```

```
r2 tnsr(config)# bfd conf-key-id 123456789
r2 tnsr(config-bfd-key)# authentication type meticulous-keyed-sha1
r2 tnsr(config-bfd-key)# secret 4a40369b4df32ed0652b548400
r2 tnsr(config-bfd-key)# exit
```

```
r1 tnsr# show bfd keys
Conf Key   Type                               Use Count
-----
123456789 meticulous-keyed-sha1 1
```

```
r2 tnsr# show bfd keys
Conf Key   Type                               Use Count
-----
123456789 meticulous-keyed-sha1 1
```

22.3.2 Configure BFD Sessions

Next, configure the BFD sessions on both routers using the authentication information configured in the previous section.

```
r1 tnsr(config)# bfd session r1_r2
r1 tnsr(config-bfd)# enable true
r1 tnsr(config-bfd)# interface TenGigabitEthernet6/0/0
r1 tnsr(config-bfd)# local address 203.0.113.2
r1 tnsr(config-bfd)# peer address 203.0.113.27
r1 tnsr(config-bfd)# desired-min-tx 1000000
r1 tnsr(config-bfd)# required-min-rx 1000000
r1 tnsr(config-bfd)# detect-multiplier 3
r1 tnsr(config-bfd)# bfd-key-id 123
r1 tnsr(config-bfd)# conf-key-id 123456789
r1 tnsr(config-bfd)# delayed true
r1 tnsr(config-bfd)# exit
r1 tnsr(config)# exit
```

Note: Note that since this node is being configured first, it has `delayed true` set, while the peer will have `false`. Since this is a new session, the difference is minimal, but when making future changes, this distinction is important. See [Setup BFD Authentication](#) for details.

```
r2 tnsr(config)# bfd session r2_r1
r2 tnsr(config-bfd)# enable true
r2 tnsr(config-bfd)# interface TenGigabitEthernet6/0/0
r2 tnsr(config-bfd)# local address 203.0.113.27
r2 tnsr(config-bfd)# peer address 203.0.113.2
```

(continues on next page)

(continued from previous page)

```
r2 tnsr(config-bfd)# desired-min-tx 1000000
r2 tnsr(config-bfd)# required-min-rx 1000000
r2 tnsr(config-bfd)# detect-multiplier 3
r2 tnsr(config-bfd)# bfd-key-id 123
r2 tnsr(config-bfd)# conf-key-id 123456789
r2 tnsr(config-bfd)# delayed false
r2 tnsr(config-bfd)# exit
r2 tnsr(config)# exit
```

22.3.3 Confirm BFD Status

With BFD configured on both nodes, check its status. The status should show a state of up and also indicate that the session is authenticated.

```
r1 tnsr# show bfd
Session Name: r1_r2
  Interface: TenGigabitEthernet6/0/0
  Local IP Addr: 203.0.113.2
  Peer  IP Addr: 203.0.113.27
  State: up
  Required Min Rx Interval: 1000000 usec
  Desired Min Tx Interval: 1000000 usec
  Detect Multiplier: 3
  BFD Key Id: 123
  Configuration Key Id: 123456789
  Authenticated: true
```

```
r2 tnsr# show bfd
Session Name: r2_r1
  Interface: TenGigabitEthernet6/0/0
  Local IP Addr: 203.0.113.27
  Peer  IP Addr: 203.0.113.2
  State: up
  Required Min Rx Interval: 1000000 usec
  Desired Min Tx Interval: 1000000 usec
  Detect Multiplier: 3
  BFD Key Id: 123
  Configuration Key Id: 123456789
  Authenticated: true
```

22.3.4 Setup OSPF

Now setup the routing protocol which will utilize the BFD status.

Note: BFD is activated by the `bfd enabled true` command on the `TenGigabitEthernet6/0/0` interface in OSPF. This is the same interface configured in BFD.

```
r1 tnsr(config)# route dynamic ospf
r1 tnsr(config-frr-ospf)# enable
r1 tnsr(config-frr-ospf)# server
r1 tnsr(config-ospf)# ospf router-id 10.2.0.1
```

(continues on next page)

(continued from previous page)

```

r1 tnsr(config-ospf)# interface TenGigabitEthernet6/0/0
r1 tnsr(config-ospf-if)# ip address * area 0.0.0.0
r1 tnsr(config-ospf-if)# ip network broadcast
r1 tnsr(config-ospf-if)# bfd enabled true
r1 tnsr(config-ospf-if)# exit
r1 tnsr(config-ospf)# int GigabitEthernet3/0/0
r1 tnsr(config-ospf-if)# ip address * area 0.0.0.0
r1 tnsr(config-ospf-if)# exit
r1 tnsr(config-ospf)# exit
r1 tnsr(config-frr-ospf)# exit

```

```

r2 tnsr(config)# route dynamic ospf
r2 tnsr(config-frr-ospf)# enable
r2 tnsr(config-frr-ospf)# server
r2 tnsr(config-ospf)# ospf router-id 10.27.0.1
r2 tnsr(config-ospf)# interface TenGigabitEthernet6/0/0
r2 tnsr(config-ospf-if)# ip address * area 0.0.0.0
r2 tnsr(config-ospf-if)# ip network broadcast
r2 tnsr(config-ospf-if)# bfd enabled true
r2 tnsr(config-ospf-if)# exit
r2 tnsr(config-ospf)# int GigabitEthernet3/0/0
r2 tnsr(config-ospf-if)# ip address * area 0.0.0.0
r2 tnsr(config-ospf-if)# exit
r2 tnsr(config-ospf)# exit
r2 tnsr(config-frr-ospf)# exit

```

22.3.5 Check OSPF Status

Check the status of OSPF to see if a neighbor relationship has been formed:

```
r1 tnsr(config)# show route dynamic ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
↪RXmtL RqstL DBsmL					
10.27.0.1	1	Full/Backup	36.415s	203.0.113.27	TenGigabitEthernet6/0/
↪0:203.0.113.2	0	0	0		

```
r2 tnsr(config)# show route dynamic ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
↪RXmtL RqstL DBsmL					
10.2.0.1	1	Full/DR	35.487s	203.0.113.2	TenGigabitEthernet6/0/
↪0:203.0.113.27	1	0	0		

22.3.6 Finish Up

Both routers are fully configured to use BFD and OSPF. If the TenGigabitEthernet6/0/0 interface fails, BFD will signal OSPF and the interface will be marked down in the OSPF daemon, and neighbors on that interface will be removed.

USER MANAGEMENT

TNSR includes a `tnsr` user by default. Administrators may create additional users to provide separate workspaces for each user. In this workspace the user may save and load configurations.

Warning: User access is controlled by NACM and the NACM default behavior varies by platform and when the TNSR installation was created. See *NETCONF Access Control Model (NACM)* for details.

23.1 User Configuration

Entering `config-auth` mode requires a username. When modifying an existing user, the username is available for autocompletion. The command will also accept a new username, which it creates when the configuration is committed. Creating a new user requires providing a means of authentication:

```
tnsr(config)# auth user <user-name>
```

A user may be deleted using the `no` form:

```
tnsr(config)# no auth user <user-name>
```

The `exit` command leaves `config-auth` mode:

```
tnsr(config-auth)# exit
tnsr(config)#
```

When exiting `config-auth` mode, TNSR commits changes to the user, which will create or update the entry for the user in the host operating system.

23.2 Authentication Methods

There are two methods for authenticating users: passwords and user keys.

23.2.1 Password Authentication

The password method takes a password entered in plain text, but stores a hashed version of the password in the configuration:

```
tnsr(config-auth)# password <plain text password>
```

Note: The password is hashed by the CLI prior to being passed to the backend. The plain text password is never stored or passed outside the specific CLI instance.

Warning: The password may be between 6 and 256 characters in length, though depending on the operating system default password hashing algorithm and key derivation behavior, the practical limit may be lower.

If the configuration is viewed using the `show configuration running` command, the hashed password will be present.

23.2.2 User Key Authentication

The second method of authentication is by user key. A user key is the same format as created by `ssh-keygen`.

To add a user key for authentication, use the `user-keys` command inside `config-auth` mode:

```
tnsr(config-auth)# user-keys <key-name>
```

The user key is read directly from the CLI. After the command is executed by pressing `Enter`, the CLI will wait for the key to be entered, typically by pasting it into the terminal or by typing. The end of input is indicated by a blank line. The normal CLI features are bypassed during this process.

NETCONF ACCESS CONTROL MODEL (NACM)

NETCONF Access Control Model (NACM) provides a means by which access can be granted to or restricted from groups in TNSR.

NACM is group-based and these groups and group membership lists are maintained in the NACM configuration.

User authentication is not handled by NACM, but by other processes depending on how the user connects. For examples, see *User Management* and *HTTP Server*.

See also:

The data model and procedures for evaluating whether a user is authorized to perform a given action are defined in [RFC 8341](#).

Warning: TNSR Does not provide protection against changing the rules in such a way that causes a loss of access. Should a lockout situation occur, see *Regaining Access if Locked Out by NACM*.

24.1 NACM Example

The example configuration in this section is the same default configuration shipped on TNSR version 18.08 mentioned in *NACM Defaults*.

Warning: In the following example, NACM is disabled first and activated at the end of the configuration. This avoids locking out the user when they are in the middle of creating the configuration, in case they unintentionally exit or commit before finishing.

```
tnsr(config)# nacm disable
tnsr(config)# nacm exec-default deny
tnsr(config)# nacm read-default deny
tnsr(config)# nacm write-default deny
tnsr(config)# nacm group admin
tnsr(config-nacm-group)# member root
tnsr(config-nacm-group)# member tnsr
tnsr(config-nacm-group)# exit
tnsr(config)# nacm rule-list admin-rules
tnsr(config-nacm-rule-list)# group admin
tnsr(config-nacm-rule-list)# rule permit-all
tnsr(config-nacm-rule)# module *
tnsr(config-nacm-rule)# access-operations *
tnsr(config-nacm-rule)# action permit
```

(continues on next page)

(continued from previous page)

```
tnsr(config-nacm-rule)# exit
tnsr(config-nacm-rule-list)# exit
tnsr(config)# nacm enable
tnsr(config)# exit
```

24.2 View NACM Configuration

The current NACM configuration can be viewed with the `show nacm` command:

```
tnsr# show nacm

NACM
====
NACM Enable: true
Default Read policy : deny
Default Write policy: deny
Default Exec policy : deny

Group: admin
-----
    root
    tnsr

Rule List: admin-rules
-----
Groups:
    admin

Name          Action Op Module Type
-----
permit-all   permit *   *
```

This may be narrowed down to only show part of the configuration.

To view all groups:

```
tnsr# show nacm group

NACM
====

Group: admin
-----
    root
    tnsr

Group: readonly
-----
    olly
    reed
```

To view a specific group, use `show nacm group <group-name>`:

```
tnsr# show nacm group admin
```

```
NACM
====
```

```
Group: admin
```

```
-----
    root
    tnsr
```

To view all rule lists:

```
tnsr# show nacm rule-list
```

```
NACM
====
```

```
Rule List: admin-rules
```

```
-----
```

```
Groups:
```

```
    admin
```

```
Name          Action Op  Module Type
```

```
-----
```

```
permit-all  permit *    *
```

```
Rule List: ro-rules
```

```
-----
```

```
Groups:
```

```
Name          Action Op  Module Type
```

```
-----
```

```
ro            permit exec *
```

```
read          deny    *    *
```

To view a specific rule list, use `show nacm rule-list <list-name>`:

```
tnsr# show nacm rule-list admin-rules
```

```
NACM
====
```

```
Rule List: admin-rules
```

```
-----
```

```
Groups:
```

```
    admin
```

```
Name          Action Op  Module Type
```

```
-----
```

```
permit-all  permit *    *
```

24.3 Enable or Disable NACM

Warning: Do not enable NACM unless the rules and groups are correctly and completely configured, otherwise access to TNSR may be cut off. If access is lost, see *Regaining Access if Locked Out by NACM*.

To enable NACM:

```
tnsr(config)# nacm enable
```

To disable NACM:

```
tnsr(config)# nacm disable
```

24.4 NACM Default Policy Actions

Alter the default policy for executing commands:

```
tnsr(config)# nacm exec-default <deny|permit>
```

Alter the default policy for reading status output:

```
tnsr(config)# nacm read-default <deny|permit>
```

Alter the default policy for writing configuration changes:

```
tnsr(config)# nacm write-default <deny|permit>
```

24.5 NACM Username Mapping

NACM does not authenticate users itself, but it does need to know the username to determine group membership.

The method of authentication determines the username as seen by NACM. For example, users authenticated by username and password (e.g. PAM auth for RESTCONF or the CLI) will have that same username in TNSR.

See also:

For more information on how users are authenticated, see [User Management](#) for CLI access and [HTTP Server](#) for access via RESTCONF.

CLI users can check their TNSR username with the `whoami` command.

NACM obeys the following rules to determine a username:

SSH Password NACM username is the same as the login username

SSH User Key NACM username is the same as the login username

HTTP Server Password NACM username is the same as the login username

HTTP Server Client Certificate NACM username is the Common Name of the user certificate (cn=subject component)

24.6 NACM Groups

To create a group, use the `nacm group <group-name>` command:

```
tnsr(config)# nacm group admin
```


This changes to the `config-nacm-group` mode where group members can be defined using the `member <username>` command:

```
tnsr(config-nacm-group)# member root
tnsr(config-nacm-group)# member tnsr
```

The username in this context is the mapped username described in [NACM Username Mapping](#).

Warning: Host operating system users that were created manually and not managed through TNSR cannot be used as group members. See [User Management](#) for information on managing users in TNSR.

To remove a member, use the `no` form of the command:

```
tnsr(config)# nacm group admin
tnsr(config-nacm-group)# no member tnsr
```

To remove a group, use `no nacm group <group-name>`:

```
tnsr(config)# no nacm group admin
```

24.7 NACM Rule Lists

NACM rules are contained inside a rule list. A rule list may contain multiple rules, and they are used in the order they are entered. Rule lists are also checked in the order they were created. Consider the order of lists and rules carefully when crafting rule lists.

Create a rule list:

```
tnsr(config)# nacm rule-list ro-rules
```

Set the group to which the rule list applies, use `group <group-name>`:

```
tnsr(config-nacm-rule-list)# group readonly
```

See also:

For information on defining groups, see [NACM Username Mapping](#).

24.8 NACM Rules

When configuring a rule list (`config-nacm-rule-list` mode), the `rule <name>` command defines a new rule:

```
tnsr(config-nacm-rule-list)# rule permit-all
```

After entering this command, the CLI will be in `config-nacm-rule` mode.

From here, a variety of behaviors for the rule can be set, including:

access-operations <name> The type of operation matched by this rule. Allowed values include:

- * Match all operations

- create** Any protocol operation that creates a new data node.

delete Any protocol operation that removes a data node.

exec Execution access to the specified protocol operation.

read Any protocol operation or notification that returns the value of a data node.

update Any protocol operation that alters an existing data node.

action <deny/permit> The action to take when this rule is matched, either `deny` to deny access or `permit` to allow access.

comment <text> Arbitrary text describing the purpose of this rule.

Next, the following types can be used to specify the restriction to be enacted by this rule:

module <*> The name of the Yang module covered by this rule, for example `netgate-nat`.

The complete list of modules can be viewed in the CLI by entering `module ?` from this mode. The [REST API documentation](#) also contains a list of modules.

path <path-name> XML path to restrict with this rule.

rpc <rpc-name> The name of an RPC call to be restricted by this rule, such as `edit-config`, `get-config`, and so on.

Warning: Users with access to modify the configuration (`edit-config`) should also be granted access to read the same paths (e.g. `get-config`). If a user only has `edit-config` access to a path, the user may receive an `access-denied` message in the CLI for that path when attempting to use a configuration command which makes a modification. This can happen because validation of certain commands requires reading the configuration to determine if the attempted command contains appropriate values.

24.8.1 NACM Rule Examples

As shown in [NACM Example](#), the following set of commands defines a rule list and then creates a rule to permit access to everything in TNSR:

```
tnsr(config)# nacm rule-list admin-rules
tnsr(config-nacm-rule-list)# group admin
tnsr(config-nacm-rule-list)# rule permit-all
tnsr(config-nacm-rule)# module *
tnsr(config-nacm-rule)# access-operations *
tnsr(config-nacm-rule)# action permit
tnsr(config-nacm-rule)# exit
tnsr(config-nacm-rule-list)# exit
```

Using the available `module` and `access-operation`, rules are possible that limit in more fine-grained ways.

This next example will allow a user in the `limited` group to see information from commands like `show`, but not make changes to the configuration:

```
tnsr(config)# nacm rule-list limited-rules
tnsr(config-nacm-rule-list)# group limited
tnsr(config-nacm-rule-list)# rule read-only
tnsr(config-nacm-rule)# module *
tnsr(config-nacm-rule)# access-operations read
tnsr(config-nacm-rule)# access-operations exec
tnsr(config-nacm-rule)# action permit
tnsr(config-nacm-rule)# exit
tnsr(config-nacm-rule-list)# exit
```

Selective restrictions are also possible with rules that limit access to specific modules while allowing access to everything else. In this example, users in the `limited` group may access any module except for NTP.

```
tnsr(config)# nacm rule-list limited-rules
tnsr(config-nacm-rule-list)# group limited
tnsr(config-nacm-rule-list)# rule no-ntp
tnsr(config-nacm-rule)# module netgate-ntp
tnsr(config-nacm-rule)# access-operations *
tnsr(config-nacm-rule)# action deny
tnsr(config-nacm-rule)# exit
tnsr(config-nacm-rule-list)# rule permit-all
tnsr(config-nacm-rule)# module *
tnsr(config-nacm-rule)# access-operations *
tnsr(config-nacm-rule)# action permit
tnsr(config-nacm-rule)# exit
tnsr(config-nacm-rule-list)# exit
```

24.9 NACM Rule Processing Order

When consulting defined rule lists, NACM acts in the following manner:

- If NACM is disabled, it skips all checks, otherwise it proceeds
- NACM consults group lists to find which groups contain this user
- NACM checks each rule list in the order they are defined
- NACM checks the group membership for each of these rule lists
- NACM compares the group defined on the rule list to the groups for this user, and if there is a match, it checks rules in the list
- NACM checks the rules in the order they are defined inside the rule list
- NACM compares the current access operation to the rule and if it matches, the rest of the rule is tested
- NACM attempts to match the following criteria, if defined on the rule:
 - The `module` on the rule name must match the requested module or `*`.
 - The `rpc-name` matches the RPC call in the request
 - The `path` matches the XML path to the requested data
- If the rule is matched, NACM consults the action on the rule and acts as indicated, either permitting or denying access
- NACM repeats these checks until there are no more rules, and then no more rule lists
- If no rules matched, NACM consults the default policies for the attempted operation and takes the indicated action

24.10 Regaining Access if Locked Out by NACM

If the NACM configuration prevents an administrator from accessing TNSR in a required way, NACM can be disabled or its configuration removed to regain access.

24.10.1 Method 1: Temporarily Disable NACM

With a complicated NACM configuration, the easiest way to regain access is to disable NACM, fix the configuration, and then enable it again. This involves disabling NACM in `/etc/tnsr.xml`, which is copied from one of the following locations, depending on which services are stopped/started: `/etc/tnsr/tnsr-none.xml`, `/etc/tnsr/tnsr-running.xml`, and `/etc/tnsr/tnsr-startup.xml`. The best practice is to edit all three files.

- Stop TNSR
- Edit `/etc/tnsr/tnsr-startup.xml`
- Locate the line with `CLICON_NACM_MODE` and change it to:

```
<CLICON_NACM_MODE>disabled</CLICON_NACM_MODE>
```

- Repeat the edit in `/etc/tnsr/tnsr-none.xml` and `/etc/tnsr/tnsr-running.xml`
- Restart TNSR
- Use the TNSR CLI to fix the broken NACM rules
- Save the new configuration
- Stop TNSR
- Edit `/etc/tnsr/tnsr-startup.xml`
- Locate the line with `CLICON_NACM_MODE` and change it to:

```
<CLICON_NACM_MODE>internal</CLICON_NACM_MODE>
```

- Repeat the edit in `/etc/tnsr/tnsr-none.xml` and `/etc/tnsr/tnsr-running.xml`
- Restart TNSR

TNSR will start with the new, fixed, NACM configuration. If access is still not working properly, repeat the process making changes to NACM until it is, or proceed to the next method to start over.

24.10.2 Method 2: Remove NACM Configuration

- Stop TNSR
- Edit `/var/tnsr/startup_db`
- Remove the entire `<nacm>...</nacm>` section from `startup_db`
- Start TNSR

TNSR will restart without any NACM configuration and it can then be reconfigured from scratch as shown in [NACM Example](#).

24.11 NACM Defaults

TNSR version 18.08 or later includes a default set of NACM rules. These rules allow members of group `admin` to have unlimited access and sets the default policies to `deny`. This configuration includes the users `tnsr` and `root` in the group `admin`.

See also:

To see the specific rules from the default configuration, see [NACM Example](#) or view the current NACM configuration as described in [View NACM Configuration](#).

For users of older installations or those who have removed the default NACM configuration, NACM defaults to disabled with no defined groups or rule lists, and with the following default policies:

```
Default Read policy : permit
Default Write policy: deny
Default Exec policy : permit
```

HTTP SERVER

TNSR includes an HTTP server, currently powered by [nginx](#). This HTTP server provides clients with access to the RESTCONF API, and there are plans to extend it to provide other services in the future.

25.1 HTTP Server Configuration

The server is configured using the `http server` command to enter `http` mode:

```
tnsr# configure
tnsr(config)# http server
tnsr(config-http)#
```

The server can be disabled with the following command:

```
tnsr(config)# no http server
```

25.1.1 Managing the HTTP Server Process

The HTTP server process can be managed using the `service` command:

```
tnsr# configure
tnsr(config)# service http <command>
```

Where `<command>` can be any of:

- start** Start the HTTP server
- stop** Stop the HTTP server
- restart** Restart (stop and then start) the HTTP server
- status** Print the status of the HTTP server process

25.2 HTTPS Encryption

The HTTP server can optionally utilize TLS (HTTPS) to secure communications between the client and server.

Warning: Though HTTPS is optional, we strongly recommend its use for optimal security.

HTTPS requires a server certificate present on the TNSR device, and this server certificate must be configured in the HTTP server:

```
tnsr(config)# http server
tnsr(config-http)# server certificate <cert-name>
```

See also:

For more information on managing certificates on TNSR, see [Public Key Infrastructure](#).

25.3 Authentication

The HTTP server supports three types of client authentication to protect access to its resources: Client certificate authentication, password authentication, and none (no authentication):

```
tnsr(config-http)# authentication type (client-certificate|password|none)
```

25.3.1 Client Certificate

The most secure means of protecting access to the HTTP server is via client certificates:

```
tnsr(config-http)# authentication type client-certificate
tnsr(config-http)# authentication client-certificate-ca <cert-name>
```

To verify client certificates, a Certificate Authority (CA) is configured in TNSR and all client certificates must be signed by this CA. The client certificate must be used by the client when attempting to connect to the HTTP server. Clients without a certificate are rejected.

See also:

For more information on managing certificates on TNSR, see [Public Key Infrastructure](#).

When using client certificates the Common Name (cn= parameter) of the client certificate is taken as the username. That username is then processed through NACM to determine group access privileges for the RESTCONF API.

25.3.2 Password

Password authentication for the HTTP server is handled via Pluggable Authentication Modules (PAM) support:

```
tnsr(config-http)# authentication type password
```

Users can be authenticated against any source supported by PAM modules in the operating system.

Once authenticated, the username is processed through NACM to determine group access privileges for the RESTCONF API.

25.3.3 None

The least secure option is to disable authentication entirely:

```
tnsr(config-http)# authentication type none
```

Warning: This option must only be used for testing and never in a production environment.

This removes all security protecting the RESTCONF API. Without authentication, any client can send requests or make changes using the API, which is extremely dangerous.

25.4 RESTCONF Server

The primary service provided by the HTTP server is the [API Endpoints](#) which uses RESTCONF. This RESTCONF service can be enabled and disabled as needed within the HTTP server configuration.

To enable access to the RESTCONF API:

```
tnsr(config-http)# enable restconf
```

To disable access to the RESTCONF API:

```
tnsr(config-http)# disable restconf
```


MONITORING

Monitoring of a TNSR system, either locally or remotely, can be accomplished in several ways:

- From the CLI, using `show` commands.
- *API Endpoints* which provide state information.
- Through the SNMP service

See also:

Refer to the [REST API documentation](#) and *RESTCONF Service Setup with Certificate-Based Authentication and NACM* for details and examples for configuration and use of the RESTCONF API.

26.1 Monitoring Interfaces

Each interface has associated counters, which enable traffic volume and error monitoring.

Note: To limit the amount of administrative traffic, VPP only updates these counters every 10 seconds.

There are several commands used to monitor interfaces: `show interface`, `show interface counters`, `interface clear counters`, and `show packet-counters`. Additionally, the counters may be retrieved using RESTCONF.

26.1.1 show interface

The `show interface` command prints important traffic volume and error counters specific to each interface. For example:

```
tnsr# show interface

Interface: TenGigabitEthernet6/0/0
  Admin status: up
  Link up, link-speed 1000 Mbps, full duplex
  Link MTU: 1500 bytes
  MAC address: 00:90:0b:7a:8a:67
  IPv4 Route Table: ipv4-VRF:0
  IPv4 addresses:
    203.0.113.2/24
  IPv6 Route Table: ipv6-VRF:0
  IPv6 addresses:
```

(continues on next page)

(continued from previous page)

```

    2001:db8:0:2::2/64
VLAN tag rewrite: disabled
counters:
  received: 9253580 bytes, 61588 packets, 0 errors
  transmitted: 628148 bytes, 5755 packets, 8 errors
  protocols: 12810 IPv4, 5101 IPv6
  50972 drops, 0 punts, 0 rx miss, 0 rx no buffer

```

The `show interface` command also supports filtering of its output. When the list is filtered, its name, description, and administrative status are printed along with the chosen output.

access-list Prints the access control lists configured on an interface

counters Prints the interface traffic counters for an interface

ip [(nat|vrrp-virtual-router)] Prints the IPv4 addresses present on the interface and the IPv4 route table used by the interface.

nat Prints the NAT role for an interface (e.g. inside or outside)

vrrp-virtual-router Prints the IPv4 VRRP status for an interface.

ipv6 [vrrp-virtual-router] Prints the IPv6 addresses present on the interface and the IPv6 route table used by the interface.

vrrp-virtual-router Prints the IPv6 VRRP status for an interface.

link Prints the link status (e.g. up or down), media type and duplex, and MTU

mac-address Prints the hardware MAC address, if present

vlan tag-rewrite Shows VLAN tag rewriting attributes for an interface.

These keywords may be used with the entire list of interfaces, for example:

```
tnsr# show interface ip
```

The filtering may also be applied to a single interface:

```
tnsr# show interface TenGigabitEthernet6/0/0 link
```

26.1.2 show interface counters

The `show interface [<if-name>] counters [verbose]` command displays detailed information on all available interface counters.

Example output:

```

tnsr(config)# show interface TenGigabitEthernet6/0/0 counters
Interface: TenGigabitEthernet6/0/0
counters:
  received: 9253580 bytes, 61588 packets, 0 errors
  transmitted: 628148 bytes, 5755 packets, 8 errors
  protocols: 12810 IPv4, 5101 IPv6
  50972 drops, 0 punts, 0 rx miss, 0 rx no buffer

```

Additional detailed packet counters for transmit and receive of unicast, multicast, and broadcast traffic may be enabled or disabled on a per-interface basis (*Interface Configuration Options*). Add the `verbose` keyword to display these statistics:

```
tnsr(config)# show interface TenGigabitEthernet6/0/0 counters verbose
Interface: TenGigabitEthernet6/0/0
detailed counters:
  received: 9258555 bytes, 61641 packets, 0 errors
  received unicast: 2464 bytes, 18 packets
  received multicast: 2464 bytes, 18 packets
  received broadcast: 622 bytes, 8 packets
  transmitted: 628676 bytes, 5761 packets, 8 errors
  transmitted unicast: 2480 bytes, 18 packets
  transmitted multicast: 2480 bytes, 18 packets
  transmitted broadcast: 0 bytes, 0 packets
  protocols: 12820 IPv4, 5105 IPv6
  51016 drops, 0 punts, 0 rx miss, 0 rx no buffer
```

Counter values take a minimum of 10 seconds to be populated with valid data.

26.1.3 clear interface counters

The interface `clear counters <name>` command clears all counters on a given interface. This command is available in config mode. If no specific interface is given, all interfaces will have their counters cleared:

```
tnsr# configure
tnsr(config)# interface clear counters
Counters cleared
tnsr(config)#
```

26.1.4 show packet-counters

The `show packet-counters` command prints packet statistics and error counters taken from the dataplane. These counters show counts of packets that have passed through various aspects of processing, such as encryption, along with various types of packet send/receive errors. The set of counters displayed will vary depending on the set of enabled features, such as NAT, IPsec, and so on.

Example output:

```
tnsr# show packet-counters
Count      Node                      Reason
624        dpdk-crypto-input         Crypto ops dequeued
624        dpdk-esp-decrypt-post     ESP post pkts
624        dpdk-esp-decrypt         ESP pkts received
622        esp-encrypt              ESP pkts received
624        ipsec-if-input           good packets received
304        ip4-input                Multicast RPF check failed
9          ip4-arp                  ARP requests sent
22        lldp-input               lldp packets received on disabled_
↪ interfaces
8          ethernet-input           no error
2          ethernet-input           unknown ethernet type
5821       ethernet-input           unknown vlan
16        arp-input                ARP request IP4 source address_
↪ learned
28        GigabitEthernet0/14/0-output interface is down
8         GigabitEthernet3/0/0-output interface is down
```

26.1.5 Interface status via API

If the *RESTCONF API* is enabled, the interface counter data may also be polled that way. For example:

Command:

```
$ curl --cert ~/tnsr/tnsr-myuser.crt \
--key ~/tnsr/tnsr-myuser.key \
--cacert ~/tnsr/tnsr-selfca.crt \
-X GET \
http://tnsr.example.com/restconf/data/netgate-interface:interfaces-state/
→interface=TenGigabitEthernet6%2F0%2F0/counters/
```

Output:

```
{
  "netgate-interface:counters": {
    "collect-time": 1563807148,
    "reset-time": 0,
    "detailed-counters": true,
    "rx-bytes": 120317,
    "rx-packets": 736,
    "rx-unicast-bytes": 19775,
    "rx-unicast-packets": 102,
    "rx-multicast-bytes": 97965,
    "rx-multicast-packets": 597,
    "rx-broadcast-bytes": 2577,
    "rx-broadcast-packets": 37,
    "rx-ip4": 175,
    "rx-ip6": 57,
    "tx-bytes": 15530,
    "tx-packets": 101,
    "tx-unicast-bytes": 15178,
    "tx-unicast-packets": 95,
    "tx-multicast-bytes": 226,
    "tx-multicast-packets": 3,
    "tx-broadcast-bytes": 126,
    "tx-broadcast-packets": 3,
    "drop": 601,
    "punt": 0,
    "rx-no-buffer": 0,
    "rx-miss": 0,
    "rx-error": 0,
    "tx-error": 21
  }
}
```

26.2 Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) service on TNSR provides a method through which the router can be monitored by a Network Monitoring System (NMS) or other software which supports SNMP.

SNMP presents information about the router to clients organized in an object identifier (OID) tree which is defined by Management Information Base (MIB) files. SNMP clients can access information by using a numerical OID or by using names looked up from MIB files.

The SNMP daemon currently supports the View-based Access Control Model (VACM). In this model, groups of communities are allowed access to SNMP information defined by views, which grant or limit their access.

Note: Future versions of TNSR will support SNMPv3 for more secure access control.

The SNMP service will respond to requests from host OS management interfaces as well as TNSR interfaces, if allowed by ACLs.

Warning: Access to the SNMP service on UDP port 161 should be limited by ACLs so that only authorized management hosts are able to reach the service.

26.2.1 Enable or Disable the SNMP Service

The SNMP server is enabled and disabled by the `server snmp (disable|enable)` command.

To enable the SNMP service:

```
tnsr(config)# server snmp enable
```

To disable the SNMP service:

```
tnsr(config)# server snmp disable
```

26.2.2 Control the SNMP Service

The SNMP service is controlled by the `service snmp (start|stop|restart|status)` command.

In most cases manual control of the service is unnecessary as the server will start and stop as needed based on the configuration.

TNSR automatically restarts the SNMP service when making changes to the SNMP configuration.

26.2.3 SNMP Configuration

The SNMP configuration is managed using the `snmp` command from `config` mode. This command has several options which are collectively used to define VACM rules to grant access to clients.

Note: TNSR automatically restarts the SNMP service when making changes to the SNMP configuration. The SNMP service will be momentarily unavailable while the service is reloading with the new configuration.

Tip: Basic system information for SNMP such as contact (`sysContact`), location (`sysLocation`), description (`sysDescr`), and hostname (`sysName`) is pulled from the values configured using the `system` command. See [Basic System Information](#) for information on setting these values.

SNMP Communities

An SNMP community in SNMPv1 and SNMPv2c is similar to a username and password in a single string. The community name is given by a client and checked against communities listed in the SNMP configuration. If the community is known, and the source of the request matches the source defined for the community, then the request continues on to have its access checked further.

Warning: SNMPv1 and SNMPv2c are not encrypted. Only allow access to the SNMP daemon from management networks or similar secure locations.

A community entry maps a traditional SNMP community name (e.g. `public`) to a VACM security name:

```
tnsr(config)# [no] snmp community community-name <community-name>
                               source (<src-prefix>|default)
                               security-name <security-name>
```

The following parameters are available:

community-name <community-name> The name for this community.

Warning: The SNMP community name should be considered as a password. Do not use an easily guessed name, and keep the community name a secret from others. Do not transmit the community name over an insecure network.

source (<src-prefix>|default) The IPv4 or IPv6 source network from which requests for this community will originate. For example, a management network.

The keyword `default` may also be used for the source, which allows a request from any source.

Warning: The best practice is to limit access by source so that only specific clients may access SNMP information. Avoid using `default` if at all possible.

security-name <security-name> The VACM security name to which this community should be mapped. This name is then used in groups.

This command may be repeated multiple times. Thus, multiple sources can set be for the same community. It can also be used to setup more complex policies such as different sources for the same community being mapped to different security names, or mapping multiple communities/sources to the same security name.

SNMP Groups

A group defines a VACM group, which is a collection of security names that have the same level of access.

```
tnsr(config)# [no] snmp group group-name <group-name>
                               security-name <security-name>
                               security-model (any|v1|v2c)
```

The following parameters are available:

group-name <group-name> The name of this group, which is used by access rules.

security-name <security-name> The security name to add as a member of this group.

Note: In SNMPv1 and SNMPv2, the security name is mapped from a community entry (*SNMP Communities*). In future versions with SNMPv3, this may also be a SNMPv3 security name (e.g. USM username, TSM identity, etc).

security-model (any|v1|v2c) The source of this security name, based on how its connection was authorized.

This command may be repeated to add multiple members to the same group.

SNMP Views

A **view** defines a subset of the entire SNMP object identifier (OID) tree. Multiple views with the same name may be defined to build a collection of OIDs to which groups may be granted access.

```
tnsr(config)# [no] snmp view view-name <view-name>
                        view-type (included|excluded)
                        oid <oid>
```

The following parameters are available:

view-name <view-name> The name of this view. Used in access rules to grant read and write access to portions of the OID tree.

view-type (included|excluded) Sets the type of view being defined.

included When set, objects under **oid** will be included in the view.

excluded When set, objects under **oid** will be excluded from the view.

oid <oid> The base **oid** under which this view either includes or excludes objects. This may be specified numerically or using names known to the SNMP daemon from MIB files.

For example, the root OID **.1** may also be given by its name **.iso**. Refer to MIB files for details.

This command may be repeated to define complex views which may include and exclude portions of the same OID hierarchy.

SNMP Access Rules

An **access** rule defines which views may be accessed by a given group. This ties together the other VACM entries, granting access to clients.

```
tnsr(config)# [no] snmp access group-name <group-name>
                        prefix (exact|prefix)
                        model (any|v1|v2c)
                        level (noauth|auth|priv)
                        read <read-view>
                        write <write-view>
```

The following parameters are available:

group-name <group-name> The name of the group being granted access, as defined by VACM group entries (*SNMP Groups*).

prefix (exact|prefix) Used by SNMPv3 to control how a context on the rule is applied to the context of the incoming connection. Since SNMPv3 is not yet supported, this must be set to **exact**.

model (any|v1|v2c) The security model of the client connection, based on how its connection was authorized.

level (auth|noauth|priv) The minimum security level at which this access rule will be allowed. Since SNMPv3 and transport security are not yet supported, this must be set to noauth.

read (<read-view>|none) The name of the view (*SNMP Views*) to which matching clients will have read access. Use none to deny read access.

write (<write-view>|none) The name of the view (*SNMP Views*) to which matching clients will have write access. Use none to deny write access.

26.2.4 SNMP Example

The following example sets up SNMP access for a single community name which can read anything under .1 (.iso) in the OID tree, and does not write access.

```
snmp community community-name tnsrmon source 10.2.4.0/24 security-name TNSRMonitor
snmp group group-name ROGroup security-name TNSRMonitor security-model v1
snmp group group-name ROGroup security-name TNSRMonitor security-model v2c
snmp view view-name systemview view-type included oid .1
snmp access group-name ROGroup prefix exact model any level noauth read systemview_
↪write none
```

Following through line by line:

First, map the SNMPv1/SNMPv2c community named `tnsrmon` to the security name `TNSRMonitor` for clients connecting from `10.2.4.0/24`, which in this example is a secure management network.

```
snmp community community-name tnsrmon source 10.2.4.0/24 security-name TNSRMonitor
```

Next, define a group named `ROGroup`, and specify that if the `TNSRMonitor` security name connects using SNMPv1, it is considered a member of this group.

```
snmp group group-name ROGroup security-name TNSRMonitor security-model v1
```

Add another entry to `ROGroup` for `TNSRMonitor` if it connects using SNMPv2c

```
snmp group group-name ROGroup security-name TNSRMonitor security-model v2c
```

Now define a view named `systemview` which includes the entire OID tree under `.1`. This could also have been specified by name, e.g. `.iso`.

```
snmp view view-name systemview view-type included oid .1
```

Finally, tie all the entries together by granting access for `ROGroup` to read from `systemview` when it connects using any security model, but do not specify a write group so that it has no write access.

```
snmp access group-name ROGroup prefix exact model any level noauth read systemview_
↪write none
```

Note: Since SNMPv3 is not yet supported, the values for `prefix` and `level` must be set as shown. See *SNMP Access Rules*.

TNSR CONFIGURATION EXAMPLE RECIPES

This section is a cookbook full of example recipes which can be used to quickly configure TNSR in a variety of ways. The use cases covered by these recipes are real-world problems encountered by Netgate customers.

These example scenarios pull together concepts discussed in more detail throughout the rest of this documentation to accomplish larger goals.

27.1 RESTCONF Service Setup with Certificate-Based Authentication and NACM

27.1.1 Use Case

RESTCONF is desirable for its ability to implement changes to TNSR remotely using the API, but allowing remote changes to TNSR also raises security concerns. When using RESTCONF, security is extremely important to protect the integrity of the router against unauthorized changes.

Note: RESTCONF deals in JSON output and input, which is easily parsed by a variety of existing libraries for programming and scripting languages.

27.1.2 Example Scenario

In this example, TNSR will be configured to allow access via RESTCONF, but the service will be protected in several key ways:

- The RESTCONF service is configured for TLS to encrypt the transport
- The RESTCONF service is configured to require a client certificate, which is validated against a private Certificate Authority known to TNSR
- NACM determines if the certificate common-name (username) is allowed access to view or make changes via RESTCONF

Item	Value
TNSR Hostname	tnsr.example.com
RESTCONF Username	myuser
NACM Group Name	admins
Additional User	anotheruser

27.1.3 TNSR Setup

Generate Certificates

Create a self-signed Certificate Authority:

```
tnsr(config)# pki private-key selfca generate
tnsr(config)# pki signing-request set common-name selfca
tnsr(config)# pki signing-request set digest sha256
tnsr(config)# pki signing-request selfca generate
tnsr(config)# pki signing-request selfca sign self enable-ca true
```

Create a certificate for the user myuser, signed by selfca:

```
tnsr(config)# pki private-key myuser generate key-length 4096
tnsr(config)# pki signing-request set common-name myuser
tnsr(config)# pki signing-request set digest sha256
tnsr(config)# pki signing-request myuser generate
tnsr(config)# pki signing-request myuser sign ca-name selfca days-valid 365 digest_
↪sha512 enable-ca false
```

Create a certificate for the RESTCONF service to use. The common-name should be the hostname of the TNSR router, which should also exist in DNS:

```
tnsr(config)# pki private-key restconf generate key-length 4096
tnsr(config)# pki signing-request set common-name tnsr.example.com
tnsr(config)# pki signing-request set digest sha256
tnsr(config)# pki signing-request restconf generate
tnsr(config)# pki signing-request restconf sign ca-name selfca days-valid 365 digest_
↪sha512 enable-ca false
```

Setup NACM

Disable NACM while making changes, to avoid locking out the account making the changes:

```
tnsr(config)# nacm disable
```

Set default policies:

```
tnsr(config)# nacm exec-default deny
tnsr(config)# nacm read-default deny
tnsr(config)# nacm write-default deny
```

Setup an admin group containing the default users plus myuser, which will match the common-name of the user certificate created above:

```
tnsr(config)# nacm group admin
tnsr(config-nacm-group)# member root
tnsr(config-nacm-group)# member tnsr
tnsr(config-nacm-group)# member myuser
tnsr(config-nacm-group)# exit
```

Setup rules to permit any action by members of the admin group:

```
tnsr(config)# nacm rule-list admin-rules
tnsr(config-nacm-rule-list)# group admin
tnsr(config-nacm-rule-list)# rule permit-all
tnsr(config-nacm-rule)# module *
tnsr(config-nacm-rule)# access-operations *
tnsr(config-nacm-rule)# action permit
tnsr(config-nacm-rule)# exit
tnsr(config-nacm-rule-list)# exit
```

Enable NACM:

```
tnsr(config)# nacm enable
tnsr(config)# exit
```

Enable RESTCONF

Enable RESTCONF and configure it for TLS and client certificate authentication:

```
tnsr(config)# http server
tnsr(config-http)# server certificate restconf
tnsr(config-http)# authentication type client-certificate
tnsr(config-http)# authentication client-certificate-ca selfca
tnsr(config-http)# enable restconf
```

27.1.4 Client Configuration

On TNSR, export the CA certificate, user certificate, and user certificate key. Place the resulting files in a secure place on a client system, in a directory with appropriate permissions, readable only by the user. Additionally, the private key file must only be readable by the user. For this example, the files will be placed in `~/tnsr/`.

First, export the CA certificate. Copy and paste this into a local file, named `tnsr-selfca.crt`:

```
tnsr# pki ca selfca get
-----BEGIN CERTIFICATE-----
[...]
-----END CERTIFICATE-----
```

Next, export the user certificate, copy and paste it and save in a local file named `tnsr-myuser.crt`:

```
tnsr# pki certificate myuser get
-----BEGIN CERTIFICATE-----
[...]
-----END CERTIFICATE-----
```

Finally, export the user certificate private key, copy and paste it and save in a local file named `tnsr-myuser.key`. Remember to protect this file so it is only readable by this user:

```
tnsr# pki private-key myuser get
-----BEGIN PRIVATE KEY-----
[...]
-----END PRIVATE KEY-----
```

This example uses `curl` to access RESTCONF, so ensure it is installed and available on the client computer.

27.1.5 Example Usage

This simple example shows fetching the contents of an ACL from RESTCONF as well as adding a new ACL entry. There are numerous possibilities here, for more details see the [REST API documentation](#).

In this example, there is an existing ACL named `blockbadhosts`. It contains several entries including a default allow rule with a sequence number of 5000.

These examples are all run from the client configured above.

Note: This is a simple demonstration using cURL and shell commands. This makes it easy to demonstrate how the service works, and how RESTCONF URLs are formed, but does not make for a good practical example.

In real-world cases these types of queries would be handled by a program or script that interacts with RESTCONF, manipulating data directly and a lot of the details will be handled by RESTCONF and JSON programming libraries.

Retrieve a specific ACL

Retrieve the entire contents of the `blockbadhosts` ACL:

Command:

```
$ curl --cert ~/tnsr/tnsr-myuser.crt \
--key ~/tnsr/tnsr-myuser.key \
--cacert ~/tnsr/tnsr-selfca.crt \
-X GET \
https://tnsr.example.com/restconf/data/netgate-acl:acl-config/acl-table/acl-
↪list=blockbadhosts
```

Output:

```
{
  "netgate-acl:acl-list": [
    {
      "acl-name": "blockbadhosts",
      "acl-description": "Block bad hosts",
      "acl-rules": {
        "acl-rule": [
          {
            "sequence": 1,
            "action": "deny",
            "ip-version": "ipv4",
            "src-ip-prefix": "203.0.113.14/32"
          },
          {
            "sequence": 2,
            "action": "deny",
            "ip-version": "ipv4",
            "src-ip-prefix": "203.0.113.15/32"
          },
          {
            "sequence": 555,
            "action": "deny",
            "ip-version": "ipv4",
            "src-ip-prefix": "5.5.5.5/32"
          }
        ]
      }
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    {
      "sequence": 5000,
      "acl-rule-description": "Default Permit",
      "action": "permit",
      "ip-version": "ipv4"
    }
  ]
}

```

The cURL parameters and RESTCONF URL can be dissected as follows:

Item	Value
cURL Client Certificate	--cert ~/tnsr/tnsr-myuser.crt
cURL Client Certificate Key	--key ~/tnsr/tnsr-myuser.key
cURL CA Cert to validate TLS	--cacert ~/tnsr/tnsr-selfca.crt
Request type (GET)	-X GET
RESTCONF Server protocol/host	https://tnsr.example.com
RESTCONF API location:	/restconf/data/
ACL config area (prefix:name)	netgate-acl:acl-config/
ACL table	acl-table/
ACL List, with restriction	acl-list=blockbadhosts

Note: Lists of items with a unique key can be restricted as shown above. The API documentation also calls this out as well, showing an optional = { name } in the query.

Retrieve a specific rule of a specific ACL

View only the default permit rule of the ACL:

Command:

```

$ curl --cert ~/tnsr/tnsr-myuser.crt \
  --key ~/tnsr/tnsr-myuser.key \
  --cacert ~/tnsr/tnsr-selfca.crt \
  -X GET \
  https://tnsr.example.com/restconf/data/netgate-acl:acl-config/acl-table/acl-
  ↳list=blockbadhosts/acl-rules/acl-rule=5000

```

Output:

```

{
  "netgate-acl:acl-rule": [
    {
      "sequence": 5000,
      "acl-rule-description": "Default Permit",
      "action": "permit",
      "ip-version": "ipv4"
    }
  ]
}

```

(continues on next page)

(continued from previous page)

```
  ]
}
```

The query is nearly identical to the previous one, with the following additional components:

Item	Value
ACL rules list	acl-rules/
ACL rule, with restriction	acl-rule=5000

Add a new rule to an existing ACL

Insert a new ACL rule entry with the following parameters:

Item	Value
Request Type	-X PUT (add content)
Content Type	-H "Content-Type: application/yang-data+json"
ACL Name	blockbadhosts
ACL Rule Sequence	10
ACL Rule Action	deny
ACL Rule Source Address	10.222.111.222/32

The new data passed in the `-d` parameter is JSON but with all whitespace removed so it can be more easily expressed on a command line.

Warning: The `Content-Type` header must be set when performing a write operation such as PUT or PATCH. The value of the header must reflect the type of data being sent. These examples use JSON, so the header is set to `application/yang-data+json`. When submitting XML, it would be `application/yang-data+xml`.

The URL is the same as if the query is retrieving the rule in question.

Warning: Note the presence of the sequence number in both the supplied JSON data and in the URL. This must match.

Command:

```
$ curl --cert ~/tnsr/tnsr-myuser.crt \
  --key ~/tnsr/tnsr-myuser.key \
  --cacert ~/tnsr/tnsr-selfca.crt \
  -H "Content-Type: application/yang-data+json" \
  -X PUT \
  -d '{"netgate-acl:acl-rule":[{"sequence": 10,"action":"deny","ip-version":"ipv4",
  ↪ "src-ip-prefix":"10.222.111.222/32"}]}' \
  https://tnsr.example.com/restconf/data/netgate-acl:acl-config/acl-table/acl-
  ↪ list=blockbadhosts/acl-rules/acl-rule=10
```

Output: This command has no output when it works successfully.

Retrieve the contents of the ACL again to see that the new rule is now present:

Command:

```
$ curl --cert ~/tnsr/tnsr-myuser.crt \
--key ~/tnsr/tnsr-myuser.key \
--cacert ~/tnsr/tnsr-selfca.crt \
-X GET \
https://tnsr.example.com/restconf/data/netgate-acl:acl-config/acl-table/acl-
↪list=blockbadhosts
```

Output:

```
{
  "netgate-acl:acl-list": [
    {
      "acl-name": "blockbadhosts",
      "acl-description": "Block bad hosts",
      "acl-rules": {
        "acl-rule": [
          {
            "sequence": 1,
            "action": "deny",
            "ip-version": "ipv4",
            "src-ip-prefix": "203.0.113.14/32"
          },
          {
            "sequence": 2,
            "action": "deny",
            "ip-version": "ipv4",
            "src-ip-prefix": "203.0.113.15/32"
          },
          {
            "sequence": 10,
            "action": "deny",
            "ip-version": "ipv4",
            "src-ip-prefix": "10.222.111.222/32"
          },
          {
            "sequence": 555,
            "action": "deny",
            "ip-version": "ipv4",
            "src-ip-prefix": "5.5.5.5/32"
          },
          {
            "sequence": 5000,
            "acl-rule-description": "Default Permit",
            "action": "permit",
            "ip-version": "ipv4"
          }
        ]
      }
    }
  ]
}
```

Use PATCH to update data

When using the PUT method, the client must supply all data in an entry to be replaced, even when only changing one small part. This makes it difficult to change, for example, the description of an ACL without sending the content of

the ACL back in the request.

The PATCH method allows individual values to be replaced without requiring all of the data to be sent. With PATCH, the client need only send the modified values in a query, along with enough information to uniquely identify the entry.

For example, to update the description of the `blockbadhosts` ACL using PATCH, the client must only include the name of the ACL and the new description. It does not need to include the entire content of the ACL and its rules as it would with a PUT request.

Item	Value
Request Type	-X PATCH (change content)
Content Type	-H "Content-Type: application/yang-data+json"
ACL Name	blockbadhosts
ACL Description	Block packets from bad hosts

The command is formatted in a similar manner to the PUT request in the previous example.

Warning: The Content-Type header must be set when performing a write operation such as PUT or PATCH. The value of the header must reflect the type of data being sent. These examples use JSON, so the header is set to `application/yang-data+json`. When submitting XML, it would be `application/yang-data+xml`.

Command:

```
$ curl --cert ~/tnsr/tnsr-myuser.crt \
  --key ~/tnsr/tnsr-myuser.key \
  --cacert ~/tnsr/tnsr-selfca.crt \
  -H "Content-Type: application/yang-data+json" \
  -X PATCH \
  -d '{"netgate-acl:acl-list":[{"acl-name": "blockbadhosts","acl-description": "Block
↪packets from bad hosts"}]}' \
  https://tnsr.example.com/restconf/data/netgate-acl:acl-config/acl-table/acl-
↪list=blockbadhosts/
```

Output: This command has no output when it works successfully.

Retrieve the contents of the ACL again to see that the new description is now present:

Command:

```
$ curl --cert ~/tnsr/tnsr-myuser.crt \
  --key ~/tnsr/tnsr-myuser.key \
  --cacert ~/tnsr/tnsr-selfca.crt \
  -X GET \
  https://tnsr.example.com/restconf/data/netgate-acl:acl-config/acl-table/acl-
↪list=blockbadhosts
```

Output:

```
{
  "netgate-acl:acl-list": [
    {
      "acl-name": "blockbadhosts",
      "acl-description": "Block packets from bad hosts",
      "acl-rules": {
        "acl-rule": [
          {
```

(continues on next page)

(continued from previous page)

```

        "sequence": 1,
        "action": "deny",
        "ip-version": "ipv4",
        "src-ip-prefix": "203.0.113.14/32"
    },
    {
        "sequence": 2,
        "action": "deny",
        "ip-version": "ipv4",
        "src-ip-prefix": "203.0.113.15/32"
    },
    {
        "sequence": 10,
        "action": "deny",
        "ip-version": "ipv4",
        "src-ip-prefix": "10.222.111.222/32"
    },
    {
        "sequence": 555,
        "action": "deny",
        "ip-version": "ipv4",
        "src-ip-prefix": "5.5.5.5/32"
    },
    {
        "sequence": 5000,
        "acl-rule-description": "Default Permit",
        "action": "permit",
        "ip-version": "ipv4"
    }
  ]
}
]
}
}
]
}

```

Remove a specific rule from an ACL

Say that entry is no longer needed and it is safe to remove. That can be done with a DELETE request for the URL corresponding to its sequence number:

Command:

```

$ curl --cert ~/tnsr/tnsr-myuser.crt \
  --key ~/tnsr/tnsr-myuser.key \
  --cacert ~/tnsr/tnsr-selfca.crt \
  -X DELETE \
  https://tnsr.example.com/restconf/data/netgate-acl:acl-config/acl-table/acl-
  ↪list=blockbadhosts/acl-rules/acl-rule=10

```

Output: This does not produce any output if it completed successfully.

Retrieve the contents of the ACL again to confirm it was removed.

27.1.6 Adding More Users

To create additional RESTCONF users, only two actions are required on TNSR: Generate a certificate for the new user, and then add the user to NACM. This example adds a new user named `anotheruser`.

Generate a new user certificate:

```
tnsr(config)# pki private-key anotheruser generate key-length 4096
tnsr(config)# pki signing-request set common-name anotheruser
tnsr(config)# pki signing-request set digest sha256
tnsr(config)# pki signing-request anotheruser generate
tnsr(config)# pki signing-request anotheruser sign ca-name selfca days-valid 365
↪ digest sha512 enable-ca false
```

Add this user to the NACM admin group:

```
tnsr(config)# nacm group admin
tnsr(config-nacm-group)# member anotheruser
tnsr(config-nacm-group)# exit
```

Then, the user certificate can be copied to a new client and used as explained previously.

27.2 TNSR IPsec Hub for pfSense

Current scenario:

HQ (hub) with 3 branch (spoke) sites, with secure interconnection between their local networks. One of the branch routers is assumed to be BGP capable. Internet access for one of the sites should be provided through the hub node.

Tip: This recipe does not contain configuration examples for IPsec cryptographic acceleration, which can greatly improve the efficiency and performance of IPsec tunnels. The availability of acceleration varies by hardware, so the specifics of acceleration configuration must be customized to the target environment.

For more information, see *IPsec Cryptographic Acceleration*

27.2.1 Input Data

The information in this section defines the local configuration which is covered in this recipe. These input values can be substituted by the actual corresponding values for a real-world implementation.

Scenario Topology

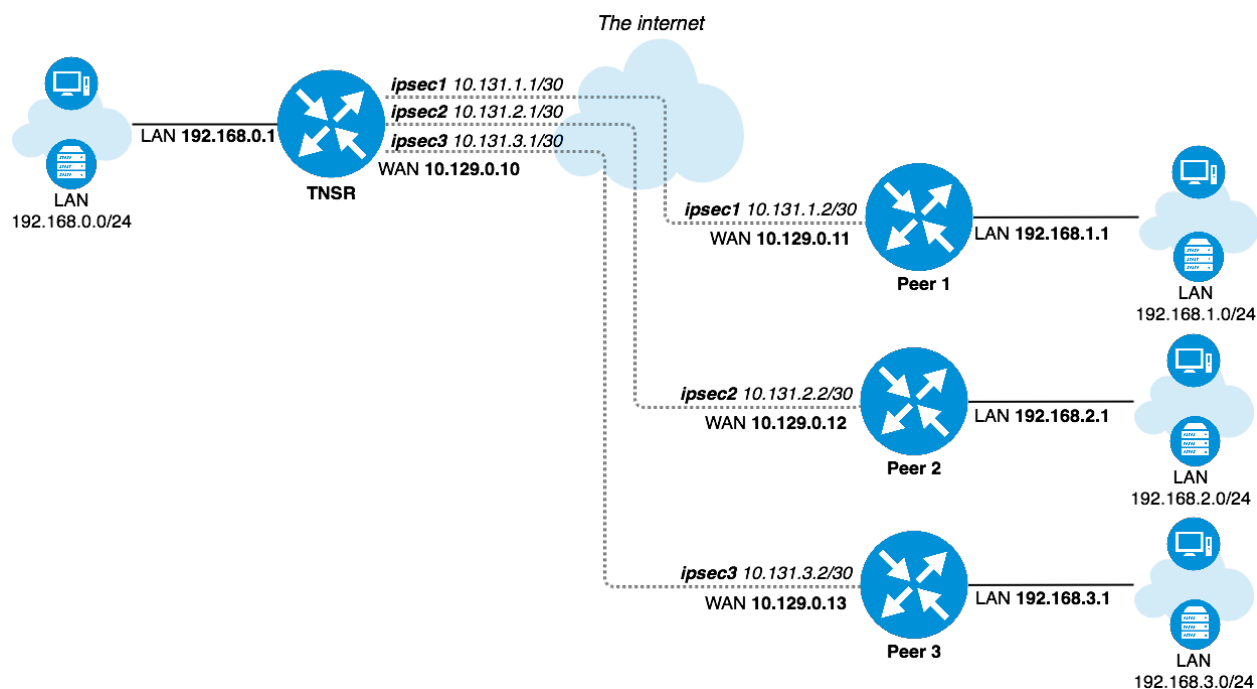


Fig. 1: TNSR IPsec Hub

TNSR and Peer Network Configuration

Table 1: TNSR Setup

Item	Value
LAN Interface	GigabitEthernet0/0
LAN Network	192.168.0.0/24
LAN IP Address static	192.168.0.1/24
WAN Interface	GigabitEthernet13/0/0
WAN IP Address DHCP	10.129.0.10/24
IPsec VTI Peer 1 IP Address	10.131.1.1/30
IPsec VTI Peer 2 IP Address	10.131.2.1/30
IPsec VTI Peer 3 IP Address	10.131.3.1/30

Table 2: Peer 1 Network Setup

Item	Value
LAN Interface	LAN
LAN Network	192.168.1.0/24
LAN IP Address static	192.168.1.1/24
WAN Interface	WAN
WAN IP Address DHCP	10.129.0.11/24
IPsec VTI TNSR IP Address	10.131.1.2/30

Table 3: Peer 2 Network Setup

Item	Value
LAN Interface	LAN
LAN Network	192.168.2.0/24
LAN IP Address static	192.168.2.1/24
WAN Interface	WAN
WAN IP Address DHCP	10.129.0.12/24
IPsec VTI TNSR IP Address	10.131.2.2/30

Table 4: Peer 3 Network Setup

Item	Value
LAN Interface	LAN
LAN Network	192.168.3.0/24
LAN IP Address static	192.168.3.1/24
WAN Interface	WAN
WAN IP Address DHCP	10.129.0.13/24
IPsec VTI TNSR IP Address	10.131.3.2/30

TNSR and Peer IPsec Configuration

General IPsec settings are the same for every node.

Table 5: IPsec IKE/Phase 1 Settings

Item	Value
Network Interface	WAN Interface
IKE type	IKEv2
Authentication method	PSK
Pre-Share Key	01234567
Local identifier	WAN IP Address
Remote identifier	Remote WAN IP Address
Encryption	AES-128-CBC
Hash	SHA1
DH group	14 (2048 bit modulus)
Lifetime	28800

Table 6: IPsec SA/Phase 2 Settings

Item	Value
Mode	Routed IPsec (VTI)
Protocol	ESP
Encryption	AES-128-CBC
Hash	SHA1
PFS group	14 (2048)
Lifetime	3600

27.2.2 Setup Details

Initial setup

It is assumed that devices have generic default setup, do not have any existing configuration errors, and are ready to be configured.

Note: In this scenario every device obtains its own static IP address on its WAN interface from an external lab gateway which is not a part of the considered scenario.

TNSR Setup

LAN settings

Setup LAN interface with static IP address:

```
tnsr tnsr# configure
tnsr tnsr(config)# interface GigabitEthernetb/0/0
tnsr tnsr(config-interface)# description LAN
tnsr tnsr(config-interface)# ip address 192.168.0.1/24
tnsr tnsr(config-interface)# enable
tnsr tnsr(config-interface)# exit
tnsr tnsr(config)# exit
```

WAN settings

Setup WAN interface for obtaining IP address via DHCP:

```
tnsr tnsr# configure
tnsr tnsr(config)# interface GigabitEthernet13/0/0
tnsr tnsr(config-interface)# description WAN
tnsr tnsr(config-interface)# dhcp client ipv4 hostname tnsr
tnsr tnsr(config-interface)# enable
tnsr tnsr(config-interface)# exit
tnsr tnsr(config)# exit
```

DHCP server

Setup DHCP server on LAN interface with following settings:

Table 7: TNSR DHCP Server Setup

Item	Value
DHCP IP address pool	192.168.0.100 to 192.168.0.199
Default gateway	TNSR LAN IP address
DNS	8.8.8.8 and 1.1.1.1

```
tnsr tnsr# configure
tnsr tnsr(config)# dhcp4 server
tnsr tnsr(config-kea-dhcp4)# description LAN DHCP
tnsr tnsr(config-kea-dhcp4)# interface listen GigabitEthernetb/0/0
```

(continues on next page)

(continued from previous page)

```
tnsr tnsr(config-kea-dhcp4)# subnet 192.168.0.0/24
tnsr tnsr(config-kea-subnet4)# interface GigabitEthernetb/0/0
tnsr tnsr(config-kea-subnet4)# pool 192.168.0.100-192.168.0.199
tnsr tnsr(config-kea-subnet4-pool)# exit
tnsr tnsr(config-kea-subnet4)# option routers
tnsr tnsr(config-kea-subnet4-opt)# data 192.168.0.1
tnsr tnsr(config-kea-subnet4-opt)# exit
tnsr tnsr(config-kea-subnet4)# option domain-name-servers
tnsr tnsr(config-kea-subnet4-opt)# data 8.8.8.8, 1.1.1.1
tnsr tnsr(config-kea-subnet4-opt)# exit
tnsr tnsr(config-kea-subnet4)# exit
tnsr tnsr(config-kea-dhcp4)# exit
tnsr tnsr(config)# dhcp4 enable
tnsr tnsr(config)# exit
```

NAT

```
tnsr tnsr# configure
tnsr tnsr(config)# nat global-options nat44 forwarding true
tnsr tnsr(config)# nat pool interface GigabitEthernet13/0/0
tnsr tnsr(config)# interface GigabitEthernetb/0/0
tnsr tnsr(config-interface)# ip nat inside
tnsr tnsr(config-interface)# exit
tnsr tnsr(config)# interface GigabitEthernet13/0/0
tnsr tnsr(config-interface)# ip nat outside
tnsr tnsr(config-interface)# exit
tnsr tnsr(config)# exit
```

Peer 1 Basic Setup

LAN settings

Setup LAN interface with static IP address.

- Navigate to **Interfaces > LAN**
- Set **IPv4 Configuration Type** to *Static IPv4*
- Set **IPv4 Address** to 192.168.1.1 and mask as 24
- Click **Save**
- Click **Apply Changes**

WAN settings

Setup WAN interface for obtaining an IP address via DHCP. This could also be a static setup, following a similar form to the LAN settings above.

- Navigate to **Interfaces > WAN**
- Set **IPv4 Configuration Type** to *DHCP*
- Click **Save**

- Click **Apply Changes**

DHCP server

Setup DHCP server on LAN interface with following settings:

Table 8: Peer 1 DHCP Server Setup

Item	Value
DHCP IP address pool	192.168.1.100 to 192.168.1.199
Default gateway	LAN IP address (pfSense Default)
DNS	LAN IP address (pfSense Default)

- Navigate to **Services > DHCP Server, LAN** tab
- Set **Range From** as 192.168.1.100 and **To** as 192.168.1.199
- Click **Save**

Peer 2 Basic Setup

LAN settings

Setup LAN interface with static IP address.

- Navigate to **Interfaces > LAN**
- Set **IPv4 Configuration Type** to *Static IPv4*
- Set **IPv4 Address** to 192.168.2.1 and mask as 24
- Click **Save**
- Click **Apply Changes**

WAN settings

Setup WAN interface for obtaining an IP address via DHCP. This could also be a static setup, following a similar form to the LAN settings above.

- Navigate to **Interfaces > WAN**
- Set **IPv4 Configuration Type** to *DHCP*
- Click **Save**
- Click **Apply Changes**

DHCP server

Setup DHCP server on LAN interface with following settings:

Table 9: Peer 2 DHCP Server Setup

Item	Value
DHCP IP address pool	192.168.2.100 to 192.168.2.199
Default gateway	LAN IP address (pfSense Default)
DNS	LAN IP address (pfSense Default)

- Navigate to **Services > DHCP Server, LAN** tab
- Set **Range From** as 192.168.2.100 and **To** as 192.168.2.199
- Click **Save**

Peer 3 Basic Setup

LAN settings

Setup LAN interface with static IP address.

- Navigate to **Interfaces > LAN**
- Set **IPv4 Configuration Type** to *Static IPv4*
- Set **IPv4 Address** to 192.168.3.1 and mask as 24
- Click **Save**
- Click **Apply Changes**

WAN settings

Setup WAN interface for obtaining an IP address via DHCP. This could also be a static setup, following a similar form to the LAN settings above.

- Navigate to **Interfaces > WAN**
- Set **IPv4 Configuration Type** to *DHCP*
- Click **Save**
- Click **Apply Changes**

DHCP server

Setup DHCP server on LAN interface with following settings:

Table 10: Peer 3 DHCP Server Setup

Item	Value
DHCP IP address pool	192.168.3.100 to 192.168.3.199
Default gateway	LAN IP address (pfSense Default)
DNS	LAN IP address (pfSense Default)

- Navigate to **Services > DHCP Server, LAN** tab
- Set **Range From** as 192.168.3.100 and **To** as 192.168.3.199
- Click **Save**

27.2.3 Access between local and remote networks via IPsec

This section describes minimal IPsec and routing settings in order to obtain secure interconnectivity between LAN networks for every device.

This document assumes that devices have generic initial setup successfully completed and are able to reach each other via WAN network.

TNSR

IPsec Configuration

IPsec setup for each pfSense node

IPsec to Peer 1

Enter config state:

```
tnsr tnsr# configure
```

Creating IPsec instance with id 1:

```
tnsr tnsr(config)# ipsec tunnel 1
tnsr tnsr(config-ipsec-tunnel)# local-address 10.129.0.10
tnsr tnsr(config-ipsec-tunnel)# remote-address 10.129.0.11
tnsr tnsr(config-ipsec-tunnel)# crypto config-type ike
```

P1 encryption settings:

```
tnsr tnsr(config-ipsec-tunnel)# crypto ike
tnsr tnsr(config-ipsec-crypto-ike)# version 2
tnsr tnsr(config-ipsec-crypto-ike)# lifetime 28800
tnsr tnsr(config-ipsec-crypto-ike)# proposal 1
tnsr tnsr(config-ike-proposal)# encryption aes128
tnsr tnsr(config-ike-proposal)# integrity sha1
tnsr tnsr(config-ike-proposal)# group modp2048
tnsr tnsr(config-ike-proposal)# exit
```

Creating peer IDs:

```
tnsr tnsr(config-ipsec-crypto-ike)# identity local
tnsr tnsr(config-ike-identity)# type address
tnsr tnsr(config-ike-identity)# value 10.129.0.10
tnsr tnsr(config-ike-identity)# exit
tnsr tnsr(config-ipsec-crypto-ike)# identity remote
tnsr tnsr(config-ike-identity)# type address
tnsr tnsr(config-ike-identity)# value 10.129.0.11
tnsr tnsr(config-ike-identity)# exit
```

Authentication:

```
tnsr tnsr(config-ipsec-crypto-ike)# authentication local
tnsr tnsr(config-ike-authentication)# round 1
tnsr tnsr(config-ike-authentication-round)# type psk
tnsr tnsr(config-ike-authentication-round)# psk 01234567
```

(continues on next page)

(continued from previous page)

```
tnsr tnsr(config-ike-authentication-round)# exit
tnsr tnsr(config-ike-authentication)# exit
tnsr tnsr(config-ipsec-crypto-ike)# authentication remote
tnsr tnsr(config-ike-authentication)# round 1
tnsr tnsr(config-ike-authentication-round)# type psk
tnsr tnsr(config-ike-authentication-round)# psk 01234567
tnsr tnsr(config-ike-authentication-round)# exit
tnsr tnsr(config-ike-authentication)# exit
```

P2 settings:

```
tnsr tnsr(config-ipsec-crypto-ike)# child 1
tnsr tnsr(config-ike-child)# lifetime 3600
tnsr tnsr(config-ike-child)# proposal 1
tnsr tnsr(config-ike-child-proposal)# encryption aes128
tnsr tnsr(config-ike-child-proposal)# integrity sha1
tnsr tnsr(config-ike-child-proposal)# group modp2048
tnsr tnsr(config-ike-child-proposal)# exit
tnsr tnsr(config-ike-child)# exit
tnsr tnsr(config-ipsec-crypto-ike)# exit
tnsr tnsr(config-ipsec-tunnel)# exit
```

Configuring tunnel interface

```
tnsr tnsr(config)# interface ipip1
tnsr tnsr(config-interface)# ip address 10.131.1.1/30
tnsr tnsr(config-interface)# exit
tnsr tnsr(config)# exit
```

IPsec to Peer 2

Enter config state:

```
tnsr tnsr# configure
```

Creating IPsec instance with id 2:

```
tnsr tnsr(config)# ipsec tunnel 2
tnsr tnsr(config-ipsec-tunnel)# local-address 10.129.0.10
tnsr tnsr(config-ipsec-tunnel)# remote-address 10.129.0.12
tnsr tnsr(config-ipsec-tunnel)# crypto config-type ike
```

P1 encryption settings:

```
tnsr tnsr(config-ipsec-tunnel)# crypto ike
tnsr tnsr(config-ipsec-crypto-ike)# version 2
tnsr tnsr(config-ipsec-crypto-ike)# lifetime 28800
tnsr tnsr(config-ipsec-crypto-ike)# proposal 1
tnsr tnsr(config-ike-proposal)# encryption aes128
tnsr tnsr(config-ike-proposal)# integrity sha1
tnsr tnsr(config-ike-proposal)# group modp2048
tnsr tnsr(config-ike-proposal)# exit
```

Creating peer ID's:

```
tnsr tnsr(config-ipsec-crypto-ike)# identity local
tnsr tnsr(config-ike-identity)# type address
tnsr tnsr(config-ike-identity)# value 10.129.0.10
tnsr tnsr(config-ike-identity)# exit
tnsr tnsr(config-ipsec-crypto-ike)# identity remote
tnsr tnsr(config-ike-identity)# type address
tnsr tnsr(config-ike-identity)# value 10.129.0.12
tnsr tnsr(config-ike-identity)# exit
```

Authentication:

```
tnsr tnsr(config-ipsec-crypto-ike)# authentication local
tnsr tnsr(config-ike-authentication)# round 1
tnsr tnsr(config-ike-authentication-round)# type psk
tnsr tnsr(config-ike-authentication-round)# psk 01234567
tnsr tnsr(config-ike-authentication-round)# exit
tnsr tnsr(config-ike-authentication)# exit
tnsr tnsr(config-ipsec-crypto-ike)# authentication remote
tnsr tnsr(config-ike-authentication)# round 1
tnsr tnsr(config-ike-authentication-round)# type psk
tnsr tnsr(config-ike-authentication-round)# psk 01234567
tnsr tnsr(config-ike-authentication-round)# exit
tnsr tnsr(config-ike-authentication)# exit
```

P2 settings:

```
tnsr tnsr(config-ipsec-crypto-ike)# child 1
tnsr tnsr(config-ike-child)# lifetime 3600
tnsr tnsr(config-ike-child)# proposal 1
tnsr tnsr(config-ike-child-proposal)# encryption aes128
tnsr tnsr(config-ike-child-proposal)# integrity sha1
tnsr tnsr(config-ike-child-proposal)# group modp2048
tnsr tnsr(config-ike-child-proposal)# exit
tnsr tnsr(config-ike-child)# exit
tnsr tnsr(config-ipsec-crypto-ike)# exit
tnsr tnsr(config-ipsec-tunnel)# exit
```

Configuring tunnel interface:

```
tnsr tnsr(config)# interface ipip2
tnsr tnsr(config-interface)# ip address 10.131.2.1/30
tnsr tnsr(config-interface)# exit
tnsr tnsr(config)# exit
```

IPsec to Peer 3

Enter config state:

```
tnsr tnsr# configure
```

Creating IPsec instance with id 3:

```
tnsr tnsr(config)# ipsec tunnel 3
tnsr tnsr(config-ipsec-tunnel)# local-address 10.129.0.10
tnsr tnsr(config-ipsec-tunnel)# remote-address 10.129.0.13
tnsr tnsr(config-ipsec-tunnel)# crypto config-type ike
```

P1 encryption settings:

```
tnsr tnsr(config-ipsec-tunnel)# crypto ike
tnsr tnsr(config-ipsec-crypto-ike)# version 2
tnsr tnsr(config-ipsec-crypto-ike)# lifetime 28800
tnsr tnsr(config-ipsec-crypto-ike)# proposal 1
tnsr tnsr(config-ike-proposal)# encryption aes128
tnsr tnsr(config-ike-proposal)# integrity sha1
tnsr tnsr(config-ike-proposal)# group modp2048
tnsr tnsr(config-ike-proposal)# exit
```

Creating peer ID's:

```
tnsr tnsr(config-ipsec-crypto-ike)# identity local
tnsr tnsr(config-ike-identity)# type address
tnsr tnsr(config-ike-identity)# value 10.129.0.10
tnsr tnsr(config-ike-identity)# exit
tnsr tnsr(config-ipsec-crypto-ike)# identity remote
tnsr tnsr(config-ike-identity)# type address
tnsr tnsr(config-ike-identity)# value 10.129.0.13
tnsr tnsr(config-ike-identity)# exit
```

Authentication:

```
tnsr tnsr(config-ipsec-crypto-ike)# authentication local
tnsr tnsr(config-ike-authentication)# round 1
tnsr tnsr(config-ike-authentication-round)# type psk
tnsr tnsr(config-ike-authentication-round)# psk 01234567
tnsr tnsr(config-ike-authentication-round)# exit
tnsr tnsr(config-ike-authentication)# exit
tnsr tnsr(config-ipsec-crypto-ike)# authentication remote
tnsr tnsr(config-ike-authentication)# round 1
tnsr tnsr(config-ike-authentication-round)# type psk
tnsr tnsr(config-ike-authentication-round)# psk 01234567
tnsr tnsr(config-ike-authentication-round)# exit
tnsr tnsr(config-ike-authentication)# exit
```

P2 settings:

```
tnsr tnsr(config-ipsec-crypto-ike)# child 1
tnsr tnsr(config-ike-child)# lifetime 3600
tnsr tnsr(config-ike-child)# proposal 1
tnsr tnsr(config-ike-child-proposal)# encryption aes128
tnsr tnsr(config-ike-child-proposal)# integrity sha1
tnsr tnsr(config-ike-child-proposal)# group modp2048
tnsr tnsr(config-ike-child-proposal)# exit
tnsr tnsr(config-ike-child)# exit
tnsr tnsr(config-ipsec-crypto-ike)# exit
tnsr tnsr(config-ipsec-tunnel)# exit
```

Configuring tunnel interface:

```
tnsr tnsr(config)# interface ipip3
tnsr tnsr(config-interface)# ip address 10.131.3.1/30
tnsr tnsr(config-interface)# exit
tnsr tnsr(config)# exit
```

Routing

This section describes routing setup. This scenario assumes one of the pfSense IPsec peers, Peer 1, uses a dynamic routing protocol (BGP) and the remaining two IPsec peers use static routing.

Peer 1 BGP Routing

Enter config state:

```
tnsr tnsr# configure
```

Defining redistributed networks, peer 2 and 3:

```
tnsr tnsr(config)# prefix-list VPN-ROUTES
tnsr tnsr(config-prefix-list)# sequence 1 permit 192.168.2.0/23 le 24
tnsr tnsr(config-prefix-list)# exit
tnsr tnsr(config)# route-map VPN-ROUTES-MAP permit sequence 1
tnsr tnsr(config-route-map)# match ip address prefix-list VPN-ROUTES
tnsr tnsr(config-route-map)# exit
```

Setup BGP instance:

```
tnsr tnsr(config)# route dynamic bgp
tnsr tnsr(config-frr-bgp)# server 65000
tnsr tnsr(config-bgp)# router-id 192.168.0.1
```

Defining neighbor:

```
tnsr tnsr(config-bgp)# neighbor 10.131.1.2
tnsr tnsr(config-bgp-neighbor)# remote-as 65001
tnsr tnsr(config-bgp-neighbor)# enable
tnsr tnsr(config-bgp-neighbor)# exit
```

Setup peer in certain address-family space:

```
tnsr tnsr(config-bgp)# address-family ipv4 unicast
tnsr tnsr(config-bgp-ip4uni)# neighbor 10.131.1.2
tnsr tnsr(config-bgp-ip4uni-nbr)# activate
tnsr tnsr(config-bgp-ip4uni-nbr)# exit
```

Defining local network in certain address-family space:

```
tnsr tnsr(config-bgp-ip4uni)# network 192.168.0.0/24
```

Defining redistributed networks

```
tnsr tnsr(config-bgp-ip4uni)# redistribute kernel route-map VPN-ROUTES-MAP
tnsr tnsr(config-bgp-ip4uni)# exit
tnsr tnsr(config-bgp)# exit
```

Enabling BGP if one is not enabled:

```
tnsr tnsr(config-frr-bgp)# enable
tnsr tnsr(config-frr-bgp)# exit
```

Better to restart service in order to be sure changes applied effectively:

```
tnsr tnsr(config)# service bgp restart
tnsr tnsr(config)# exit
```

Peer 2 Static Routing

```
tnsr tnsr# configure
tnsr tnsr(config)# route ipv4 table ipv4-VRF:0
tnsr tnsr(config-route-table-v4)# route 192.168.2.0/24
tnsr tnsr(config-rttbl4-next-hop)# next-hop 0 via 10.131.2.2 ipip2
tnsr tnsr(config-rttbl4-next-hop)# exit
tnsr tnsr(config-route-table-v4)# exit
tnsr tnsr(config)# exit
```

Peer 3 Static Routing

```
tnsr tnsr# configure
tnsr tnsr(config)# route ipv4 table ipv4-VRF:0
tnsr tnsr(config-route-table-v4)# route 192.168.3.0/24
tnsr tnsr(config-rttbl4-next-hop)# next-hop 0 via 10.131.3.2 ipip3
tnsr tnsr(config-rttbl4-next-hop)# exit
tnsr tnsr(config-route-table-v4)# exit
tnsr tnsr(config)# exit
```

Peer 1 Setup

IPsec Settings

Phase 1

- Navigate to **VPN > IPsec**
- Click **Add P1**
- Set **Key Exchange version** to *IKEv2*
- Set **Internet Protocol** to *IPv4*
- Set **Interface** to *WAN*
- Set **Remote Gateway** to *10.129.0.10*
- Set **Authentication Method** to *Mutual PSK*
- Set **My identifier** to *My IP address*
- Set **Peer identifier** to *Peer IP address*
- Set **Pre-Shared Key** to *01234567*
- Set **Encryption**:
 - **Algorithm** to *AES*
 - **Key length** to *128 bit*

- **Hash** to *SHA1*
 - **DH Group** to *14 (2048 bit)*
- Set **Lifetime** as 28800
- Click **Save**

Phase 2

- On the newly created Phase 1 entry, click **Show Phase 2 Entries**
- Click **Add P2**
- Set **Mode** to *Routed (VTI)*
- Set **Local Network** to 10.131.2.2 and mask 30
- Set **Remote Network** to 10.131.2.1
- Set **Protocol** to *ESP*
- Set **Encryption Algorithms** to *AES* and *128 bit*
- Uncheck all other **Encryption Algorithms** entries
- Set **Hash Algorithms** to *SHA1*
- Uncheck all other **Hash Algorithms** entries
- Set **PFS key group** to *14 (2048 bit)*
- Set **Lifetime** as 3600
- Click **Save**
- Click **Apply Changes**

Interface

- Navigate to **Interfaces > Interface Assignments**
- From the **Available network ports** list, choose *ipsecNNNN (IPsec VTI)* (The ID number will vary)
- Click **Add**
- Note the newly created interface name, such as OPTX
- Navigate to **Interfaces > OPTX**
- Check **Enable**
- Click **Save**
- Click **Apply Changes**

Routing

- Navigate to **System > Package Manager** and install the FRR package
- Browse to **Services > FRR Global/Zebra**
- Check **Enable FRR**

- Set **Master Password** to any value

Note: This is a requirement for the zebra management daemon to run, this password is not used by clients.

- Check **Enable logging**
- Set **Router ID** to 192.168.1.1
In this case, it is the LAN interface IP address, assuming it will be always be available for routing between LAN subnets.
- Click **Save**
- Navigate to the **[BGP]** tab
- Check **Enable BGP Routing**
- Check **Log Adjacency Changes**
- Set **Local AS** to 65001
- Set **Router ID** to 192.168.1.1
- Set **Networks to Distribute** to 192.168.1.0/24
- Navigate to the **Neighbors** tab
- Click **Add**
- Set **Name/Address** to 10.131.1.1 (TNSR VTI interface IP address)
- Set **Remote AS** to 65000
- Click **Save**

At this point, routes to 192.168.0.0/24, 192.168.2.0/24, and 192.168.3.0/24 will be learned by BGP and installed in the routing table. If it is not so, check **Status > FRR** on the **BGP** tab. That page contains useful BGP troubleshooting information. Additionally, check the routing log at **Status > System Logs** on the **Routing** tab under **System**.

Firewall

To allow connections into the local LAN from remote IPsec sites, create necessary pass rules under **Firewall > Rules** on the **IPsec** tab. These rules would have a **Source** set to the remote LAN or whichever network is the source of the traffic to allow.

For simplicity, this example has a rule to pass IPv4 traffic from any source to any destination since the only IPsec interface traffic will be from 192.168.0.0/22.

NAT

TNSR will perform NAT for this peer, so outbound NAT is not necessary. It may be left at the default, which will not touch IPsec traffic, or outbound NAT may be disabled entirely which will also prevent LAN subnet traffic from exiting out the WAN unintentionally.

Peer 2 Setup

IPsec Settings

Phase 1

- Navigate to **VPN > IPsec**
- Click **Add P1**
- Set **Key Exchange version** to *IKEv2*
- Set **Internet Protocol** to *IPv4*
- Set **Interface** to *WAN*
- Set **Remote Gateway** to *10.129.0.10*
- Set **Authentication Method** to *Mutual PSK*
- Set **My identifier** to *My IP address*
- Set **Peer identifier** to *Peer IP address*
- Set **Pre-Shared Key** to *01234567*
- Set **Encryption**:
 - **Algorithm** to *AES*
 - **Key length** to *128 bit*
 - **Hash** to *SHA1*
 - **DH Group** to *14 (2048 bit)*
- Set **Lifetime** as *28800*
- Click **Save**

Phase 2

- On the newly created Phase 1 entry, click **Show Phase 2 Entries**
- Click **Add P2**
- Set **Mode** to *Routed (VTI)*
- Set **Local Network** to *10.131.3.2* and mask *30*
- Set **Remote Network** to *10.131.3.1*
- Set **Protocol** to *ESP*
- Set **Encryption Algorithms** to *AES* and *128 bit*
- Uncheck all other **Encryption Algorithms** entries
- Set **Hash Algorithms** to *SHA1*
- Uncheck all other **Hash Algorithms** entries
- Set **PFS key group** to *14 (2048 bit)*
- Set **Lifetime** as *3600*

- Click **Save**
- Click **Apply Changes**

Interface

- Navigate to **Interfaces > Interface Assignments**
- From the **Available network ports** list, choose *ipsecNNNN (IPsec VTI)* (The ID number will vary)
- Click **Add**
- Note the newly created interface name, such as OPTX
- Navigate to **Interfaces > OPTX**
- Check **Enable**
- Click **Save**
- Click **Apply Changes**

Routing

- Navigate to **System > Routing, Static Routes** tab
- Click **Add**
- Set **Destination network** to 192.168.0.0 and mask 23
- Set **Gateway** to the newly created VTI interface gateway, which has an address of 10.131.2.1
- Click **Save**
- Click **Add**
- Set **Destination network** to 192.168.3.0 and mask 24
- Set **Gateway** to the newly created VTI interface gateway, which has an address of 10.131.2.1
- Click **Save**
- Click **Apply Changes**

Firewall

To allow connections into the local LAN from remote IPsec sites, create necessary pass rules under **Firewall > Rules** on the **IPsec** tab. These rules would have a **Source** set to the remote LAN or whichever network is the source of the traffic to allow.

For simplicity, this example has a rule to pass IPv4 traffic from any source to any destination since the only IPsec interface traffic will be from 192.168.0.0/22.

NAT

TNSR will perform NAT for this peer, so outbound NAT is not necessary. It may be left at the default, which will not touch IPsec traffic, or outbound NAT may be disabled entirely which will also prevent LAN subnet traffic from exiting out the WAN unintentionally.

Peer 3 Setup

IPsec Settings

Phase 1

- Navigate to **VPN > IPsec**
- Click **Add P1**
- Set **Key Exchange version** to *IKEv2*
- Set **Internet Protocol** to *IPv4*
- Set **Interface** to *WAN*
- Set **Remote Gateway** to *10.129.0.10*
- Set **Authentication Method** to *Mutual PSK*
- Set **My identifier** to *My IP address*
- Set **Peer identifier** to *Peer IP address*
- Set **Pre-Shared Key** to *01234567*
- Set **Encryption**:
 - **Algorithm** to *AES*
 - **Key length** to *128 bit*
 - **Hash** to *SHA1*
 - **DH Group** to *14 (2048 bit)*
- Set **Lifetime** as *28800*
- Click **Save**

Phase 2

- On the newly created Phase 1 entry, click **Show Phase 2 Entries**
- Click **Add P2**
- Set **Mode** to *Routed (VTI)*
- Set **Local Network** to *10.131.4.2* and mask *30*
- Set **Remote Network** to *10.131.4.1*
- Set **Protocol** to *ESP*
- Set **Encryption Algorithms** to *AES* and *128 bit*
- Uncheck all other **Encryption Algorithms** entries
- Set **Hash Algorithms** to *SHA1*
- Uncheck all other **Hash Algorithms** entries
- Set **PFS key group** to *14 (2048 bit)*
- Set **Lifetime** as *3600*

- Click **Save**
- Click **Apply Changes**

Interface

- Navigate to **Interfaces > Interface Assignments**
- From the **Available network ports** list, choose *ipsecNNNN (IPsec VTI)* (The ID number will vary)
- Click **Add**
- Note the newly created interface name, such as OPTX
- Navigate to **Interfaces > OPTX**
- Check **Enable**
- Click **Save**
- Click **Apply Changes**

Routing

- Navigate to **System > Routing, Static Routes** tab
- Click **Add**
- Set **Destination network** to 192.168.0.0 and mask 23
- Set **Gateway** to the newly created VTI interface gateway, which has an address of 10.131.3.1
- Click **Save**
- Click **Add**
- Set **Destination network** to 192.168.2.0 and mask 24
- Set **Gateway** to the newly created VTI interface gateway, which has an address of 10.131.3.1
- Click **Save**
- Click **Apply Changes**

Firewall

To allow connections into the local LAN from remote IPsec sites, create necessary pass rules under **Firewall > Rules** on the **IPsec** tab. These rules would have a **Source** set to the remote LAN or whichever network is the source of the traffic to allow.

For simplicity, this example has a rule to pass IPv4 traffic from any source to any destination since the only IPsec interface traffic will be from 192.168.0.0/22.

NAT

TNSR will perform NAT for this peer, so outbound NAT is not necessary. It may be left at the default, which will not touch IPsec traffic, or outbound NAT may be disabled entirely which will also prevent LAN subnet traffic from exiting out the WAN unintentionally.

Access to the internet for remote network

This section describes minimal routing and NAT settings which provide access to the Internet for one of the remote networks. In current case this is Peer 1 that exchanges routing information with TNSR via BGP.

This document assumes that devices have IPsec setup successfully completed, able to reach each other via IPsec tunnel using path information from the dynamic routing protocol.

TNSR

NAT/PAT

Setup NAT for remote network, in this case PAT is used.

Note: Defining NAT inside interface for internet traffic sourced from Peer 1. Outside interface and PAT were defined earlier.

```
tnsr tnsr# configure
tnsr tnsr(config)# interface ipipl
tnsr tnsr(config-interface)# ip nat inside
tnsr tnsr(config-interface)# exit
```

Peer 1 Policy Route

Routing

Setup access to the internet via IPsec VTI interface with a policy-based routing rule.

- Navigate to **Firewall > Rules**
- Create (or modify existing default pass ipv4 LAN any) rule:
 - Set **Address Family** to **IPv4**
 - Set **Protocol** to **ANY**
 - Set **Source** to **LAN net**
 - Set **Destination** to **ANY**
 - Click **Display Advanced**
 - Set **Gateway** to **<IPsec interface name>_VTIV4**
 - Click **Save**

Note: VTI on pfSense does not support `reply-to`. Despite this policy routing rule on Peer1 which covers all traffic, there must also be kernel routes to remote LANs for the return traffic to find the way back.

27.3 Edge Router Speaking eBGP with Static Redistribution for IPv4 And IPv6

27.3.1 Use Case

Especially in cases where an enterprise is multi-homed with its own block of network addresses, it may become necessary to configure dynamic routing between network service providers. This is accomplished by use of external BGP (eBGP).

In this use case, the enterprise will use TNSR to speak eBGP with two network service providers, in order to exchange routes which may be redistributed from static/connected routing.

27.3.2 Example Scenario

In this example, the enterprise using TNSR will have a fictitious autonomous system number (ASN) of 65505. The network service providers in this example will have ASNs of 65510 and 65520. The enterprise using TNSR will redistribute a single /24 network from static into BGP. That network will then be advertised to each of the service providers. The service providers will announce a full routing table to the TNSR instance.

Scenario Topology

Table 11: BGP Router Setup Parameters

Item	Value
TNSR Autonomous System Number	65505
ISP_A Autonomous System Number	65510
ISP_B Autonomous System Number	65520
IPv4 Network to be announced	192.0.2.0/24
IPv6 Network to be announced	2001:db8:a100:1005::/64
TNSR to ISP_A IPv4 Network Address	203.0.113.8/30
TNSR to ISP_A IPv6 Global Address	2001:db8:fa00:ffaa::/64
TNSR to ISP_B IPv4 Network Address	100.64.0.48/30
TNSR to ISP_B IPv6 Global Address	2001:db8:fb00:ffbb::/64

27.3.3 TNSR Configuration Steps

Step 1: Configure Interfaces

```
tnsr# conf
tnsr(config)# interface GigabitEthernet0/13/0
tnsr(config-interface)# description "To ISP A"
tnsr(config-interface)# ip address 203.0.113.9/30
tnsr(config-interface)# ipv6 address 2001:db8:1000:aaaa::2/64
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)#
tnsr(config)# interface GigabitEthernet0/14/0
tnsr(config-interface)# description "To ISP B"
tnsr(config-interface)# ip address 100.64.0.49/30
tnsr(config-interface)# ipv6 address 2001:db8:9999:ffff::2/64
```

(continues on next page)

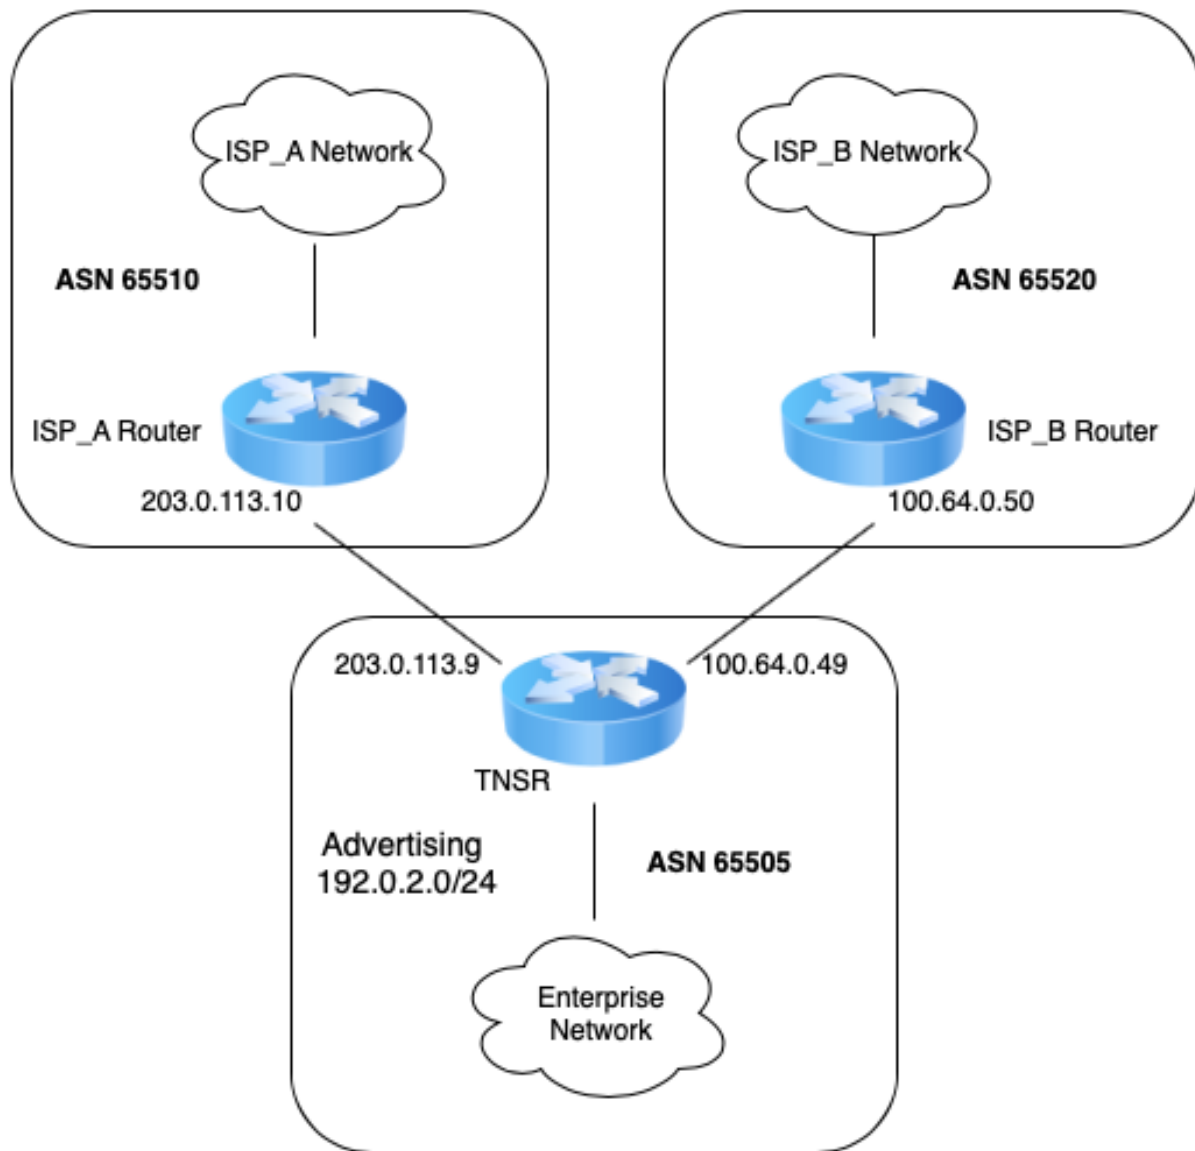
Example: IPv4

Fig. 2: TNSR BGP Router (IPv4)

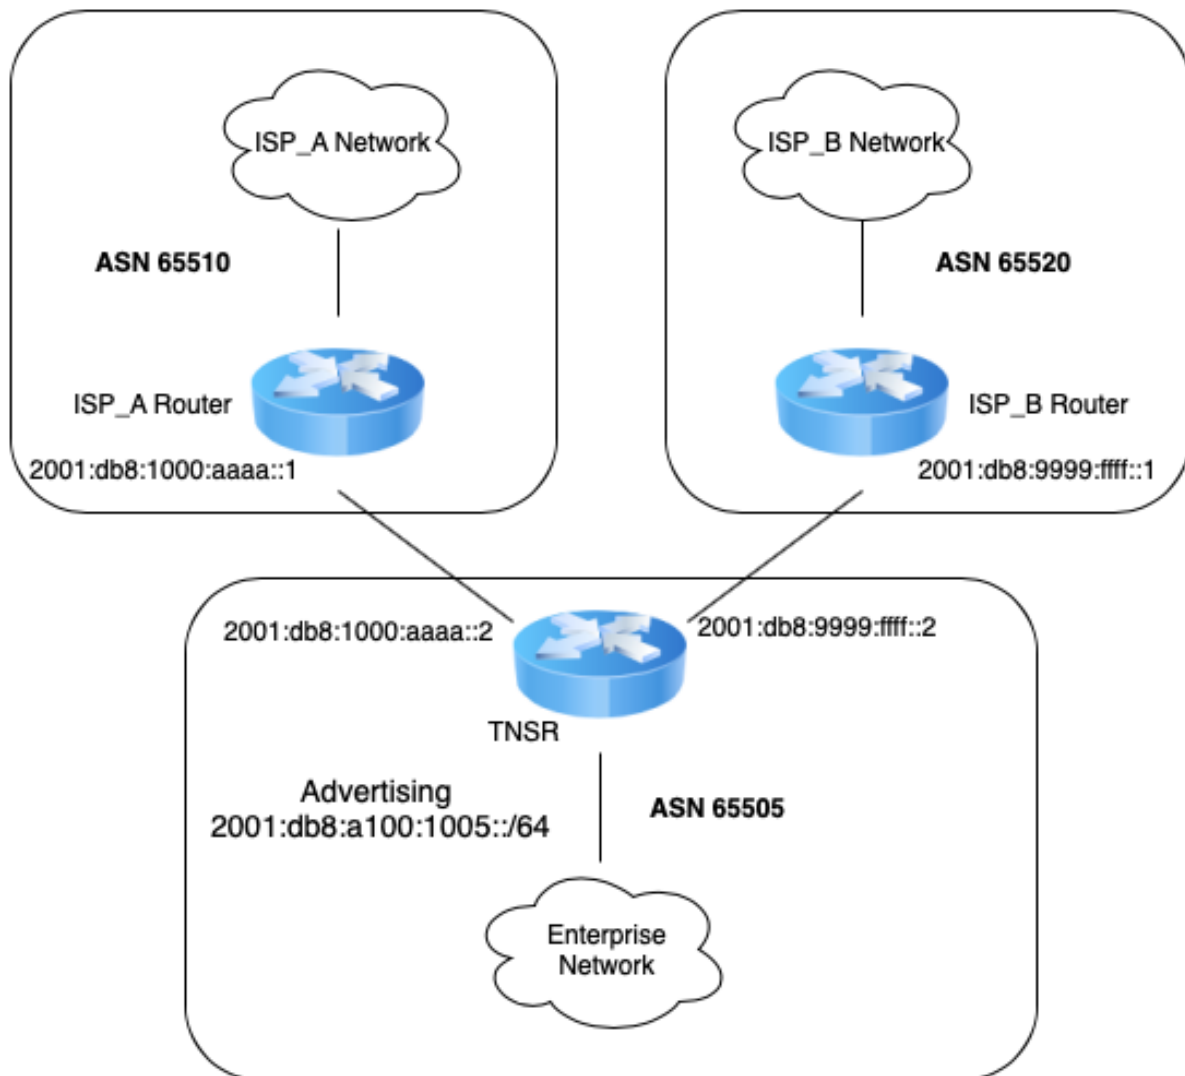
Example: IPv6

Fig. 3: TNSR BGP Router (IPv6)

(continued from previous page)

```
tnsr(config-interface)# enable
tnsr(config-interface)# exit
tnsr(config)#
```

Step 2: Enable BGP

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# enable
tnsr(config-frr-bgp)# exit
tnsr(config)#
```

Step 3: Create prefix-lists for route export via BGP

```
tnsr(config)# route dynamic prefix-list EXPORT_IPv4
tnsr(config-prefix-list)# description "IPv4 Routes to Export"
tnsr(config-prefix-list)# seq 10 permit 192.0.2.0/24
tnsr(config-prefix-list)# exit
tnsr(config)#
tnsr(config)# route dynamic prefix-list EXPORT_IPv6
tnsr(config-prefix-list)# description "IPv6 Routes to Export"
tnsr(config-prefix-list)# seq 10 permit 2001:db8:a100:1005::/64
tnsr(config-prefix-list)# exit
tnsr(config)#
```

Step 4: Create static route for networks to be advertised in BGP

```
tnsr(config)# route ipv4 table ipv4-VRF:0
tnsr(config-route-table-v4)# route 192.0.2.0/24
tnsr(config-rttbl4-next-hop)# next-hop 1 via local
tnsr(config-rttbl4-next-hop)# exit
tnsr(config-route-table-v4)# exit
```

```
tnsr(config)# route ipv6 table ipv6-VRF:0
tnsr(config-route-table-v6)# route 2001:db8:a100:1005::/64
tnsr(config-rttbl6-next-hop)# next-hop 1 via local
tnsr(config-rttbl6-next-hop)# exit
tnsr(config-route-table-v6)# exit
tnsr(config)#
```

Step 5: Configure BGP global options

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)# server 65505
tnsr(config-bgp)# router-id 203.0.113.9
tnsr(config-bgp)# address-family ipv4 unicast
tnsr(config-bgp-ip4uni)# redistribute kernel
tnsr(config-bgp-ip4uni)# exit
tnsr(config-bgp)# address-family ipv6 unicast
tnsr(config-bgp-ip4uni)# redistribute kernel
```

(continues on next page)

(continued from previous page)

```
tnsr(config-bgp-ip4uni)# exit
tnsr(config-bgp)#
```

Step 6: Configure BGP global neighbor options

```
tnsr(config-bgp)# neighbor 203.0.113.10
tnsr(config-bgp-neighbor)# remote-as 65510
tnsr(config-bgp-neighbor)# description "ISP_A IPv4"
tnsr(config-bgp-neighbor)# enable
tnsr(config-bgp-neighbor)# exit
```

```
tnsr(config-bgp)# neighbor 2001:db8:1000:aaaa::1
tnsr(config-bgp-neighbor)# remote-as 65510
tnsr(config-bgp-neighbor)# description "ISP_A IPv6"
tnsr(config-bgp-neighbor)# enable
tnsr(config-bgp-neighbor)# exit
```

```
tnsr(config-bgp)# neighbor 100.64.0.50
tnsr(config-bgp-neighbor)# remote-as 65520
tnsr(config-bgp-neighbor)# description "ISP_B IPv4"
tnsr(config-bgp-neighbor)# enable
tnsr(config-bgp-neighbor)# exit
```

```
tnsr(config-bgp)# neighbor 2001:db8:9999:ffff::1
tnsr(config-bgp-neighbor)# remote-as 65520
tnsr(config-bgp-neighbor)# description "ISP_B IPv6"
tnsr(config-bgp-neighbor)# enable
tnsr(config-bgp-neighbor)# exit
tnsr(config-bgp)#
```

Step 7: Configure BGP neighbor address-family IPv4 unicast options

```
tnsr(config-bgp)# address-family ipv4 unicast
tnsr(config-bgp-ip4uni)# neighbor 203.0.113.10
tnsr(config-bgp-ip4uni-nbr)# prefix-list EXPORT_IPv4 out
tnsr(config-bgp-ip4uni-nbr)# activate
tnsr(config-bgp-ip4uni-nbr)# exit
tnsr(config-bgp-ip4uni)# neighbor 100.64.0.50
tnsr(config-bgp-ip4uni-nbr)# prefix-list EXPORT_IPv4 out
tnsr(config-bgp-ip4uni-nbr)# activate
tnsr(config-bgp-ip4uni-nbr)# exit
tnsr(config-bgp-ip4uni)# exit
tnsr(config-bgp)#
```

Step 8: Configure BGP neighbor address-family IPv6 unicast options

```
tnsr(config-bgp)# address-family ipv6 unicast
tnsr(config-bgp-ip4uni)# neighbor 2001:db8:1000:aaaa::1
tnsr(config-bgp-ip4uni-nbr)# prefix-list EXPORT_IPv6 out
tnsr(config-bgp-ip4uni-nbr)# activate
```

(continues on next page)

(continued from previous page)

```

tnsr(config-bgp-ip4uni-nbr)# exit
tnsr(config-bgp-ip4uni)# neighbor 2001:db8:9999:ffff::1
tnsr(config-bgp-ip4uni-nbr)# prefix-list EXPORT_IPv6 out
tnsr(config-bgp-ip4uni-nbr)# activate
tnsr(config-bgp-ip4uni-nbr)# exit
tnsr(config-bgp-ip4uni)# exit
tnsr(config-bgp)# exit
tnsr(config-frr-bgp)# exit
tnsr(config)#

```

27.3.4 JSON Configuration

Listing 1: Download: tnsr-bgp-edge-router.json

```

1 {
2   "data": {
3     "bgp-config": {
4       "global-options": {
5         "enable": true
6       },
7       "routers": {
8         "router": [
9           {
10            "asn": 65505,
11            "router-id": "203.0.113.9",
12            "address-families": {
13              "address-family": [
14                {
15                  "family": "ipv4",
16                  "subfamily": "labeled-unicast"
17                },
18                {
19                  "family": "ipv4",
20                  "subfamily": "multicast"
21                },
22                {
23                  "family": "ipv4",
24                  "subfamily": "unicast",
25                  "neighbors": {
26                    "neighbor": [
27                      {
28                        "peer": "100.64.0.50",
29                        "activate": true,
30                        "prefix-list-out": "EXPORT_IPv4"
31                      },
32                      {
33                        "peer": "203.0.113.10",
34                        "activate": true,
35                        "prefix-list-out": "EXPORT_IPv4"
36                      }
37                    ]
38                  },
39                  "redistributions": {
40                    "named-sources": {
41                      "route-source": [

```

(continues on next page)

(continued from previous page)

```

42         {
43             "source": "kernel",
44             "present": true
45         }
46     ]
47 }
48 }
49 },
50 {
51     "family": "ipv4",
52     "subfamily": "vpn"
53 },
54 {
55     "family": "ipv6",
56     "subfamily": "labeled-unicast"
57 },
58 {
59     "family": "ipv6",
60     "subfamily": "multicast"
61 },
62 {
63     "family": "ipv6",
64     "subfamily": "unicast",
65     "neighbors": {
66         "neighbor": [
67             {
68                 "peer": "2001:db8:1000:aaaa::1",
69                 "activate": true,
70                 "prefix-list-out": "EXPORT_IPv6"
71             },
72             {
73                 "peer": "2001:db8:9999:ffff::1",
74                 "activate": true,
75                 "prefix-list-out": "EXPORT_IPv6"
76             }
77         ]
78     },
79     "redistributions": {
80         "named-sources": {
81             "route-source": [
82                 {
83                     "source": "kernel",
84                     "present": true
85                 }
86             ]
87         }
88     }
89 },
90 {
91     "family": "ipv6",
92     "subfamily": "vpn"
93 },
94 {
95     "family": "l2vpn",
96     "subfamily": "evpn"
97 },
98 {

```

(continues on next page)

(continued from previous page)

```

99         "family": "vpn4",
100         "subfamily": "unicast"
101     },
102     {
103         "family": "vpn6",
104         "subfamily": "unicast"
105     }
106 ]
107 },
108 "neighbors": {
109     "neighbor": [
110         {
111             "peer": "100.64.0.50",
112             "capability-negotiate": true,
113             "description": "<![CDATA[\"ISP_B IPv4\"]]>",
114             "interface": "GigabitEthernet0/14/0",
115             "remote-asn": 65520,
116             "enable": true
117         },
118         {
119             "peer": "2001:db8:1000:aaaa::1",
120             "capability-negotiate": true,
121             "description": "<![CDATA[\"ISP_A IPv6\"]]>",
122             "interface": "GigabitEthernet0/13/0",
123             "remote-asn": 65510,
124             "enable": true
125         },
126         {
127             "peer": "2001:db8:9999:ffff::1",
128             "capability-negotiate": true,
129             "description": "<![CDATA[\"ISP_B IPv6\"]]>",
130             "interface": "GigabitEthernet0/14/0",
131             "remote-asn": 65520,
132             "enable": true
133         },
134         {
135             "peer": "203.0.113.10",
136             "capability-negotiate": true,
137             "description": "<![CDATA[\"ISP_A IPv4\"]]>",
138             "interface": "GigabitEthernet0/13/0",
139             "remote-asn": 65510,
140             "enable": true
141         }
142     ]
143 }
144 }
145 ]
146 }
147 },
148 "interfaces-config": {
149     "interface": [
150         {
151             "name": "GigabitEthernet0/13/0",
152             "description": "<![CDATA[\"To ISP A\"]]>",
153             "enabled": true,
154             "ipv4": {
155                 "enabled": true,

```

(continues on next page)

(continued from previous page)

```

156     "forwarding": false,
157     "address": {
158         "ip": "203.0.113.9/30"
159     }
160 },
161 "ipv6": {
162     "enabled": true,
163     "forwarding": false,
164     "address": {
165         "ip": "2001:db8:1000:aaaa::2/64"
166     }
167 },
168 },
169 {
170     "name": "GigabitEthernet0/14/0",
171     "description": "<![CDATA[\"To ISP B\"]>",
172     "enabled": true,
173     "ipv4": {
174         "enabled": true,
175         "forwarding": false,
176         "address": {
177             "ip": "100.64.0.49/30"
178         }
179     },
180     "ipv6": {
181         "enabled": true,
182         "forwarding": false,
183         "address": {
184             "ip": "2001:db8:9999:ffff::2/64"
185         }
186     }
187 },
188 {
189     "name": "GigabitEthernet0/15/0",
190     "enabled": true,
191     "ipv4": {
192         "enabled": true,
193         "forwarding": false,
194         "address": {
195             "ip": "10.255.255.19/24"
196         }
197     }
198 }
199 ]
200 },
201 "http-config": {
202     "restconf": {
203         "enable": true
204     },
205     "authentication": {
206         "auth-type": "none"
207     }
208 },
209 "prefix-list-config": {
210     "prefix-lists": {
211         "list": [
212             {

```

(continues on next page)

(continued from previous page)

```

213     "name": "EXPORT_IPv4",
214     "description": "<![CDATA[\"IPv4 Routes to Export\"]]>",
215     "rules": {
216         "rule": [
217             {
218                 "sequence": 10,
219                 "action": "permit",
220                 "prefix": "192.0.2.0/24"
221             }
222         ]
223     },
224 },
225 {
226     "name": "EXPORT_IPv6",
227     "description": "<![CDATA[\"IPv6 Routes to Export\"]]>",
228     "rules": {
229         "rule": [
230             {
231                 "sequence": 10,
232                 "action": "permit",
233                 "prefix": "2001:db8:a100:1005::/64"
234             }
235         ]
236     },
237 },
238 ]
239 },
240 },
241 "route-table-config": {
242     "static-routes": {
243         "route-table": [
244             {
245                 "name": "ipv4-VRF:0",
246                 "address-family": "ipv4",
247                 "ipv4-routes": {
248                     "route": [
249                         {
250                             "destination-prefix": "192.0.2.0/24",
251                             "next-hop": {
252                                 "hop": [
253                                     {
254                                         "hop-id": 1,
255                                         "local": true
256                                     }
257                                 ]
258                             }
259                         }
260                     ]
261                 }
262             },
263             {
264                 "name": "ipv6-VRF:0",
265                 "address-family": "ipv6",
266                 "ipv6-routes": {
267                     "route": [
268                         {
269                             "destination-prefix": "2001:db8:a100:1005::/64",

```

(continues on next page)

(continued from previous page)

```

270         "next-hop": {
271             "hop": [
272                 {
273                     "hop-id": 1,
274                     "local": true
275                 }
276             ]
277         }
278     }
279 ]
280 }
281 }
282 ]
283 }
284 }
285 }
286 }

```

27.4 Service Provider Route Reflectors and Client for iBGP IPv4

27.4.1 Use Case

In large service provider networks it is necessary to divide the routing functionality into two or more layers: a backbone layer and a gateway layer. This allows backbone routers to be focused on core routing and switching to/from other areas of the routing domain, and gateway routers may then be focused on interconnecting other service provider customers.

27.4.2 Example Scenario

In this example, the service provider will have a fictitious autonomous system number (ASN) of 65505. Each network POP, of which only one will be detailed here, will feature 2 backbone routers which will be configured as route-reflectors. These backbone routers will be participating in BGP Cluster ID 100. Other POPs will likely be different Cluster IDs.

There will also be a single gateway router which will be a client of the backbone route-reflectors. Of course, in real world scenarios there would likely be many more gateway routers, each serving a full complement of customers.

Table 12: BGP Route Reflector Setup Parameters

Item	Value
TNSR Autonomous System Number	65505
IPv4 Networks to be announced	192.0.2.0/24, 203.0.113.0/24
BGP Route-Reflector Cluster ID	100

Scenario Topology

27.4.3 TNSR Configuration Steps

Step 1: Configure Interfaces

RR1:

Example: IPv4

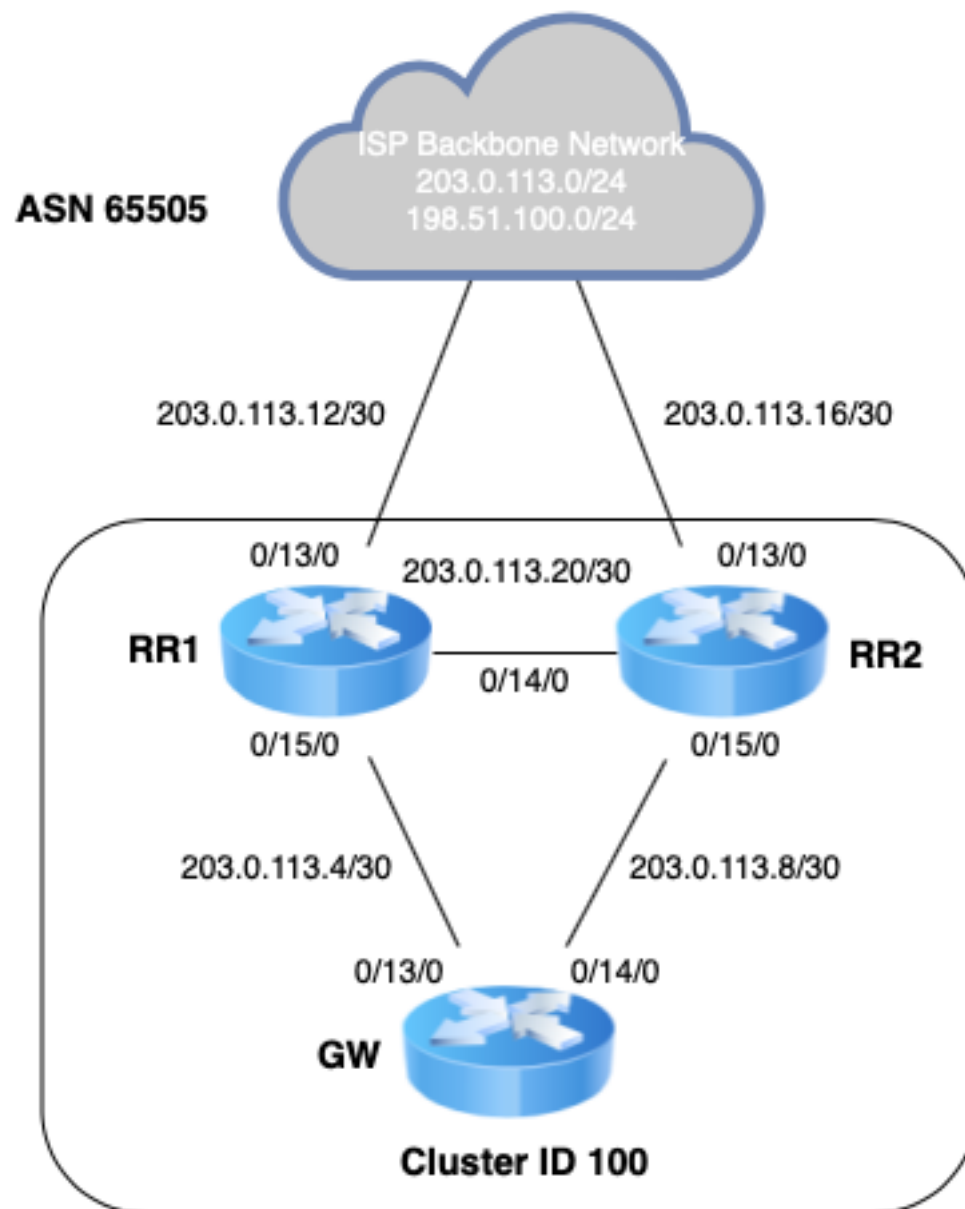


Fig. 4: TNSR BGP Route Reflector

```
rr1 tnsr# conf
rr1 tnsr(config)# interface GigabitEthernet0/13/0
rr1 tnsr(config-interface)# description "To Backbone Network"
rr1 tnsr(config-interface)# ip address 203.0.113.13/30
rr1 tnsr(config-interface)# enable
rr1 tnsr(config-interface)# exit
rr1 tnsr(config)# interface GigabitEthernet0/14/0
rr1 tnsr(config-interface)# description "To RR2 Router"
rr1 tnsr(config-interface)# ip address 203.0.113.21/30
rr1 tnsr(config-interface)# enable
rr1 tnsr(config-interface)# exit
rr1 tnsr(config)# interface GigabitEthernet0/15/0
rr1 tnsr(config-interface)# description "To GW router"
rr1 tnsr(config-interface)# ip address 203.0.113.5/30
rr1 tnsr(config-interface)# enable
rr1 tnsr(config-interface)# exit
rr1 tnsr(config)#
```

RR2:

```
rr2 tnsr# conf
rr2 tnsr(config)# interface GigabitEthernet0/13/0
rr2 tnsr(config-interface)# description "To Backbone Network"
rr2 tnsr(config-interface)# ip address 203.0.113.17/30
rr2 tnsr(config-interface)# enable
rr2 tnsr(config-interface)# exit
rr2 tnsr(config)# interface GigabitEthernet0/14/0
rr2 tnsr(config-interface)# description "To RR1 Router"
rr2 tnsr(config-interface)# ip address 203.0.113.22/30
rr2 tnsr(config-interface)# enable
rr2 tnsr(config-interface)# exit
rr2 tnsr(config)# interface GigabitEthernet0/15/0
rr2 tnsr(config-interface)# description "To GW router"
rr2 tnsr(config-interface)# ip address 203.0.113.9/30
rr2 tnsr(config-interface)# enable
rr2 tnsr(config-interface)# exit
rr2 tnsr(config)#
```

GW:

```
gw tnsr# conf
gw tnsr(config)# interface GigabitEthernet0/13/0
gw tnsr(config-interface)# description "To RR1 Router"
gw tnsr(config-interface)# ip address 203.0.113.6/30
gw tnsr(config-interface)# enable
gw tnsr(config-interface)# exit
gw tnsr(config)# interface GigabitEthernet0/14/0
gw tnsr(config-interface)# description "To RR2 Router"
gw tnsr(config-interface)# ip address 203.0.113.10/30
gw tnsr(config-interface)# enable
gw tnsr(config-interface)# exit
gw tnsr(config)# interface GigabitEthernet0/15/0
gw tnsr(config-interface)# desc "To Customer Router"
gw tnsr(config-interface)# ip address 203.0.113.25/30
gw tnsr(config-interface)# enable
gw tnsr(config-interface)# exit
gw tnsr(config)#
```

Step 2: Enable BGP

RR1:

```
rr1 tnsr(config)# route dynamic bgp
rr1 tnsr(config-frr-bgp)# enable
rr1 tnsr(config-frr-bgp)# exit
rr1 tnsr(config)#
```

RR2:

```
rr2 tnsr(config)# route dynamic bgp
rr2 tnsr(config-frr-bgp)# enable
rr2 tnsr(config-frr-bgp)# exit
rr2 tnsr(config)#
```

GW:

```
gw tnsr(config)# route dynamic bgp
gw tnsr(config-frr-bgp)# enable
gw tnsr(config-frr-bgp)# exit
gw tnsr(config)#
```

Step 3: Create prefix-lists for route import into BGP on Route-Reflectors

RR1:

```
rr1 tnsr(config)# route dynamic prefix-list REDISTRIBUTE_IPv4
rr1 tnsr(config-prefix-list)# description "IPv4 Routes to Import"
rr1 tnsr(config-prefix-list)# seq 10 permit 192.0.2.0/24
rr1 tnsr(config-prefix-list)# seq 20 permit 203.0.113.0/24
rr1 tnsr(config-prefix-list)# exit
rr1 tnsr(config)#
```

RR2:

```
rr2 tnsr(config)# route dynamic prefix-list REDISTRIBUTE_IPv4
rr2 tnsr(config-prefix-list)# description "IPv4 Routes to Import"
rr2 tnsr(config-prefix-list)# seq 10 permit 192.0.2.0/24
rr2 tnsr(config-prefix-list)# seq 20 permit 203.0.113.0/24
rr2 tnsr(config-prefix-list)# exit
rr2 tnsr(config)#
```

Step 4: Create route-map for route import into iBGP on route-reflectors

RR1:

```
rr1 tnsr(config)# route dynamic route-map REDISTRIBUTE_IPv4 permit sequence 10
rr1 tnsr(config-route-map)# match ip address prefix-list REDISTRIBUTE_IPv4
rr1 tnsr(config-route-map)# set origin igp
rr1 tnsr(config-route-map)# exit
rr1 tnsr(config)#
```

RR2:

```
rr2 tnsr(config)# route dynamic route-map REDISTRIBUTE_IPv4 permit sequence 10
rr2 tnsr(config-route-map)# match ip address prefix-list REDISTRIBUTE_IPv4
rr2 tnsr(config-route-map)# set origin igp
rr2 tnsr(config-route-map)# exit
rr2 tnsr(config)#
```

Step 5: Create static route for networks to be advertised in BGP

RR1:

```
rr1 tnsr(config)# route ipv4 table ipv4-VRF:0
rr1 tnsr(config-route-table-v4)# route 192.0.2.0/24
rr1 tnsr(config-rttbl4-next-hop)# next-hop 1 via local
rr1 tnsr(config-rttbl4-next-hop)# exit
rr1 tnsr(config-route-table-v4)# route 203.0.113.0/24
rr1 tnsr(config-rttbl4-next-hop)# next-hop 1 via local
rr1 tnsr(config-rttbl4-next-hop)# exit
rr1 tnsr(config-route-table-v4)# exit
rr1 tnsr(config)#
```

RR2:

```
rr2 tnsr(config)# route ipv4 table ipv4-VRF:0
rr2 tnsr(config-route-table-v4)# route 192.0.2.0/24
rr2 tnsr(config-rttbl4-next-hop)# next-hop 1 via local
rr2 tnsr(config-rttbl4-next-hop)# exit
rr2 tnsr(config-route-table-v4)# route 203.0.113.0/24
rr2 tnsr(config-rttbl4-next-hop)# next-hop 1 via local
rr2 tnsr(config-rttbl4-next-hop)# exit
rr2 tnsr(config-route-table-v4)# exit
rr2 tnsr(config)#
```

Step 6: Configure BGP global options

RR1:

```
rr1 tnsr(config)# route dynamic bgp
rr1 (config-frr-bgp)# server 65505
rr1 tnsr(config-bgp)# router-id 203.0.113.21
rr1 tnsr(config-bgp)# cluster-id 100
rr1 tnsr(config-bgp)# address-family ipv4 unicast
rr1 tnsr(config-bgp-ip4uni)# redistribute kernel route-map REDISTRIBUTE_IPv4
rr1 tnsr(config-bgp-ip4uni)# exit
rr1 tnsr(config-bgp)#
```

RR2:

```
rr1 tnsr(config)# route dynamic bgp
rr1 (config-frr-bgp)# server 65505
rr2 tnsr(config-bgp)# router-id 203.0.113.22
rr2 tnsr(config-bgp)# cluster-id 100
rr2 tnsr(config-bgp)# address-family ipv4 unicast
rr2 tnsr(config-bgp-ip4uni)# redistribute kernel route-map REDISTRIBUTE_IPv4
rr2 tnsr(config-bgp-ip4uni)# exit
rr2 tnsr(config-bgp)#
```

GW:

```
gw tnsr(config)# route dynamic bgp
gw (config-frr-bgp)# server 65505
gw tnsr(config-bgp)# router-id 203.0.113.6
gw tnsr(config-bgp)#
```

Step 7: Configure iBGP peer-group for backbone route-reflectors and add neighbor

RR1:

```
rr1 tnsr(config-bgp)# neighbor iBGP
rr1 tnsr(config-bgp-neighbor)# remote-as 65505
rr1 tnsr(config-bgp-neighbor)# description "iBGP Sessions"
rr1 tnsr(config-bgp-neighbor)# update-source GigabitEthernet0/14/0
rr1 tnsr(config-bgp-neighbor)# enable
rr1 tnsr(config-bgp-neighbor)# exit
rr1 tnsr(config-bgp)# neighbor 203.0.113.22
rr1 tnsr(config-bgp-neighbor)# peer-group iBGP
rr1 tnsr(config-bgp-neighbor)# enable
rr1 tnsr(config-bgp-neighbor)# exit
```

RR2:

```
rr2 tnsr(config-bgp)# neighbor iBGP
rr2 tnsr(config-bgp-neighbor)# remote-as 65505
rr2 tnsr(config-bgp-neighbor)# description "iBGP Sessions"
rr2 tnsr(config-bgp-neighbor)# update-source GigabitEthernet0/14/0
rr2 tnsr(config-bgp-neighbor)# enable
rr2 tnsr(config-bgp-neighbor)# exit
rr2 tnsr(config-bgp)# neighbor 203.0.113.21
rr2 tnsr(config-bgp-neighbor)# peer-group iBGP
rr2 tnsr(config-bgp-neighbor)# enable
rr2 tnsr(config-bgp-neighbor)# exit
```

Step 8: Configure RR-CLIENT peer-group for route-reflector clients and add neighbor

RR1:

```
rr1 tnsr(config-bgp)# neighbor RR-CLIENT
rr1 tnsr(config-bgp-neighbor)# remote-as 65505
rr1 tnsr(config-bgp-neighbor)# description "RR-Client Sessions"
rr1 tnsr(config-bgp-neighbor)# update-source GigabitEthernet0/15/0
rr1 tnsr(config-bgp-neighbor)# enable
rr1 tnsr(config-bgp-neighbor)# exit
rr1 tnsr(config-bgp)# neighbor 203.0.113.6
rr1 tnsr(config-bgp-neighbor)# peer-group RR-CLIENT
rr1 tnsr(config-bgp-neighbor)# enable
rr1 tnsr(config-bgp-neighbor)# exit
rr1 tnsr(config-bgp)#
```

RR2:

```
rr2 tnsr(config-bgp)# neighbor RR-CLIENT
rr2 tnsr(config-bgp-neighbor)# remote-as 65505
```

(continues on next page)

(continued from previous page)

```
rr2 tnsr(config-bgp-neighbor)# description "RR-Client Sessions"
rr2 tnsr(config-bgp-neighbor)# update-source GigabitEthernet0/15/0
rr2 tnsr(config-bgp-neighbor)# enable
rr2 tnsr(config-bgp-neighbor)# exit
rr2 tnsr(config-bgp)# neighbor 203.0.113.10
rr2 tnsr(config-bgp-neighbor)# peer-group RR-CLIENT
rr2 tnsr(config-bgp-neighbor)# enable
rr2 tnsr(config-bgp-neighbor)# exit
rr2 tnsr(config-bgp)#
```

Step 9: Configure both peer-group address-family options on route-reflectors

RR1:

```
rr1 tnsr(config-bgp)# address-family ipv4 unicast
rr1 tnsr(config-bgp-ip4uni)# neighbor iBGP
rr1 tnsr(config-bgp-ip4uni-nbr)# next-hop-self
rr1 tnsr(config-bgp-ip4uni-nbr)# activate
rr1 tnsr(config-bgp-ip4uni-nbr)# exit
rr1 tnsr(config-bgp-ip4uni)# neighbor RR-CLIENT
rr1 tnsr(config-bgp-ip4uni-nbr)# route-reflector-client
rr1 tnsr(config-bgp-ip4uni-nbr)# activate
rr1 tnsr(config-bgp-ip4uni-nbr)# exit
rr1 tnsr(config-bgp-ip4uni)# exit
rr1 tnsr(config-bgp)#
```

RR2:

```
rr2 tnsr(config-bgp)# address-family ipv4 unicast
rr2 tnsr(config-bgp-ip4uni)# neighbor iBGP
rr2 tnsr(config-bgp-ip4uni-nbr)# next-hop-self
rr2 tnsr(config-bgp-ip4uni-nbr)# activate
rr2 tnsr(config-bgp-ip4uni-nbr)# exit
rr2 tnsr(config-bgp-ip4uni)# neighbor RR-CLIENT
rr2 tnsr(config-bgp-ip4uni-nbr)# route-reflector-client
rr2 tnsr(config-bgp-ip4uni-nbr)# activate
rr2 tnsr(config-bgp-ip4uni-nbr)# exit
rr2 tnsr(config-bgp-ip4uni)# exit
rr2 tnsr(config-bgp)#
```

Step 10: Configure iBGP on gateway router to both route-reflectors

GW:

```
gw tnsr(config-bgp)# neighbor 203.0.113.5
gw tnsr(config-bgp-neighbor)# remote-as 65505
gw tnsr(config-bgp-neighbor)# description "RR1 Session"
gw tnsr(config-bgp-neighbor)# update-source GigabitEthernet0/13/0
gw tnsr(config-bgp-neighbor)# enable
gw tnsr(config-bgp-neighbor)# exit
gw tnsr(config-bgp)# neighbor 203.0.113.9
gw tnsr(config-bgp-neighbor)# remote-as 65505
gw tnsr(config-bgp-neighbor)# description "RR2 Session"
```

(continues on next page)

(continued from previous page)

```

gw tnsr(config-bgp-neighbor)# update-source GigabitEthernet0/14/0
gw tnsr(config-bgp-neighbor)# enable
gw tnsr(config-bgp-neighbor)# exit
gw tnsr(config-bgp)# address-family ipv4 unicast
gw tnsr(config-bgp-ip4uni)# neighbor 203.0.113.5
gw tnsr(config-bgp-ip4uni-nbr)# activate
gw tnsr(config-bgp-ip4uni-nbr)# exit
gw tnsr(config-bgp-ip4uni)# neighbor 203.0.113.9
gw tnsr(config-bgp-ip4uni-nbr)# activate
gw tnsr(config-bgp-ip4uni-nbr)# exit
gw tnsr(config-bgp-ip4uni)# exit
gw tnsr(config-bgp)#

```

27.4.4 JSON Configuration

RR1

Listing 2: Download: tnsr-bgp-router-reflector-rr1.
json

```

1 {
2   "data": {
3     "bgp-config": {
4       "global-options": {
5         "enable": true
6       },
7       "routers": {
8         "router": [
9           {
10            "asn": 65505,
11            "cluster-id": "100",
12            "router-id": "203.0.113.21",
13            "address-families": {
14              "address-family": [
15                {
16                  "family": "ipv4",
17                  "subfamily": "labeled-unicast"
18                },
19                {
20                  "family": "ipv4",
21                  "subfamily": "multicast"
22                },
23                {
24                  "family": "ipv4",
25                  "subfamily": "unicast",
26                  "neighbors": {
27                    "neighbor": [
28                      {
29                        "peer": "RR-CLIENT",
30                        "activate": true,
31                        "route-reflector-client": true
32                      },
33                      {
34                        "peer": "iBGP",

```

(continues on next page)

(continued from previous page)

```

35         "activate": true,
36         "next-hop-self": true
37     }
38 ]
39 },
40 "redistributions": {
41     "named-sources": {
42         "route-source": [
43             {
44                 "source": "kernel",
45                 "route-map": "REDISTRIBUTE_IPv4"
46             }
47         ]
48     }
49 },
50 },
51 {
52     "family": "ipv4",
53     "subfamily": "vpn"
54 },
55 {
56     "family": "ipv6",
57     "subfamily": "labeled-unicast"
58 },
59 {
60     "family": "ipv6",
61     "subfamily": "multicast"
62 },
63 {
64     "family": "ipv6",
65     "subfamily": "unicast"
66 },
67 {
68     "family": "ipv6",
69     "subfamily": "vpn"
70 },
71 {
72     "family": "l2vpn",
73     "subfamily": "evpn"
74 },
75 {
76     "family": "vpn4",
77     "subfamily": "unicast"
78 },
79 {
80     "family": "vpn6",
81     "subfamily": "unicast"
82 }
83 ]
84 },
85 "neighbors": {
86     "neighbor": [
87         {
88             "peer": "203.0.113.22",
89             "capability-negotiate": true,
90             "peer-group-name": "iBGP",
91             "enable": true

```

(continues on next page)

(continued from previous page)

```

92         },
93         {
94             "peer": "203.0.113.6",
95             "capability-negotiate": true,
96             "peer-group-name": "RR-CLIENT",
97             "enable": true
98         },
99         {
100             "peer": "RR-CLIENT",
101             "capability-negotiate": true,
102             "description": "<![CDATA[\"RR-Client Sessions\"]]>",
103             "remote-asn": 65505,
104             "enable": true,
105             "update-source": "GigabitEthernet0/15/0"
106         },
107         {
108             "peer": "iBGP",
109             "capability-negotiate": true,
110             "description": "<![CDATA[\"iBGP Sessions\"]]>",
111             "remote-asn": 65505,
112             "enable": true,
113             "update-source": "GigabitEthernet0/14/0"
114         }
115     ]
116 }
117 }
118 ]
119 }
120 },
121 "interfaces-config": {
122     "interface": [
123         {
124             "name": "GigabitEthernet0/13/0",
125             "description": "<![CDATA[\"To Backbone Network\"]]>",
126             "enabled": true,
127             "ipv4": {
128                 "enabled": true,
129                 "forwarding": false,
130                 "address": {
131                     "ip": "203.0.113.13/30"
132                 }
133             },
134             "ipv6": {
135                 "enabled": true,
136                 "forwarding": false
137             }
138         },
139         {
140             "name": "GigabitEthernet0/14/0",
141             "description": "<![CDATA[\"To RR2 Router\"]]>",
142             "enabled": true,
143             "ipv4": {
144                 "enabled": true,
145                 "forwarding": false,
146                 "address": {
147                     "ip": "203.0.113.21/30"
148                 }
149             }
150         }
151     ]
152 }

```

(continues on next page)

(continued from previous page)

```

149     },
150     "ipv6": {
151         "enabled": true,
152         "forwarding": false
153     }
154 },
155 {
156     "name": "GigabitEthernet0/15/0",
157     "description": "<![CDATA[\"To GW router\"]>",
158     "enabled": true,
159     "ipv4": {
160         "enabled": true,
161         "forwarding": false,
162         "address": {
163             "ip": "203.0.113.5/30"
164         }
165     },
166     "ipv6": {
167         "enabled": true,
168         "forwarding": false
169     }
170 }
171 ]
172 },
173 "prefix-list-config": {
174     "prefix-lists": {
175         "list": [
176             {
177                 "name": "REDISTRIBUTE_IPv4",
178                 "description": "<![CDATA[\"IPv4 Routes to Import\"]>",
179                 "rules": {
180                     "rule": [
181                         {
182                             "sequence": 10,
183                             "action": "permit",
184                             "prefix": "192.0.2.0/24"
185                         },
186                         {
187                             "sequence": 20,
188                             "action": "permit",
189                             "prefix": "203.0.113.0/24"
190                         }
191                     ]
192                 }
193             }
194         ]
195     }
196 },
197 "route-map-config": {
198     "route-maps": {
199         "map": [
200             {
201                 "name": "REDISTRIBUTE_IPv4",
202                 "rules": {
203                     "rule": [
204                         {
205                             "sequence": 10,

```

(continues on next page)

(continued from previous page)

```

206         "policy": "permit",
207         "match": {
208             "ip-address-prefix-list": "REDISTRIBUTE_IPv4"
209         },
210         "set": {
211             "origin": "igp"
212         }
213     }
214 ]
215 }
216 }
217 ]
218 }
219 },
220 "route-table-config": {
221     "static-routes": {
222         "route-table": [
223             {
224                 "name": "ipv4-VRF:0",
225                 "address-family": "ipv4",
226                 "ipv4-routes": {
227                     "route": [
228                         {
229                             "destination-prefix": "192.0.2.0/24",
230                             "next-hop": {
231                                 "hop": [
232                                     {
233                                         "hop-id": 1,
234                                         "local": true
235                                     }
236                                 ]
237                             }
238                         },
239                         {
240                             "destination-prefix": "203.0.113.0/24",
241                             "next-hop": {
242                                 "hop": [
243                                     {
244                                         "hop-id": 1,
245                                         "local": true
246                                     }
247                                 ]
248                             }
249                         }
250                     ]
251                 }
252             }
253         ]
254     }
255 }
256 }
257 }

```

RR2

Listing 3: Download: tnsr-bgp-router-reflector-rr2.
json

```
1 {
2   "data": {
3     "bgp-config": {
4       "global-options": {
5         "enable": true
6       },
7       "routers": {
8         "router": [
9           {
10            "asn": 65505,
11            "cluster-id": "100",
12            "router-id": "203.0.113.22",
13            "address-families": {
14              "address-family": [
15                {
16                  "family": "ipv4",
17                  "subfamily": "unicast",
18                  "neighbors": {
19                    "neighbor": [
20                      {
21                        "peer": "RR-CLIENT",
22                        "activate": true,
23                        "route-reflector-client": true
24                      },
25                      {
26                        "peer": "iBGP",
27                        "activate": true,
28                        "next-hop-self": true
29                      }
30                    ]
31                  },
32                  "redistributions": {
33                    "named-sources": {
34                      "route-source": [
35                        {
36                          "source": "kernel",
37                          "route-map": "REDISTRIBUTE_IPv4"
38                        }
39                      ]
40                    }
41                  }
42                },
43                {
44                  "family": "ipv6",
45                  "subfamily": "unicast",
46                  "redistributions": null
47                }
48              ]
49            },
50            "neighbors": {
51              "neighbor": [
52                {
53                  "peer": "203.0.113.10",
54                  "capability-negotiate": true,
```

(continues on next page)

(continued from previous page)

```

55         "peer-group-name": "RR-CLIENT",
56         "enable": true
57     },
58     {
59         "peer": "203.0.113.21",
60         "capability-negotiate": true,
61         "peer-group-name": "iBGP",
62         "enable": true
63     },
64     {
65         "peer": "RR-CLIENT",
66         "capability-negotiate": true,
67         "description": "<![CDATA[\"RR-Client Sessions\"]]>",
68         "remote-asn": 65505,
69         "enable": true,
70         "update-source": "GigabitEthernet0/15/0"
71     },
72     {
73         "peer": "iBGP",
74         "capability-negotiate": true,
75         "description": "<![CDATA[\"iBGP Sessions\"]]>",
76         "remote-asn": 65505,
77         "enable": true,
78         "update-source": "GigabitEthernet0/14/0"
79     }
80 ]
81 }
82 }
83 ]
84 }
85 },
86 "interfaces-config": {
87     "interface": [
88         {
89             "name": "GigabitEthernet0/13/0",
90             "description": "<![CDATA[\"To Backbone Network\"]]>",
91             "enabled": true,
92             "ipv4": {
93                 "enabled": true,
94                 "forwarding": false,
95                 "address": {
96                     "ip": "203.0.113.17/30"
97                 }
98             },
99             "ipv6": {
100                 "enabled": true,
101                 "forwarding": false
102             }
103         },
104         {
105             "name": "GigabitEthernet0/14/0",
106             "description": "<![CDATA[\"To RR1 Router\"]]>",
107             "enabled": true,
108             "ipv4": {
109                 "enabled": true,
110                 "forwarding": false,
111                 "address": {

```

(continues on next page)

(continued from previous page)

```

112         "ip": "203.0.113.22/30"
113     }
114 },
115     "ipv6": {
116         "enabled": true,
117         "forwarding": false
118     }
119 },
120 {
121     "name": "GigabitEthernet0/15/0",
122     "description": "<![CDATA[\"To GW router\"]]>",
123     "enabled": true,
124     "ipv4": {
125         "enabled": true,
126         "forwarding": false,
127         "address": {
128             "ip": "203.0.113.9/30"
129         }
130     },
131     "ipv6": {
132         "enabled": true,
133         "forwarding": false
134     }
135 }
136 ]
137 },
138 "prefix-list-config": {
139     "prefix-lists": {
140         "list": [
141             {
142                 "name": "REDISTRIBUTE_IPv4",
143                 "description": "<![CDATA[\"IPv4 Routes to Import\"]]>",
144                 "rules": {
145                     "rule": [
146                         {
147                             "sequence": 10,
148                             "action": "permit",
149                             "prefix": "192.0.2.0/24"
150                         },
151                         {
152                             "sequence": 20,
153                             "action": "permit",
154                             "prefix": "203.0.113.0/24"
155                         }
156                     ]
157                 }
158             }
159         ]
160     }
161 },
162 "route-map-config": {
163     "route-maps": {
164         "map": [
165             {
166                 "name": "REDISTRIBUTE_IPv4",
167                 "rules": {
168                     "rule": [

```

(continues on next page)

(continued from previous page)

```

169         {
170             "sequence": 10,
171             "policy": "permit",
172             "match": {
173                 "ip-address-prefix-list": "REDISTRIBUTE_IPv4"
174             },
175             "set": {
176                 "origin": "igp"
177             }
178         }
179     ]
180 }
181 }
182 ]
183 }
184 },
185 "route-table-config": {
186     "static-routes": {
187         "route-table": [
188             {
189                 "name": "ipv4-VRF:0",
190                 "address-family": "ipv4",
191                 "ipv4-routes": {
192                     "route": [
193                         {
194                             "destination-prefix": "192.0.2.0/24",
195                             "next-hop": {
196                                 "hop": [
197                                     {
198                                         "hop-id": 1,
199                                         "local": true
200                                     }
201                                 ]
202                             }
203                         },
204                         {
205                             "destination-prefix": "203.0.113.0/24",
206                             "next-hop": {
207                                 "hop": [
208                                     {
209                                         "hop-id": 1,
210                                         "local": true
211                                     }
212                                 ]
213                             }
214                         }
215                     ]
216                 }
217             }
218         ]
219     }
220 }
221 }
222 }

```

GW

Listing 4: Download: tnsr-bgp-router-reflector-gw.json

```

1 {
2   "data": {
3     "bgp-config": {
4       "global-options": {
5         "enable": true
6       },
7       "routers": {
8         "router": [
9           {
10            "asn": 65505,
11            "router-id": "203.0.113.6",
12            "address-families": {
13              "address-family": [
14                {
15                  "family": "ipv4",
16                  "subfamily": "labeled-unicast"
17                },
18                {
19                  "family": "ipv4",
20                  "subfamily": "multicast"
21                },
22                {
23                  "family": "ipv4",
24                  "subfamily": "unicast",
25                  "neighbors": {
26                    "neighbor": [
27                      {
28                        "peer": "203.0.113.5",
29                        "activate": true
30                      },
31                      {
32                        "peer": "203.0.113.9",
33                        "activate": true
34                      }
35                    ]
36                  }
37                },
38                {
39                  "family": "ipv4",
40                  "subfamily": "vpn"
41                },
42                {
43                  "family": "ipv6",
44                  "subfamily": "labeled-unicast"
45                },
46                {
47                  "family": "ipv6",
48                  "subfamily": "multicast"
49                },
50                {
51                  "family": "ipv6",
52                  "subfamily": "unicast"
53                }

```

(continues on next page)

(continued from previous page)

```

54         {
55             "family": "ipv6",
56             "subfamily": "vpn"
57         },
58         {
59             "family": "l2vpn",
60             "subfamily": "evpn"
61         },
62         {
63             "family": "vpn4",
64             "subfamily": "unicast"
65         },
66         {
67             "family": "vpn6",
68             "subfamily": "unicast"
69         }
70     ]
71 },
72 "neighbors": {
73     "neighbor": [
74         {
75             "peer": "203.0.113.5",
76             "capability-negotiate": true,
77             "description": "<![CDATA[\"RR1 Session\"]]>",
78             "remote-asn": 65505,
79             "enable": true,
80             "update-source": "GigabitEthernet0/13/0"
81         },
82         {
83             "peer": "203.0.113.9",
84             "capability-negotiate": true,
85             "description": "<![CDATA[\"RR2 Session\"]]>",
86             "remote-asn": 65505,
87             "enable": true,
88             "update-source": "GigabitEthernet0/14/0"
89         }
90     ]
91 }
92 }
93 ]
94 }
95 },
96 "interfaces-config": {
97     "interface": [
98         {
99             "name": "GigabitEthernet0/13/0",
100             "description": "<![CDATA[\"To RR1 Router\"]]>",
101             "enabled": true,
102             "ipv4": {
103                 "enabled": true,
104                 "forwarding": false,
105                 "address": {
106                     "ip": "203.0.113.6/30"
107                 }
108             },
109             "ipv6": {
110                 "enabled": true,

```

(continues on next page)

(continued from previous page)

```

111         "forwarding": false
112     },
113 },
114 {
115     "name": "GigabitEthernet0/14/0",
116     "description": "<![CDATA[\"To RR2 Router\"]>",
117     "enabled": true,
118     "ipv4": {
119         "enabled": true,
120         "forwarding": false,
121         "address": {
122             "ip": "203.0.113.10/30"
123         }
124     },
125     "ipv6": {
126         "enabled": true,
127         "forwarding": false
128     }
129 },
130 {
131     "name": "GigabitEthernet0/15/0",
132     "description": "<![CDATA[\"To Customer Router\"]>",
133     "enabled": true,
134     "ipv4": {
135         "enabled": true,
136         "forwarding": false,
137         "address": {
138             "ip": "203.0.113.25/30"
139         }
140     },
141     "ipv6": {
142         "enabled": true,
143         "forwarding": false
144     }
145 }
146 ]
147 }
148 }
149 }

```

27.5 LAN + WAN with NAT (Basic SOHO Router Including DHCP and DNS Resolver)

27.5.1 Use Case

A typical use case for TNSR is a device that sits between a local area network (LAN) in an office or home and a wide area network (WAN) such as the Internet.

At a minimum, such a TNSR instance routes traffic between the LAN and the WAN. In many cases, it provides additional services that are useful for a LAN, including:

- DHCP to provide hosts in the LAN with IP addresses.
- DNS to respond to name resolution queries from hosts in the LAN

- NAT (Network Address Translation), to map one public IPv4 address to internal (private) IP addresses assigned to hosts on the LAN.

27.5.2 Example Scenario

This example configures TNSR with basic the basic functions mentioned earlier: DHCP, DNS, and NAT

Item	Value
Local PC	DHCP: 172.16.1.100/24
TNSR Local Interface	GigabitEthernet0/14/2
TNSR Local Address	172.16.1.1/24
TNSR Internet Interface	GigabitEthernet0/14/1
TNSR Internet Address	203.0.113.2/24
Remote DNS	8.8.8.8, 8.8.4.4




Fig. 5: Basic SOHO Router Example

27.5.3 TNSR Configuration

Basic Connectivity

First, there is the basic interface configuration of TNSR to handle IP connectivity:

```
tnsr(config)# interface GigabitEthernet0/14/2
tnsr(config-interface)# ip address 172.16.1.1/24
tnsr(config-interface)# description Local
tnsr(config-interface)# enable
tnsr(config-interface)# exit
```

```
tnsr(config)# interface GigabitEthernet0/14/1
tnsr(config-interface)# ip address 203.0.113.2/24
tnsr(config-interface)# description Internet
tnsr(config-interface)# enable
tnsr(config-interface)# exit
```

DHCP

Next, configure the DHCP server and DHCP pool on TNSR:

```
tnsr(config)# dhcp4 server
tnsr(config-kea-dhcp4)# description LAN DHCP Server
tnsr(config-kea-dhcp4)# interface listen GigabitEthernet0/14/2
tnsr(config-kea-dhcp4)# option domain-name
tnsr(config-kea-dhcp4-opt)# data example.com
tnsr(config-kea-dhcp4-opt)# exit
tnsr(config-kea-dhcp4)# subnet 172.16.1.0/24
tnsr(config-kea-subnet4)# pool 172.16.1.100-172.16.1.245
tnsr(config-kea-subnet4-pool)# exit
tnsr(config-kea-subnet4)# interface GigabitEthernet0/14/2
tnsr(config-kea-subnet4)# option domain-name-servers
tnsr(config-kea-subnet4-opt)# data 172.16.1.1
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-subnet4)# option routers
tnsr(config-kea-subnet4-opt)# data 172.16.1.1
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-dhcp4)# exit
tnsr(config)# dhcp4 enable
```

The above example configures `example.com` as the domain name supplied to all clients. For the specific subnet in the example, the TNSR IP address inside the subnet is supplied by DHCP as the default gateway for clients, and DHCP will instruct clients to use the DNS Resolver daemon on TNSR at `172.16.1.1` for DNS.

Outbound NAT

Now configure Outbound NAT:

```
tnsr(config)# nat pool addresses 203.0.113.2
tnsr(config)# interface GigabitEthernet0/14/1
tnsr(config-interface)# ip nat outside
tnsr(config-interface)# exit
tnsr(config)# interface GigabitEthernet0/14/2
tnsr(config-interface)# ip nat inside
tnsr(config-interface)# exit
tnsr(config)# nat global-options nat44 forwarding true
tnsr(config)#
```

DNS Resolver

Finally, configure a DNS Resolver in forwarding mode:

```
tnsr# configure
tnsr(config)# unbound server
tnsr(config-unbound)# interface 127.0.0.1
tnsr(config-unbound)# interface 172.16.1.1
tnsr(config-unbound)# outgoing-interface 203.0.113.2
tnsr(config-unbound)# access-control 172.16.1.0/24 allow
tnsr(config-unbound)# forward-zone .
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.8.8
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.4.4
tnsr(config-unbound-fwd-zone)# exit
tnsr(config-unbound)# exit
tnsr(config)# unbound enable
```

This example enables the Unbound DNS service and configures it to listen on localhost as well as 172.16.1.1 (GigabitEthernet0/14/2, labeled LAN in the example). It uses 203.0.113.2, which is the example WAN interface address, for outgoing queries. The example also allows clients inside that subnet, 172.16.1.0/24, to perform DNS queries and receive responses. It will send all DNS queries to the upstream DNS servers 8.8.8.8 and 8.8.4.4.

27.5.4 Local PC Configuration

No configuration is necessary on the Local PC, it will pull all its required settings from DHCP.

27.6 Using Access Control Lists (ACLs)

27.6.1 Use Case

A standard ACL works with IPv4 or IPv6 traffic at layer 3. The name of an ACL is arbitrary so it may be named in a way that makes its purpose obvious.

ACLs consist of one or more rules, defined by a sequence number that determines the order in which the rules are applied. A common practice is to start numbering at a value higher than 0 or 1, and to leave gaps in the sequence so that rules may be added later. For example, the first rule could be 10, followed by 20.

27.6.2 Example Scenario

This example configures TNSR with an ACL that allows SSH, ICMP and HTTP/HTTPs connections only from a specific Remote Admin Host:

Item	Value
Local PC	DHCP: 172.16.1.100/24
TNSR Local Interface	GigabitEthernet0/14/2
TNSR Local Address	172.16.1.1/24
TNSR Internet Interface	GigabitEthernet0/14/1
TNSR Internet Address	203.0.113.2/24
Remote Admin Host	208.123.73.10/24




Fig. 6: ACL Example Scenario

27.6.3 TNSR Configuration

```
tnsr(config)# acl WAN_protecting_acl
tnsr(config-acl)# rule 10
tnsr(config-acl-rule)# action permit
tnsr(config-acl-rule)# ip-version ipv4
tnsr(config-acl-rule)# destination address 203.0.113.2/32
tnsr(config-acl-rule)# destination port 22
tnsr(config-acl-rule)# source ip address 208.123.73.10/32
tnsr(config-acl-rule)# protocol tcp
tnsr(config-acl-rule)# exit
tnsr(config-acl)# rule 20
tnsr(config-acl-rule)# action permit
tnsr(config-acl-rule)# ip-version ipv4
tnsr(config-acl-rule)# destination address 203.0.113.2/32
tnsr(config-acl-rule)# destination port 80
tnsr(config-acl-rule)# source ip address 208.123.73.10/32
tnsr(config-acl-rule)# protocol tcp
tnsr(config-acl-rule)# exit
tnsr(config-acl)# rule 30
tnsr(config-acl-rule)# action permit
tnsr(config-acl-rule)# ip-version ipv4
tnsr(config-acl-rule)# destination address 203.0.113.2/32
tnsr(config-acl-rule)# destination port 443
tnsr(config-acl-rule)# source ip address 208.123.73.10/32
tnsr(config-acl-rule)# protocol tcp
tnsr(config-acl-rule)# exit
tnsr(config-acl)# rule 40
tnsr(config-acl-rule)# action deny
tnsr(config-acl-rule)# ip-version ipv4
tnsr(config-acl-rule)# destination port 22
tnsr(config-acl-rule)# protocol tcp
tnsr(config-acl-rule)# exit
tnsr(config-acl)# rule 50
tnsr(config-acl-rule)# action deny
tnsr(config-acl-rule)# ip-version ipv4
tnsr(config-acl-rule)# destination port 80
tnsr(config-acl-rule)# protocol tcp
tnsr(config-acl-rule)# exit
tnsr(config-acl)# rule 60
tnsr(config-acl-rule)# action deny
tnsr(config-acl-rule)# ip-version ipv4
tnsr(config-acl-rule)# destination port 443
tnsr(config-acl-rule)# protocol tcp
tnsr(config-acl-rule)# exit
tnsr(config-acl)# exit
tnsr(config-acl)# rule 70
tnsr(config-acl-rule)# action permit
tnsr(config-acl-rule)# ip-version ipv4
tnsr(config-acl-rule)# exit
tnsr(config)# int GigabitEthernet0/14/1
tnsr(config-interface)# access-list input acl WAN_protecting_acl sequence 10
tnsr(config-interface)# exit
tnsr(config)#
```

Rules 10-30 allow SSH, HTTP and HTTPs access to the WAN IP address from the Remote Admin Host. Then Rules 40-60 block SSH, HTTPS and HTTPs on the WAN IP address from all other IP addresses. Finally, rule 70 allows all other incoming traffic.

27.7 Inter-VLAN Routing

27.7.1 Use Case

Inter-VLAN routing is a process of forwarding network traffic from one VLAN to another VLAN using a router or layer 3 device.

27.7.2 Example Scenario

This example configures TNSR with VLANs:

Item	Value
TNSR Internet Interface	GigabitEthernet0/14/1
TNSR Internet Address	203.0.113.2/24
TNSR Local Interface	GigabitEthernet0/14/2
TNSR VLAN 10 Interface	GigabitEthernet0/14/2.10
TNSR VLAN 10 Address	172.16.10.1/24
TNSR VLAN 20 Interface	GigabitEthernet0/14/2.20
TNSR VLAN 20 Address	172.16.20.1/24




Fig. 7: Inter-VLAN Routing Example

27.7.3 TNSR Configuration

A few pieces of information are necessary to create a VLAN subinterface (“subif”):

- The parent interface which will carry the tagged traffic, e.g. GigabitEthernet3/0/0
- The subinterface ID number, which is a positive integer that uniquely identifies this subif on the parent interface. It is commonly set to the same value as the VLAN tag
- The VLAN tag used by the subif to tag outgoing traffic, and to use for identifying incoming traffic bound for this subif. This is an integer in the range 1–4095, inclusive. This VLAN must also be tagged on the corresponding switch configuration for the port used by the parent interface.

Create Subinterfaces

First, create subinterfaces for VLAN 10 and VLAN 20:

```
tnsr(config)# interface subif GigabitEthernet0/14/2 10
tnsr(config-subif)# dot1q 10
tnsr(config-subif)# exact-match
tnsr(config-subif)# exit
```

```
tnsr(config)# interface subif GigabitEthernet0/14/2 20
tnsr(config-subif)# dot1q 20
tnsr(config-subif)# exact-match
tnsr(config-subif)# exit
```

The subif interface appears with the parent interface name and the subif id, joined by a ..

Configure Interfaces

At this point, subinterface behaves identically to a regular interface in that it may have an IP address, routing, and so on:

```
tnsr(config)# interface GigabitEthernet0/14/2.10
tnsr(config-interface)# ip address 172.16.10.1/24
tnsr(config-interface)# description VLAN10
tnsr(config-interface)# enable
tnsr(config-interface)# exit
```

```
tnsr(config)# interface GigabitEthernet0/14/2.20
tnsr(config-interface)# ip address 172.16.20.1/24
tnsr(config-interface)# description VLAN20
tnsr(config-interface)# enable
tnsr(config-interface)# exit
```

Configure DHCP

Next, configure the DHCP server and DHCP pool on TNSR for each VLAN.

For VLAN 10:

```
tnsr(config)# dhcp4 server
tnsr(config-kea-dhcp4)# description LAN DHCP Server
tnsr(config-kea-dhcp4)# interface listen GigabitEthernet0/14/2.10
tnsr(config-kea-dhcp4)# option domain-name
tnsr(config-kea-dhcp4-opt)# data example.com
tnsr(config-kea-dhcp4-opt)# exit
tnsr(config-kea-dhcp4)# subnet 172.16.10.0/24
tnsr(config-kea-subnet4)# pool 172.16.10.100-172.16.10.245
tnsr(config-kea-subnet4-pool)# exit
tnsr(config-kea-subnet4)# interface GigabitEthernet0/14/2.10
tnsr(config-kea-subnet4)# option domain-name-servers
tnsr(config-kea-subnet4-opt)# data 172.16.10.1
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-subnet4)# option routers
tnsr(config-kea-subnet4-opt)# data 172.16.10.1
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-dhcp4)# exit
```

And for VLAN 20:

```
tnsr(config)# dhcp4 server
tnsr(config-kea-dhcp4)# interface listen GigabitEthernet0/14/2.20
tnsr(config-kea-dhcp4)# subnet 172.16.20.0/24
tnsr(config-kea-subnet4)# pool 172.16.20.100-172.16.20.245
tnsr(config-kea-subnet4-pool)# exit
tnsr(config-kea-subnet4)# interface GigabitEthernet0/14/2.20
tnsr(config-kea-subnet4)# option domain-name-servers
tnsr(config-kea-subnet4-opt)# data 172.16.20.1
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-subnet4)# option routers
tnsr(config-kea-subnet4-opt)# data 172.16.20.1
tnsr(config-kea-subnet4-opt)# exit
tnsr(config-kea-dhcp4)# exit
tnsr(config)# dhcp4 enable
```

Configure Outbound NAT

Now configure Outbound NAT:

```
tnsr(config)# nat pool addresses 203.0.113.2
tnsr(config)# interface GigabitEthernet0/14/1
tnsr(config-interface)# ip nat outside
tnsr(config-interface)# exit
tnsr(config)# interface GigabitEthernet0/14/2.10
tnsr(config-interface)# ip nat inside
tnsr(config-interface)# exit
tnsr(config)# interface GigabitEthernet0/14/2.20
tnsr(config-interface)# ip nat inside
tnsr(config-interface)# exit
tnsr(config)# nat global-options nat44 forwarding true
tnsr(config)#
```

Configure DNS Resolver

Finally, configure a DNS Resolver in forwarding mode:

```

tnsr# configure
tnsr(config)# unbound server
tnsr(config-unbound)# interface 127.0.0.1
tnsr(config-unbound)# interface 172.16.10.1
tnsr(config-unbound)# interface 172.16.20.1
tnsr(config-unbound)# outgoing-interface 203.0.113.2
tnsr(config-unbound)# access-control 172.16.10.0/24 allow
tnsr(config-unbound)# access-control 172.16.20.0/24 allow
tnsr(config-unbound)# forward-zone .
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.8.8
tnsr(config-unbound-fwd-zone)# nameserver address 8.8.4.4
tnsr(config-unbound-fwd-zone)# exit
tnsr(config-unbound)# exit
tnsr(config)# unbound enable

```

Now there are two VLANs on the physical “LAN” port and interface GigabitEthernet0/14/2 now works as trunk port between TNSR and downstream L2/L3 switch.

This switch must be configured to match the expected VLAN tags and it must also have access ports configured for clients on each VLAN.

27.8 GRE ERSPAN Example Use Case

Encapsulated Remote Switched Port Analyzer (ERSPAN) is a type of GRE tunnel which allows a remote Intrusion Detection System (IDS) or similar packet inspection device to receive copies of packets from a local interface. This operates similar to a local mirror or span port on a switch, but in a remote capacity.

A typical use case for this is central packet inspection or a case where a remote site has plenty of bandwidth available, but no suitable local hardware for inspecting packets.

On TNSR, this is accomplished by configuring an ERSPAN GRE tunnel and then configuring a span to link the ERSPAN tunnel a local interface. From that point on, a copy of every packet on the interface being spanned is sent across GRE.

Note: The receiving end does not need to support ERSPAN, a standard GRE tunnel will suffice.

See also:

In environments which do not allow GRE traffic, such as Azure, VXLAN interfaces may be used instead. See [VXLAN SPAN Example](#).

27.8.1 Example Scenario

In this example, copies of packets from a local TNSR interface will be copied to a remote IDS for inspection.

Item	Value
Local Server:	172.29.193.47/24
TNSR Local Interface:	VirtualFunctionEthernet0/6/0
TNSR Local Address:	172.29.193.60/24
TNSR Internet Interface:	VirtualFunctionEthernet0/7/0
TNSR Internet Address:	172.29.194.142/24
IDS Address:	172.29.194.90/24

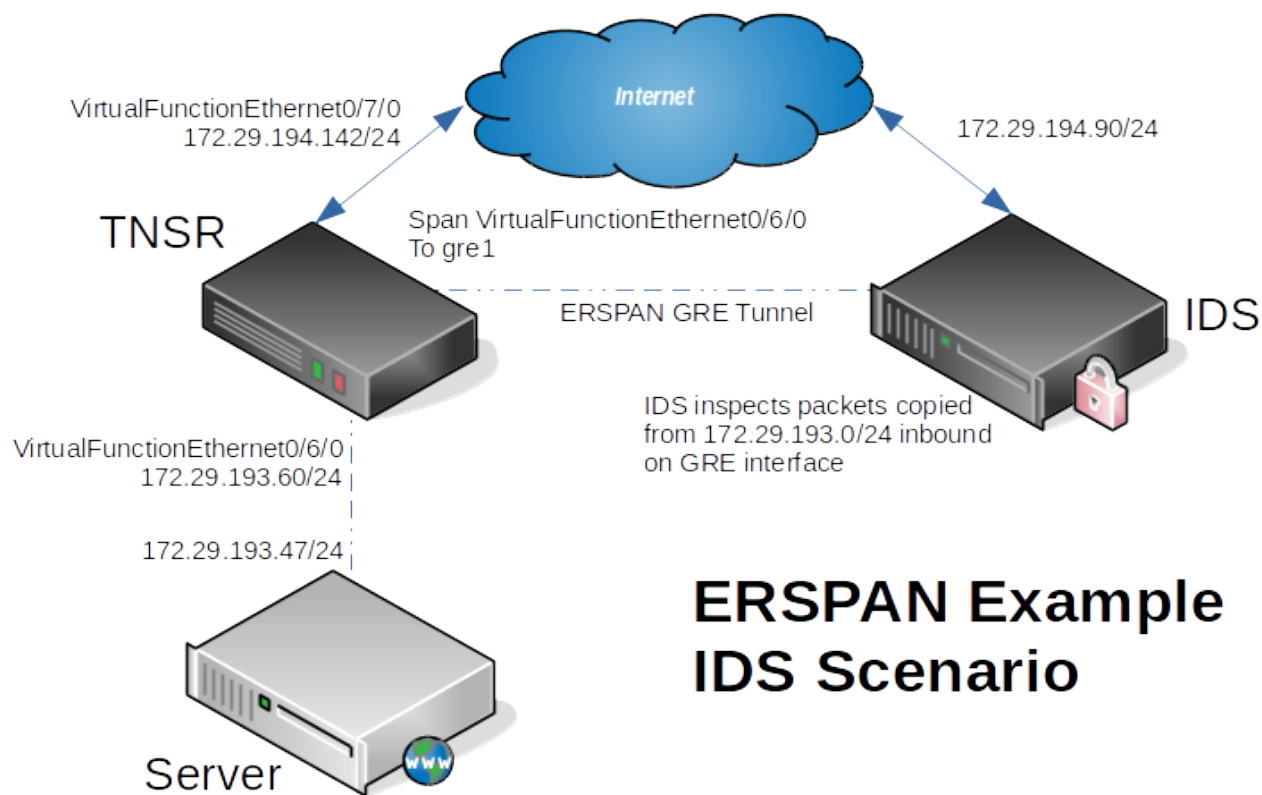


Fig. 8: ERSPAN Example

27.8.2 TNSR Configuration

First, there is the basic interface configuration of TNSR to handle IP connectivity:

```
tnsr(config)# interface VirtualFunctionEthernet0/6/0
tnsr(config-interface)# ip address 172.29.193.160/24
tnsr(config-interface)# description Local
tnsr(config-interface)# enable
tnsr(config-interface)# exit

tnsr(config)# interface VirtualFunctionEthernet0/7/0
tnsr(config-interface)# ip address 172.29.194.142/24
tnsr(config-interface)# description Internet
tnsr(config-interface)# enable
tnsr(config-interface)# exit
```

Next, configure the GRE tunnel on TNSR:

```
tnsr(config)# gre gre1
tnsr(config-gre)# destination 172.29.194.90
tnsr(config-gre)# source 172.29.194.142
tnsr(config-gre)# tunnel-type erspan session-id 1
tnsr(config-gre)# instance 1
tnsr(config-gre)# exit

tnsr(config)# interface gre1
tnsr(config-interface)# enable
```

(continues on next page)

(continued from previous page)

```
tnsr(config-interface) # exit
```

Finally, configure a SPAN that ties the local interface to the GRE interface:

```
tnsr(config) # span VirtualFunctionEthernet0/6/0
tnsr(config-span) # onto gre1 hw both
tnsr(config-span) # exit
```

27.8.3 Server Configuration

No configuration is necessary on the server. Any packet it sends which flows through TNSR will automatically be copied across the ERSPAN tunnel to the IDS.

27.8.4 IDS Configuration

The IDS must support GRE interfaces and also must support inspecting packets on GRE interfaces. The IDS does not need to explicitly support ERSPAN to receive copies of packets from TNSR.

At a minimum, take the following steps on the IDS:

- Configure a GRE tunnel between the IDS and TNSR, it does not need to have an address internal to the GRE tunnel.
- Configure the IDS software to inspect packets on the GRE interface

27.9 OSPF Router with Multiple Areas and Summarization

27.9.1 Example Scenario

This recipe demonstrates two routers which handle traffic for multiple local networks. Though it is a simple configuration, multiple areas are used so that routes for each site may be summarized.

Summarization reduces the number of routes that each neighbor must advertise and reduces the number of routes that each neighbor must maintain in its local database. As networks grow, this becomes an important factor when resources are constrained. This example allows for significant future expansion with little or no increase in OSPF database complexity for peers.

In modern networking environments, most implementations like TNSR are capable of handling many thousands of routes in a single area. Even so, using multiple areas with summarization can be easier for administrators to manage and troubleshoot.

Since each of these routers is connected to more than one area, each becomes an Area Border Router (ABR). As such, they are capable of route summarization using Type 3 Link State Advertisement (LSA) messages.

Note: This example ignores external connectivity, only focusing on the relationship between two routers and their component networks.

Additionally, since each of these routers is not connected to other routers outside the backbone network, their local areas can be considered stub areas and the local interfaces can be configured as passive interfaces.

See also:

For a simpler example involving a single area, see *OSPF Example*.

Scenario Topology

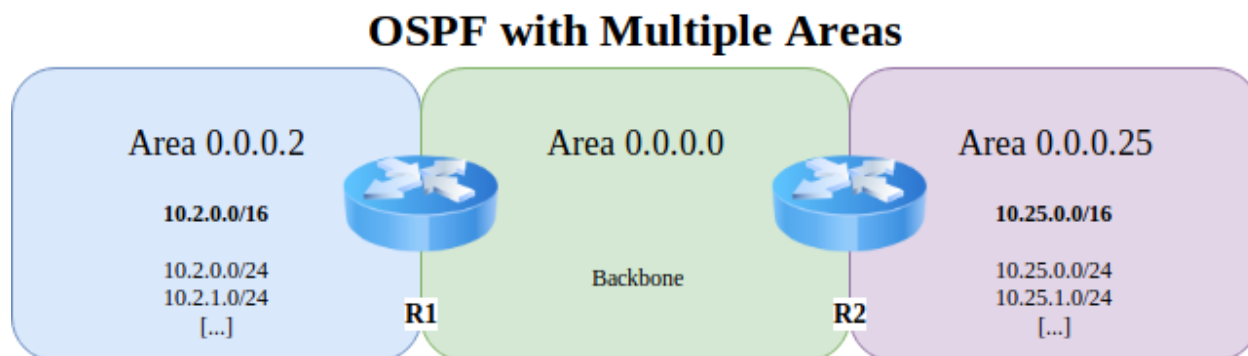


Fig. 9: TNSR OSPF with ABR

Scenario Information

Table 13: Shared Information

Item	Value
OSPF Backbone Area	0.0.0.0
Backbone Network	172.16.0.0/24

Table 14: Router 1 Information

Item	Value
Backbone Address	172.16.0.2/24
Local Router ID	10.2.0.1
Local OSPF Area	0.0.0.2
Active Interfaces (Cost)	TenGigabitEthernet6/0/0 (5)
Passive Interfaces	TenGigabitEthernet6/0/1, TenGigabitEthernet8/0/0
Local Networks	10.2.0.0/24, 10.2.1.0/24
Local Network Summary	10.2.0.0/16

Table 15: Router 2 Information

Item	Value
Backbone Address	172.16.0.25/24
Local Router ID	10.25.0.1
Local OSPF Area	0.0.0.25
Active Interfaces (Cost)	GigabitEthernet3/0/0 (5)
Passive Interfaces	GigabitEthernet0/13/0, GigabitEthernet0/14/0
Local Networks	10.25.0.0/24, 10.25.1.0/24
Local Network Summary	10.25.0.0/16

27.9.2 TNSR Configuration Steps

Configure Interfaces on R1

```
r1 tnsr# conf
r1 tnsr(config)# interface TenGigabitEthernet6/0/0
r1 tnsr(config-interface)# description "To Backbone"
r1 tnsr(config-interface)# ip address 172.16.0.2/24
r1 tnsr(config-interface)# mtu 1500
r1 tnsr(config-interface)# enable
r1 tnsr(config-interface)# exit
r1 tnsr(config)#
r1 tnsr(config)# interface TenGigabitEthernet6/0/1
r1 tnsr(config-interface)# description "Local Network 1"
r1 tnsr(config-interface)# ip address 10.2.0.1/24
r1 tnsr(config-interface)# enable
r1 tnsr(config-interface)# exit
r1 tnsr(config)#
r1 tnsr(config)# interface TenGigabitEthernet8/0/0
r1 tnsr(config-interface)# description "Local Network 2"
r1 tnsr(config-interface)# ip address 10.2.1.1/24
r1 tnsr(config-interface)# enable
r1 tnsr(config-interface)# exit
r1 tnsr(config)#
```

Configure Interfaces on R2

```
r2 tnsr# conf
r2 tnsr(config)# interface GigabitEthernet3/0/0
r2 tnsr(config-interface)# description "To Backbone"
r2 tnsr(config-interface)# ip address 172.16.0.25/24
r1 tnsr(config-interface)# mtu 1500
r2 tnsr(config-interface)# enable
r2 tnsr(config-interface)# exit
r2 tnsr(config)#
r2 tnsr(config)# interface GigabitEthernet0/13/0
r2 tnsr(config-interface)# description "Local Network 1"
r2 tnsr(config-interface)# ip address 10.25.0.1/24
r2 tnsr(config-interface)# enable
r2 tnsr(config-interface)# exit
r2 tnsr(config)#
r2 tnsr(config)# interface GigabitEthernet0/14/0
r2 tnsr(config-interface)# description "Local Network 2"
r2 tnsr(config-interface)# ip address 10.25.1.1/24
r2 tnsr(config-interface)# enable
r2 tnsr(config-interface)# exit
r2 tnsr(config)#
```

Configure OSPF on R1

```
r1 tnsr(config)# route dynamic ospf
r1 tnsr(config-frr-ospf)# server
r1 tnsr(config-ospf)# ospf router-id 10.2.0.1
```

(continues on next page)

(continued from previous page)

```
r1 tnsr(config-ospf)# passive-interface TenGigabitEthernet6/0/1
r1 tnsr(config-ospf)# passive-interface TenGigabitEthernet8/0/0
r1 tnsr(config-ospf)# interface TenGigabitEthernet6/0/1
r1 tnsr(config-ospf-if)# ip address * area 0.0.0.2
r1 tnsr(config-ospf)# interface TenGigabitEthernet8/0/0
r1 tnsr(config-ospf-if)# ip address * area 0.0.0.2
r1 tnsr(config-ospf)# interface TenGigabitEthernet6/0/0
r1 tnsr(config-ospf-if)# ip address * cost 5
r1 tnsr(config-ospf-if)# ip address * area 0.0.0.0
r1 tnsr(config-ospf-if)# exit
r1 tnsr(config-ospf)# area 0.0.0.2
r1 tnsr(config-ospf-area)# stub
r1 tnsr(config-ospf-area)# range 10.2.0.0/16
r1 tnsr(config-ospf-area)# exit
r1 tnsr(config-ospf)# exit
r1 tnsr(config-frr-ospf)# enable
r1 tnsr(config-frr-ospf)# exit
r1 tnsr(config)#
```

Configure OSPF on R2

```
r2 tnsr(config)# route dynamic ospf
r2 tnsr(config-frr-ospf)# server
r2 tnsr(config-ospf)# ospf router-id 10.25.0.1
r2 tnsr(config-ospf)# passive-interface GigabitEthernet0/13/0
r2 tnsr(config-ospf)# passive-interface GigabitEthernet0/14/0
r2 tnsr(config-ospf)# interface GigabitEthernet0/13/0
r2 tnsr(config-ospf-if)# ip address * area 0.0.0.25
r2 tnsr(config-ospf-if)# exit
r2 tnsr(config-ospf)# interface GigabitEthernet0/14/0
r2 tnsr(config-ospf-if)# ip address * area 0.0.0.25
r2 tnsr(config-ospf-if)# exit
r2 tnsr(config-ospf)# interface GigabitEthernet3/0/0
r2 tnsr(config-ospf-if)# ip address * cost 5
r2 tnsr(config-ospf-if)# ip address * area 0.0.0.0
r2 tnsr(config-ospf-if)# exit
r2 tnsr(config-ospf)# area 0.0.0.25
r2 tnsr(config-ospf-area)# stub
r2 tnsr(config-ospf-area)# range 10.25.0.0/16
r2 tnsr(config-ospf-area)# exit
r2 tnsr(config-ospf)# exit
r2 tnsr(config-frr-ospf)# enable
r2 tnsr(config-frr-ospf)# exit
```

Notes

This scenario can easily be adjusted to connect with other local routers handling additional networks inside the local ranges. To do so, remove the `stub` configuration for the local area and `passive-interface` directives for interfaces which will communicate with local routers. Then configure the other routers as needed.

27.10 TNSR Remote Office With Existing IPsec Hub

In this example, remote offices with a [Netgate SG-5100](#) running TNSR will be configured for site-to-site VPN to an existing IPsec Head-End at a central headquarters location.

Workers at remote offices will need Direct Internet Access (DIA) and corporate intranet access from their location using IPsec IKEv2 with Pre-Shared Key and secured crypto methods (AES128/SHA1/DH2048). Direct Internet Access also needs to be made available to a guest network through distinct VLANs so that guest and staff devices can be isolated.

The TNSR Remote Office Deployment will be completed in the following high level steps:

27.10.1 Step 1: Prepare for Deployment

Before the deployment can begin, it is important to gather all needed hardware, software, and parameters in advance.

Prerequisites

- Create a reference diagram that shows the logical topology.
- Review the *[TNSR Zero-to-Ping documentation](#)*.
- Whenever possible, get a static public IP address for the remote office TNSR.
- If needed, set up DHCP pass-through (public IP) or set a DMZ host on the ISP Modem.
- Management (web login, admin access) of the ISP modem.

Notes

- IPsec tunnels work best on a static and public IP address.
- NAT-T is not currently supported on TNSR.
- Some *dynamic* ISP services may provide a relatively stable IP address but this will break the IPsec tunnel if the IP address changes.

Reference Diagram

Create and maintain a reference diagram to support the deployment, as shown in the example below:

Remote Office Deployment Parameters

Define and document the deployment parameters for the initial remote office setup, as shown in the example below:

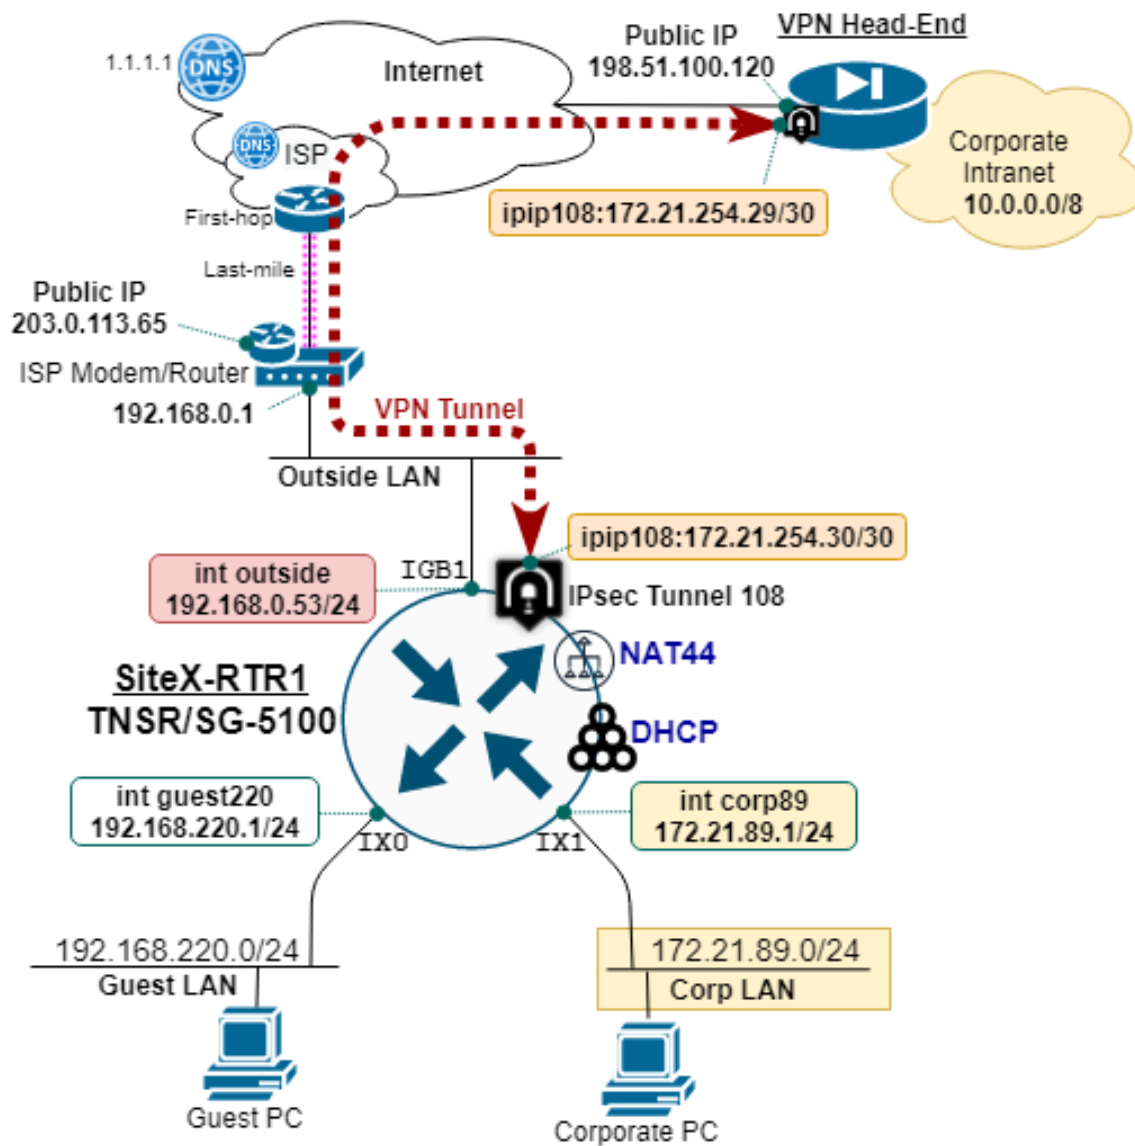


Fig. 10: TNSR remote office reference diagram

Table 16: Base Deployment Parameters

Parameter	Value
TNSR Hostname	siteX-rtr1
TNSR Outside Interface IP	192.168.0.53/24
TNSR Outside Public IP	203.0.113.65/24
Guest LAN Name	guest220
Guest LAN IP	192.168.220.1/24
Guest DHCP Range	100-199
Guest DNS IP	1.1.1.1,9.9.9.9
Corporate LAN Name	corp89
Corporate LAN IP	172.21.89.1/24
Corporate DHCP Range	100-199
Corporate DNS IP	10.10.10.75,1.1.1.1

In this example, the Remote Office is deployed behind an ISP cable modem performing NAT.

- TNSR Outside IP is different than the ISP Modem Public IP address
- ISP cable modem provides NAT for inside devices, which includes the TNSR outside interface
- IPsec uses IP Protocol 50 (ESP) to transfer encrypted payload
 - ISP cable modem may need to be configured for DMZ host set to TNSR outside IP address or inbound ESP traffic may be dropped

IPsec VPN Tunnel Parameters

Define and document the parameters for the corporate IPsec tunnel and IP routing configuration, as shown in the example below:

Table 17: IPsec VPN Tunnel Parameters

Parameter	Value
TNSR Outside IP	192.168.0.53/24
TNSR Public IP	203.0.113.65/24
IPsec Tunnel Peer IP	198.51.100.120
IPsec Tunnel ID	108
IPsec IKEv2 Crypto	AES128/SHA1/DH14
IPsec IKEv2 Authen	PRE-SHARED-KEY
IPsec Child SA Crypto	AES128GCM16/DH14
IPsec Tunnel IP	172.21.254.30/30
IPsec Tunnel Next-hop	172.21.254.29
Corporate IP Block	10.0.0.0/8

27.10.2 Step 2: Initial Setup Tasks

Use the following diagram to support the initial setup tasks.

Ethernet Ports Connections

Use RJ45 ethernet cables to connect the **SG-5100 network ports** to the **outside**, **corp89**, and **guest220** networks.

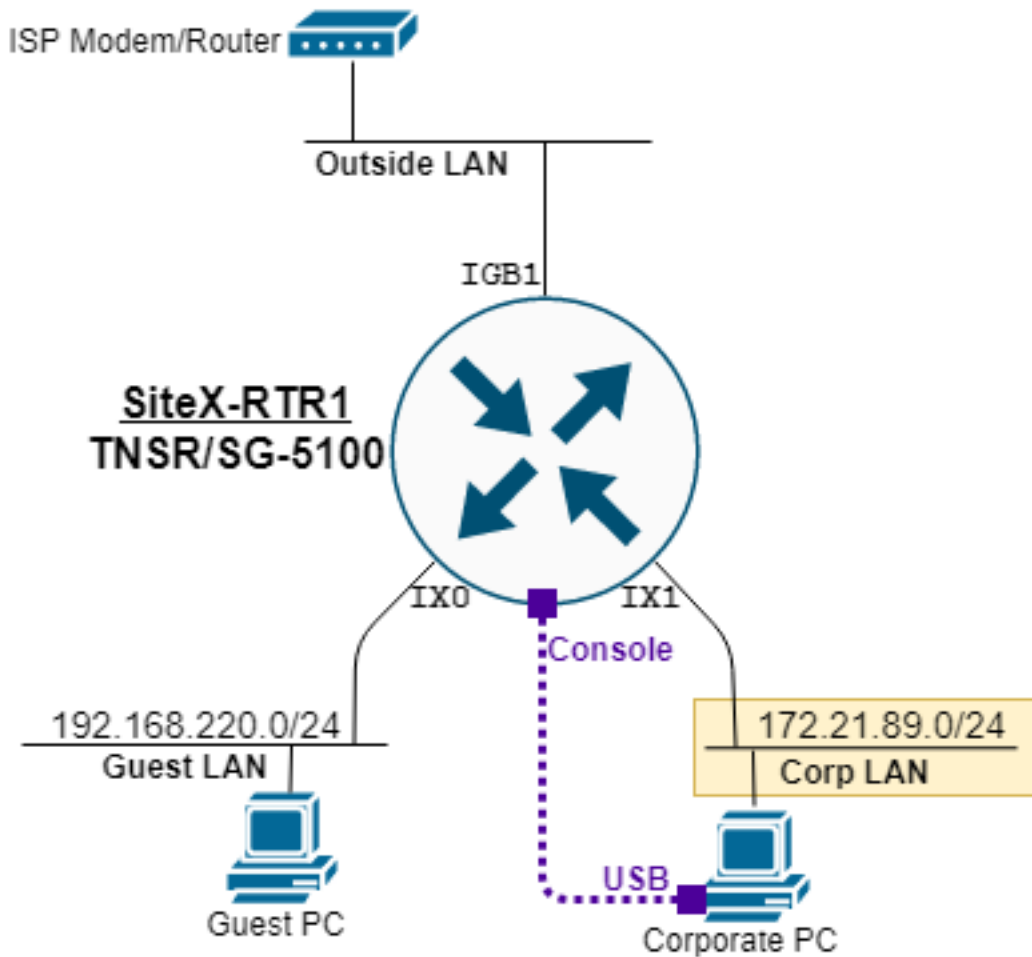


Fig. 11: TNSR remote office setup diagram

Table 18: Assign (and label) SG-5100 Interfaces

Port Label	VPP Name	CentOS Name	Assignment/name
IGB0	n/a	enp3s0	HostOS Interface
IGB1	GigabitEthernet4/0/0	vpp1	outside
IX0	TenGigabitEthernet6/0/0	vpp2	guest220
IX1	TenGigabitEthernet6/0/1	vpp3	corp89
IX2	TenGigabitEthernet8/0/0	vpp4	unassigned
IX3	TenGigabitEthernet8/0/1	vpp5	unassigned

HostOS Interface Notes

- Configuration of the HostOS interface is outside the scope of this guide, see [Edit the Host Interface Configuration](#).
- Do not connect the HostOS interface to the same subnet, or broadcast domain, used by TNSR VPP interfaces as it may produced unexpected results.
- More information at [Host Interfaces](#).

Initial TNSR Setup

Boot up the SG-5100 appliance and [connect to the console](#). Once connected, hit return to get a clear login prompt, and log in with the default TNSR user credentials.

Change Default Password

It is important to change the default password on the `tnsr` user before proceeding:

```
localhost tnsr# shell passwd
Changing password for user tnsr.
Changing password for tnsr.
(current) UNIX password:
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
localhost tnsr#
```

Set Hostname

Set the TNSR hostname and save the configuration:

```
localhost tnsr# config
localhost tnsr(config)# system name sitex-rtr1
sitex-rtr1 tnsr(config)# copy configuration running startup
sitex-rtr1 tnsr(config)# exit
sitex-rtr1 tnsr#
```

Configure Dataplane

Enable selected network and crypto devices on the TNSR dataplane and then restart it, as shown in the config fragment below:

```
dataplane dpdk dev 0000:04:00.0 network name outside
dataplane dpdk dev 0000:06:00.0 network name guest220
dataplane dpdk dev 0000:06:00.1 network name corp89
dataplane dpdk dev 0000:01:00.0 crypto
#
service dataplane restart
```

Note: When enabling the crypto hardware device (QAT) on the console port or during system boot, it is normal for a number of log messages to display as it initializes, these can typically be ignored. For example:

```
[ 836.798096] c3xxxvf 0000:01:01.4: enabling device (0000 -> 0002)
[ 836.804235] DMAR: 64bit 0000:01:01.4 uses identity mapping
[ 836.839343] c3xxxvf 0000:01:01.0: Failed to register crypto algs
[ 836.859227] c3xxxvf: probe of 0000:01:01.0 failed with error -14
[ 836.865313] c3xxxvf 0000:01:01.1: enabling device (0000 -> 0002)
[ 836.871718] DMAR: 64bit 0000:01:01.1 uses identity mapping
[ 836.877853] c3xxxvf 0000:01:01.4: Failed to register crypto algs
[ 836.897244] c3xxxvf: probe of 0000:01:01.4 failed with error -14
```

Inspect an interface with the `show interface` command. As seen below, the **outside** interface is still in **Admin down** state, no IP address is assigned, and no packets have been seen.

```
sitex-rt1 tnsr(config)# show interface outside
Interface: outside
  Admin status: down
  Link down, unknown duplex
  Link MTU: 9206 bytes
  MAC address: 00:90:0b:7c:06:0d
  IPv4 Route Table: ipv4-VRF:0
  IPv6 Route Table: ipv6-VRF:0
  VLAN tag rewrite: disable
  counters:
    received: 0 bytes, 0 packets, 0 errors
    transmitted: 0 bytes, 0 packets, 0 errors
    protocols: 0 IPv4, 0 IPv6
    0 drops, 0 punts, 0 rx miss, 0 rx no buffer
```

Save and Reboot

Save the configuration:

```
copy configuration running startup
```

And then reboot:

```
shell sudo reboot
```

Watch the console logs as the system boots up, then log in as the `tnsr` user with the new password that was set earlier.

27.10.3 Step 3: TNSR IP Configuration

Use the following diagram to support the configuration to provide IP connectivity for the remote office.

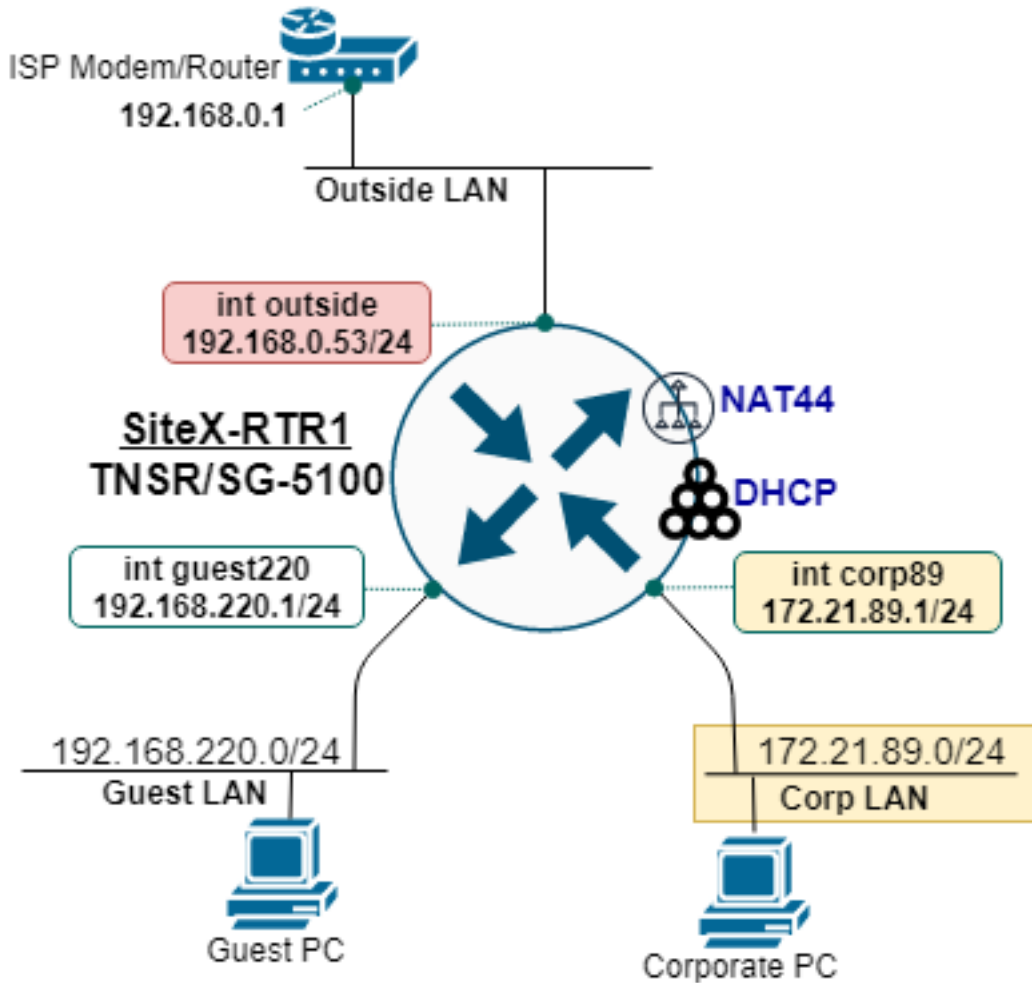


Fig. 12: TNSR remote office IP diagram

There are two inside IP subnets: **corp89** and **guest220**.

- TNSR provides DHCP service to both inside subnets, **corp89** and **guest220**, each using a DHCP address range of 100–199.
- DNS on the **corp89** DHCP includes an internal DNS address for access to the intranet.

Table 19: IP Assignments

Interface Name	IP Address	IP Gateway	NAT
outside	192.168.0.53/24	192.168.0.1	outside/pool
corp89	172.21.89.1/24	n/a	inside
guest220	192.168.220.1/24	n/a	inside

Configure Interface IP Addressing and NAT

With the settings from above, enable IP and NAT on the **corp89**, **guest220**, and **outside** interfaces:

```
# outside interface
interface outside
  ip address 192.168.0.53/24
  ip nat outside
  enable
  exit
#
# corp89 interface
interface corp89
  ip address 172.21.89.1/24
  ip nat inside
  enable
  exit
#
# guest220 interface
interface guest220
  ip address 192.168.220.1/24
  ip nat inside
  enable
  exit
#
# enable outside interface IP as NAT pool (PAT)
nat pool interface outside
```

Use `show interface` and `show interface ip` to inspect the status of the interfaces.

Configure DHCP Pools

Configure the DHCP server for both inside interfaces, **corp89** and **guest220**:

```
dhcp4 server
  interface listen corp89
  subnet 172.21.89.0/24
    pool 172.21.89.100-172.21.89.199
  exit
  interface corp89
    option routers
      data 172.21.89.1
    exit
    option domain-name-server
      data 10.10.10.75,1.1.1.1
    exit
  exit
# next DHCP for guest220
interface listen guest220
subnet 192.168.220.0/24
  pool 192.168.220.100-192.168.220.199
  exit
interface guest220
  option routers
    data 192.168.220.1
  exit
  option domain-name-server
```

(continues on next page)

(continued from previous page)

```
data 1.1.1.1,9.9.9.9
exit
exit
exit
# enable the DHCP4 server
dhcp4 enable
```

Devices on the **corp89** and **guest220** interfaces should now be able to pull an IP address from DHCP and use the internet via NAT.

Configure IP Default Route

When using a static IP configuration on the outside interface, it is necessary to add an IP Default Route.

```
route ipv4 table ipv4-VRF:0
route 0.0.0.0/0
next-hop 0 via 192.168.0.1 outside
exit
exit
```

Devices on **corp89** and **guest220** should now be able to ping the internet.

27.10.4 Step 4: Protect the Outside Interface with ACLs

Before proceeding with the VPN IPsec site-to-site tunnel, it is critical to apply Access Control Lists (ACLs) to the **outside** interface.

The reason being, if the **outside** interface is exposed to the internet, it will be frequently probed by bots attempting to login using weak credentials. This can be seen by inspecting `lastb` from the `root` user:

```
sudo lastb | head -100
```

Multiple ACLs can be applied to an input or output queue on an interface, as ordered by sequence. This offers a modular and scalable approach to ACLs for a given interface.

Output ACL - Reflect

The reflect ACL is a special action that permits output traffic and also permits the return, or input, traffic to match the IP flow.

Create an ACL named `outbound-reflect` and apply it:

```
acl outbound-reflect
rule 5
desc reflect permit outbound traffic AND permit return traffic on input
action reflect
ip-version ipv4
exit
exit
#
# Apply to interface as output ACL
interface outside
access-list output acl outbound-reflect sequence 10
exit
```

Input ACL - DHCP Response

If using `dhcp client ipv4` on the outside interface, be sure to permit DHCP responses on destination port UDP 68 by creating an ACL named `dhcp-outside` and applying it:

```
acl dhcp-outside
  rule 1
    desc DHCP Response to client on outside interface
    action permit
    ip-version ipv4
    protocol udp
    destination port 68
    exit
  exit
#
# Apply ACL to interface Access-List
interface outside
  access-list input acl dhcp-outside sequence 5
  exit
```

Input ACL - SSH-Outside

To only permit inbound SSH access from specified IP hosts, create an ACL rule named `ssh-outside` and apply it. In this example, **rule 221** permits a block of IP addresses from corporate headquarters and **rule 222** permits a single IP address for assistance from a service provider:

```
acl ssh-outside
  rule 221
    desc Allow SSH from HQ
    action permit
    ip-version ipv4
    protocol tcp
    destination port 22
    source address 198.51.100.0/24
    exit
  rule 222
    desc Allow SSH from service provider
    action permit
    ip-version ipv4
    protocol tcp
    destination port 22
    source address 192.0.2.88/32
    exit
  exit
#
# Apply to outside interface as input ACL
interface outside
  access-list input acl ssh-outside sequence 10
  exit
```

Then validate that only the specified IP addresses are able to SSH to the outside IP address of TNSR.

Note: A NAT static mapping from outside addresses to inside addresses on port 22 (SSH) may be required.

```
nat static mapping tcp local 172.21.89.1 22 external outside 22 out-to-in-only
```

NAT port forwarding is covered in *Step 6: Port Forwarding with NAT*.

Input ACL - IPsec-Outside

Configure an ACL, named `ipsec-outside`, to permit three (3) types of IPsec traffic:

1. IP Protocol 50; Encapsulating Security Payload (ESP)
2. IP Protocol UDP; Destination Port 500: IKEv2 Message Exchange
3. IP Protocol UDP; Destination Port 4500: NAT-T (**not supported in current TNSR release**)

```
acl ipsec-outside
  rule 11
    desc Permit ESP
    action permit
    ip-version ipv4
    source address 198.51.100.120/32
    protocol 50
    exit
  rule 12
    desc Permit UDP500 - IKEv2
    action permit
    ip-version ipv4
    source address 198.51.100.120/32
    protocol udp
    destination port 500
    exit
  rule 13
    desc Permit UDP4500 - NAT-T Future
    action permit
    ip-version ipv4
    source address 198.51.100.120/32
    protocol udp
    destination port 4500
    exit
  exit
# Apply ACL to interface Access-List
interface outside
  access-list input acl ipsec-outside sequence 20
  exit
```

27.10.5 Step 5: Corporate VPN with IPsec Tunnel

An IPsec IKEv2 VPN tunnel is configured between the the remote office and the VPN head end at the corporate office.

IPsec Deployment Parameters

The IPsec tunnel is built with the following parameters:

- IP address of each tunnel endpoint, both the remote office and VPN head end.
- Local and Remote Identity and Pre-Shared Key (PSK).
- Compatible Phase 1 Proposal: AES128, SH1, DH14.

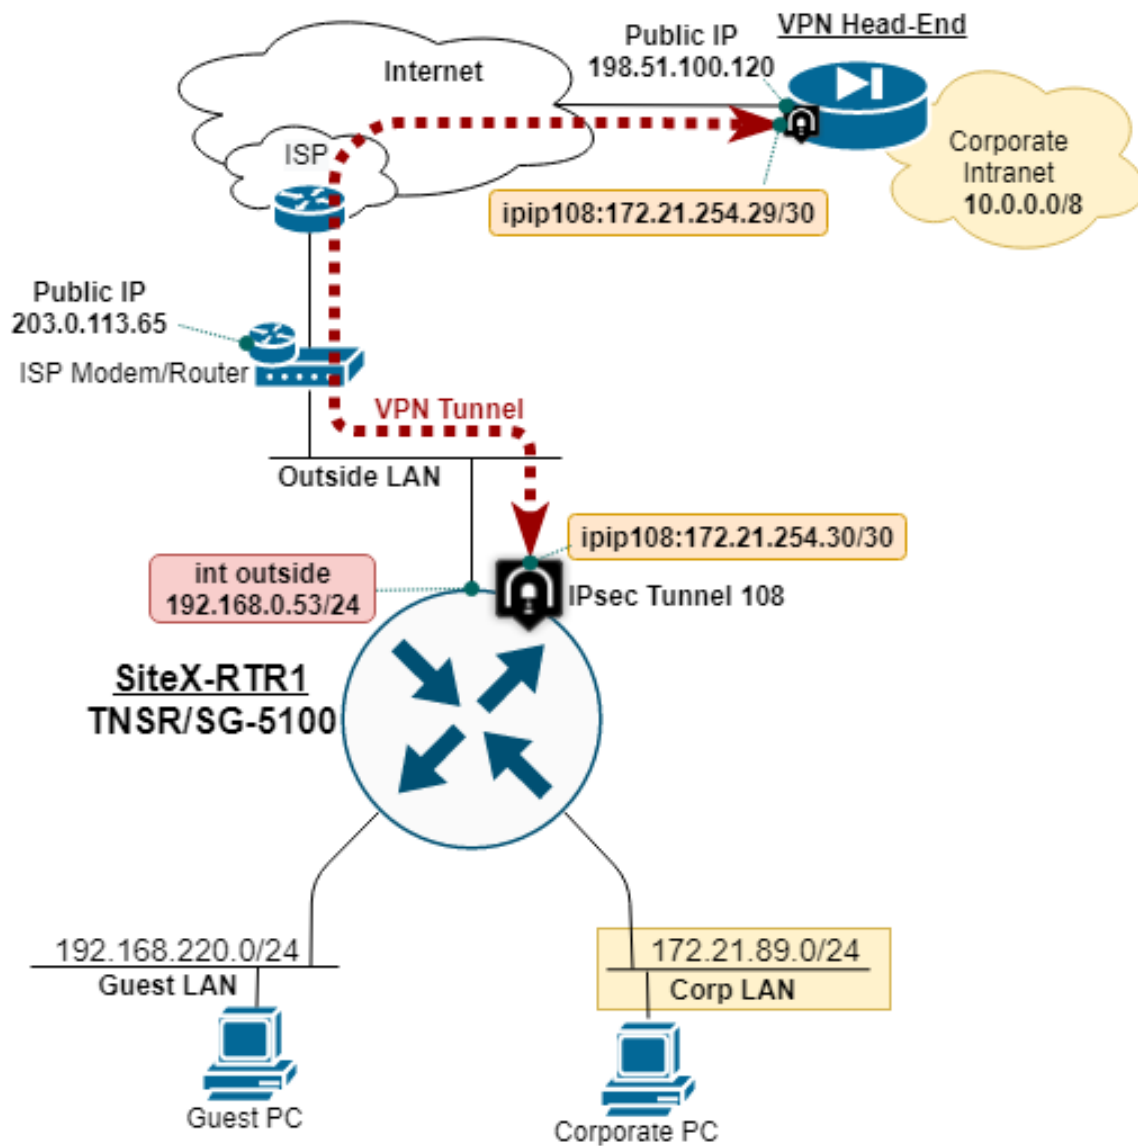


Fig. 13: TNSR remote office IPsec VPN diagram

- Compatible Phase 2 Proposal: AES128GCM16, DH14.
- IP address assigned to the tunnel interface, typically a /30 subnet.
- IP route to direct corporate traffic over IPsec tunnel via next-hop IP.

The parameters for this deployment were captured at [IPsec VPN Tunnel Parameters](#).

ISP NAT Inbound Forwarding of ESP Packets

In this example, the Remote Office TNSR is behind NAT on the ISP cable mode, and using a Private IP address.

- TNSR Outside IP: 192.168.0.53
- TNSR Public IP: 203.0.113.65

The DMZ Host feature on the ISP cable modem is used to direct inbound ESP, IP Protocol 50, traffic to TNSR.

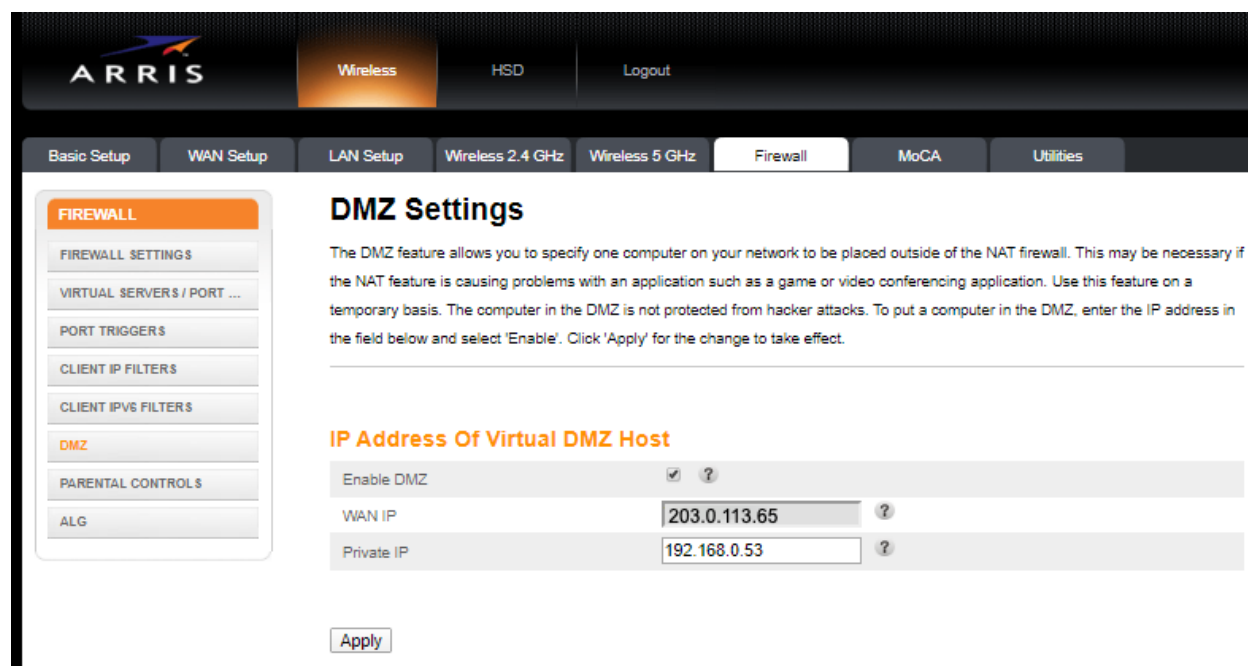


Fig. 14: ISP Modem - DMZ Host

When possible, use a Public IP address directly on TNSR, without external NAT.

TNSR NAT Inbound Forwarding of ESP Packets

ESP packets use IP Protocol 50 for data transfer over the secured tunnel. TNSR NAT provides the ability to configure static mappings for IP protocols ICMP, UDP, and TCP. If using NAT on TNSR, and IPsec on the outside interface, you must enable global NAT forwarding to permit input ESP packets on the outside interface.

IKEv2 message exchange use UDP destination port 500 which if NAT static mapped will permit IKEv2 negotiation, but the ESP packets are needed to carry traffic.

```
nat global-options nat44 forwarding true
```

Configure IPsec Tunnel

Configure the IPsec tunnel on the remote office TNSR:

```
ipsec tunnel 108
  local-address 192.168.0.53
  remote-address 198.51.100.120
  crypto config-type ike
  crypto ike
    version 2
    lifetime 28800
    proposal 1
      encryption aes128
      integrity sha1
      group modp2048
    exit
  identity local
    type address
    value 203.0.113.65
  exit
  identity remote
    type address
    value 198.51.100.120
  exit
  authentication local
    round 1
      type psk
      psk AD78PPQMP00
    exit
  exit
  authentication remote
    round 1
      type psk
      psk AD78PPQMP00
    exit
  exit
  child 1
    lifetime 3600
    proposal 1
      encryption aes128gcm16
      no integrity
      group modp2048
    exit
  exit
  exit
  exit
#
```

Configure IP Tunnel Address

Configure a static IP route to direct corporate traffic over the IPsec tunnel via the next-hop:

```
interface ipip108
  ip address 172.21.254.30/30
  exit
#
```

To test, ping the other end of the IPsec tunnel with:

```
ping 172.21.254.29 source 172.21.254.30 count 5
```

Configure IP Route to Corporate

Configure a static IP route to direct corporate traffic over the IPsec tunnel via the next-hop on the IPsec tunnel:

```
route ipv4 table ipv4-VRF:0
route 10.0.0.0/8
    next-hop 0 via 172.21.254.29 ipip108
exit
exit
#
```

Save changes:

```
configuration copy running startup
```

27.10.6 Step 6: Port Forwarding with NAT

This section shows an example of a single inbound *Network Address Translation* (NAT) port forward, also known as NAT pinholes or port mapping, to access an internal web host.

Although VPN connections are preferred, sometimes it is desired, or necessary, to provide direct access to an internal networked device.

Use good judgement AND action when permitting access to the network connected device from the outside.

Good practices on internet accessible devices:

- Change all default passwords
- Only provide access to needed ports
- Update firmware to latest and periodically
- Include source addresses in your ACLs (access control lists) rules whenever possible
- Don't use low grade-dog-food networked products - if the vendor does not put their name on the product, definitely avoid it

With due diligence (see above), NAT port forwarding can be used to provide specific outside access to inside networked devices.

In this recipe, we setup NAT port forwarding to an internal system with a web interface to provide remote access to a support technician.

NAT Port Forwarding

Define NAT port forwarding rule:

```
nat static mapping tcp local 172.21.89.12 8443 external outside 8443
```

Permit Port Forward Traffic with ACL

Traffic that is port forwarded by NAT must also be permitted by the outside access control list (ACL). The ACL is created and applied to input queue access-list on the outside interface.

```
acl http-outside
  rule 10
    desc Permit from Corp to TCP-8443
    action permit
    ip-version ipv4
    source address 198.51.100.0/24
    protocol tcp
    destination port 8443
    exit
  exit
#
# Apply ACL to interface Access-List
interface outside
  access-list input acl http-outside sequence 101
  exit
```

The internal web host is should now be accessible from permitted IP addresses. Test to confirm that the configuration is correct.

27.10.7 Step 7: SNMP Monitoring

SNMP is used to monitor the remote office from an external location.

Configure SNMP Server

Configure and enable the SNMP server on TNSR, as shown in the example below:

```
snmp community community-name PlzzaGuy source 198.51.100.0/24 security-name tnsrsnmp
snmp group group-name ROGroup security-name tnsrsnmp security-model v2c
snmp view view-name systemview view-type included oid .1
snmp access group-name ROGroup prefix exact model any level noauth read systemview
↪write none
#
server snmp enable
```

Enable Port Forward to SNMP

If needed, setup static mapping from an outside UDP port to the SNMP port on TNSR:

```
nat static mapping udp local 192.168.0.53 161 external outside 161
```

Permit SNMP Polling with ACL

To permit access to the SNMP server from the outside, create an ACL rule named `snmp-outside` and apply it:


```
acl snmp-outside
  rule 11
    desc Permit to SNMP
    action permit
    ip-version ipv4
    protocol udp
    source address 198.51.100.0/24
    destination port 161
    exit
  exit
#
# Apply ACL to interface Access-List
interface outside
  access-list input acl snmp-outside sequence 102
  exit
```

ADVANCED CONFIGURATION

The items in this section can be used to control lower-level behavior of the dataplane and host operating system in various ways. These can help to increase performance and efficiency for large workloads.

28.1 Dataplane Configuration

For the majority of cases the default dataplane configuration is sufficient, but certain cases may require adjustments. These are often covered in more detail throughout the documentation, and relevant sections will be linked where appropriate.

These commands are all available in `config` mode (*Configuration Mode*).

Warning: The dataplane service requires a restart to enable configuration changes described in this section. After making changes, restart the dataplane from `config` mode using the following command:

```
tnsr# configure
tnsr(config)# service dataplane restart
```

28.1.1 Buffers

The commands in this section control the amount of memory pre-allocated by the dataplane for buffers.

Buffers per NUMA

Systems with multiple CPU sockets and Non-uniform memory access (NUMA) capabilities may need specific tuning to ensure that enough buffer space is available for the number of separate NUMA nodes. The number of NUMA nodes is typically the number of populated CPU sockets. Specifically, the scenarios which require tuning typically involve a large number of interfaces combined with multiple CPU worker threads.

Note: This refers to separate hardware CPUs, not a single CPU with multiple cores.

The dataplane buffers `buffers-per-numa <buffers-per-numa>` command allocates the given number of buffers for each CPU socket (e.g. 16384).

Default Data Size

The dataplane buffers `default-data-size` `<default-data-size>` controls the default size of each buffer, in bytes (e.g. 2048).

28.1.2 CPU Workers and Affinity

The dataplane has a variety of commands to fine-tune how it uses available CPU resources on the host. These commands control CPU cores TNSR will use, both the number of cores and specific cores.

See also:

Cores defined here may also be pinned to interface receive (RX) queues, provided that cores are defined using either the `corelist-workers` or `coremask-workers` methods. See [Interface Configuration Options](#) for details.

Worker Configuration

dataplane cpu corelist-workers `<first>` `[- <last>]` Defines a specific list of CPU cores to be used by the dataplane. The command supports adding single cores to the list at a time, or ranges of cores. Run the command multiple times with different core numbers or ranges to define the full list of cores to utilize. When removing items with `no`, the command accepts a specific core to remove from the list.

dataplane cpu coremask-workers `<mask>` Similar to `corelist-workers`, but the cores are defined as a hexadecimal mask instead of a list. For example, `0x0000000000C0000C`

dataplane cpu main-core `<n>` Assigns the main dataplane process to a specific CPU core.

dataplane cpu scheduler-policy (`batch|fifo|idle|other|rr`) Defines a specific scheduler policy for worker thread processor usage allocation

batch Scheduling batch processes. Uses dynamic priorities based on `nice` values in the host OS, but always gives the thread a small scheduling penalty so that other processes take precedence.

fifo First in-first out scheduling. Will preempt other types of threads and threads with a lower priority.

idle Scheduling very low priority jobs.

other Default Linux time-sharing scheduling. Uses dynamic priorities based on `nice` values in the host OS, similar to `batch` but without the built-in penalty.

rr Round-robin scheduling. Similar to `fifo` but each thread is time-limited

dataplane cpu scheduler-priority `<n>` For the `fifo` and `rr` scheduler policies, this number sets the priority of processes for the dataplane. It can be any number between 1 (low) and 99 (high).

dataplane cpu skip-cores `<n>` Defines the number of cores to skip when creating additional worker threads, in the range of 1 to the highest available core number. The first `<n>` cores will not be used by worker threads.

Note: This does not affect the core used by the main thread, which is set by `dataplane cpu main-core` `<n>`.

Warning: This option is incompatible with *interface RX queue core pinning*. To utilize interface RX queue core pinning, define a list of cores using either `corelist-workers` or `coremask-workers` instead.

dataplane cpu workers <n> Defines the number of worker threads to create for the dataplane.

Note: The number of worker threads is in addition to the main process. For example, with a worker count of 4, the dataplane will use one main process with four worker threads, for a total of five threads.

Warning: This option is incompatible with *interface RX queue core pinning*. To utilize interface RX queue core pinning, define a list of cores using either `corelist-workers` or `coremask-workers` instead.

Worker Example

This example sets four additional worker threads, and instructs the dataplane to skip one core when assigning worker threads to cores:

```
tnsr(config)# dataplane cpu workers 4
tnsr(config)# dataplane cpu skip-cores 1
tnsr(config)# service dataplane restart
```

Worker Status

The `show dataplane cpu threads` command displays the current dataplane process list, including the core usage and process IDs. This output corresponds to the example above:

```
tnsr(config)# show dataplane cpu threads
ID Name      Type    PID   LCore Core Socket
-- --
0 vpp_main    2330    1     0     0
1 vpp_wk_0    workers 2346    2     2     0
2 vpp_wk_1    workers 2347    3     3     0
3 vpp_wk_2    workers 2348    4     4     0
4 vpp_wk_3    workers 2349    5     5     0
```

The output includes the following columns:

id Dataplane thread ID.

name Name of the dataplane process.

type The type of thread, which will be blank for the main process.

pid The host OS process ID for each thread.

LCore The logical core used by the process.

Core The CPU core used by the process.

Socket The CPU socket associated with the core used by the process.

28.1.3 DPDK Configuration

Commands in this section configure hardware settings for DPDK devices.

dataplane dpdk dev <pci-id> (cryptolcrypto-vf|network) [num-rx-queues [<rq>]] [num-tx-queues [<tq>]] [num-rx-desc [<rd>]] [num-tx-desc [<td>]] Configures a specific dataplane device for use by TNSR.

cryptolcrypto-vf Configures QAT devices for cryptographic acceleration. See *Setup QAT Compatible Hardware* for details.

network Configures network interface devices, see *Setup NICs in Dataplane* for details.

num-rx-queues [<rq>] num-tx-queues [<tq>] Receive and transmit queue sizes for this device.

num-rx-desc [<rd>] [num-tx-desc [<td>]] Receive and transmit descriptor sizes for this device. Certain network cards, such as Fortville models, may need the descriptors set to 2048 to avoid dropping packets at high loads.

dataplane dpdk iova-mode (palva) Manually configures the IO Virtual Addresses (IOVA) mode used by DPDK when performing hardware IO from user space. Hardware must use IO addresses, but it cannot utilize user space virtual addresses directly. These IO addresses can be either physical addresses (PA) or virtual addresses (VA). No matter which mode is set, these are abstracted to TNSR as IOVA addresses so it does not need to use them directly.

In most cases the default IOVA mode selected by DPDK is optimal.

Warning: When the `vfio-pci` UIO driver is active, IOVA must be explicitly set to `pa` since the automatic selection of `va` will fail with that driver.

See also:

For more detail on IOVA, consult [the DPDK documentation](#).

pa Physical Address mode. IOVA addresses used by DPDK correspond to physical addresses, and both physical and virtual memory layouts match. This mode is safest from the perspective of the hardware, and is the mode chosen by default. Most hardware supports PA mode at a minimum.

The primary downside of PA mode is that memory fragmentation in physical space must also be reflected in virtual memory space.

va Virtual Address mode. IOVA addresses do not follow the layout of physical memory; Physical memory is changed to match the virtual memory instead. Because virtual memory appears as one continuous segment, large memory allocations are more likely to succeed.

The primary downside of VA mode is that it relies on kernel support and the availability of IOMMU.

dataplane dpdk no-tx-checksum-offload Disables transmit checksum offloading for network devices.

dataplane dpdk no-multi-seg Disables multi-segment buffers for network devices. Can improve performance, but disables jumbo MTU support. Recommended for Mellanox devices.

dataplane dpdk num-crypto-mbufs <num> Sets the number of memory buffers used by the dataplane for cryptographic tasks, in the range 1–4294967295. Higher values can improve throughput when the dataplane encrypt/decrypt nodes are processing data.

dataplane dpdk uio-driver [<driver-name>] Configures the UIO driver for interfaces. See *Setup NICs in Dataplane*.

dataplane dpdk vdev <sw-dev-type> Defines a software device to be used by the dataplane, such as:

aesni_gcm AESNI GCM cryptodev

aesni_mb AESNI multibuffer cryptodev

28.1.4 Memory

Commands in this section configure memory allocation for the dataplane.

dataplane (i)l(p)6 heap-size [<size>] Defines the amount of memory to be allocated for the dataplane FIB. The default is 32MB. For more information, see [Working with Large BGP Tables](#).

Note: When tuning this value, also consider increasing the [Statistics Segment](#) heap-size.

dataplane ip6 hash-buckets [<size>] Defines the number of IPv6 forwarding table hash buckets. The default is 65536.

28.1.5 NAT

Commands in this section configure dataplane NAT behavior.

dataplane nat dslite-ce Enables DS-Lite CE mode.

dataplane nat max-translations-per-user <n> Defines the number of NAT translation entries to allow for each IP address. The default value is 10240, but it can be set to any integer value between 1–262144. The ideal value depends entirely on the environment and number of sessions per IP address involved in NAT. This includes traffic sourced from TNSR itself address as well, not only internal source IP addresses.

dataplane nat mode (deterministic|endpoint-dependent|simple) Configures the operating NAT mode. See [Dataplane NAT Modes](#).

dataplane nat mode-options simple (out2in-dp|static-mapping-only) Configures options for the NAT mode. See [Dataplane NAT Modes](#).

NAT Memory

Memory available for NAT functions can also be tuned to scale for larger operations. The following parameters are available:

dataplane nat user hash buckets <size> Number of buckets in NAT user lookup hash table. Can be from 1–65535, default 128.

dataplane nat user hash memory <size> Memory size of NAT user lookup hash table. Can be from 1–4294967295, default 67108864 (64MiB).

dataplane nat translation hash buckets <size> Number of buckets in session lookup hash tables. Can be from 1–65535, default 1024.

dataplane nat translation hash memory <size> Memory size of session lookup hash tables. Can be from 1–4294967295, default 134217728 (128MiB).

With the default user hash memory, each user hash bucket can contain approximately 512 active elements (“sessions”). To determine the total number of supported NAT sessions, multiply:

128 (user hash buckets) x 512 (max elements per user hash bucket) = 65,536 NAT sessions

To support more than 65,536 NAT sessions, NAT user hash memory must be increased along with NAT user hash buckets. In the case of user hash, a single client may consume many elements/sessions, limited by the `nat max-translations-per-user` option mentioned previously in this section.

The `nat translation` options are similar to the `nat user` options, but are utilized for endpoint-dependent NAT lookup tables.

28.1.6 Statistics Segment

These commands configure the statistics segment parameters for the dataplane. This feature enables local access to dataplane statistics via shared memory.

See also:

For more information on how to make use of this feature, see the VPP documentation for the example [stat_client](#).

dataplane statseg heap-size <heap-size>[kKmMgG] Size of shared memory allocation for stats segment, in bytes. This value can be suffixed with K (kilobytes), M (megabytes), or G (gigabytes) in upper or lowercase. Default value is 96M.

Note: This value may need to be increased to accommodate large amounts of routes in routing tables. The default value of 96M can safely accommodate approximately one million routes.

The statistics segment is used to maintain counters for routes, and when multiple worker threads are used, these counters are maintained in each thread. Each counter consumes 16 bytes, and there are two counters for each route. When computing these memory requirements, also keep in mind that the main thread counts in addition to each worker thread. For example, with two worker threads, there are actually three threads total.

The total memory required for route counters alone will be: `<routes> * <threads> * 2 counters * 16 Bytes`. Additionally, when new memory is being allocated, it must be in a contiguous segment approximately 1.5x the size calculated above. This can negatively impact memory allocation in cases where usage of the statistics segment has become fragmented after repeated allocations and reallocations. All these factors combined mean that when using a large number of routes with multiple worker threads, this value should be given a generous increase over expected normal values.

The dataplane may crash and state that it is out of memory if this value is set too low.

dataplane statseg per-node-counters enable Enables per-graph-node performance statistics.

dataplane statseg socket-name <socket-name> Absolute path to UNIX domain socket for stats segment. The default path is `/run/vpp/stats.sock`.

28.2 Host Memory Management Configuration

TNSR has commands to tweak a few common host OS memory management parameters.

These are:

sysctl vm nr_hugepages <u64> Virtual memory, maximum number of huge pages. This controls allocations of huge areas of contiguous memory, which is used to keep TNSR in memory, rather than

swapping. Each huge page is 2MB by default, and the default number of huge pages is 1024. Multiplying the values yields 2GB of RAM set aside. This value can be tweaked lower for systems with less memory or higher for systems with more available memory and larger workloads.

sysctl vm max_map_count <u64> Virtual memory, maximum map count. This controls the number of memory map areas available to a given process. With workloads requiring larger amounts of memory, this may need increased to allow sufficient levels of memory allocation operations to succeed. The default value is 3096.

sysctl kernel shmmem <u64> Maximum size, in bytes, of a single shared memory segment in the kernel. Default value is 2147483648 (2GB).

To view the current active values of these parameters, use `show sysctl`:

```
tnsr# show sysctl
vm/nr_hugepages = 1024
vm/max_map_count = 3096
kernel/shmmax = 2147483648
```

28.3 IP Reassembly

IP reassembly deals with packet fragments, which can be problematic for certain routing, access control, and related tasks which require packet header information. When a packet is fragmented, only the first fragment carries full header information such as TCP/UDP port data. This can lead to problems with processing fragmented packets which involve NAT, for example, which requires that port data.

28.3.1 IP Reassembly Types

TNSR supports two types of IP reassembly: Full Reassembly and Shallow Virtual Reassembly.

Full Reassembly

Full reassembly is more common and the type of reassembly found most often networking platforms. When performing full reassembly, the router waits until all fragments of a packet arrive, and then it acts on that packet to apply ACLs, NAT, and so on, before delivering the packet further.

This means that fragmented packets must be held in a buffer, consuming memory, for a long enough time to allow later fragments to arrive in a reasonable window. This not only consumes memory required for the buffer, but adds latency since all fragments must arrive before the entire packet can be acted upon.

Shallow Virtual Reassembly

Shallow Virtual Reassembly (SVR) does not reassemble fragmented packets, but retains the L4 information so that later fragments can have operations applied using the L4 data from the initial fragment. For example, MAP-T and MAP-E BR rely on the destination port of incoming IPv4 packets to determine the destination CE. Tracking the ports from the first fragment and populating that data into the buffer opaque data of later fragments allows MAP to figure out the correct destination address for future fragments without having to reassemble the entire packet.

Note: Some features of TNSR, such as NAT and MAP, require SVR to function and when those features are active, TNSR will implicitly enable SVR.

28.3.2 IP Reassembly Options

IP reassembly is enabled on a per-interface basis using the `ip reassembly enable (IPv4)` or `ipv6 reassembly enable (IPv6)` commands from within `config-interface` mode. The type may also be set to `full` or `virtual` in a similar manner, see [Interface Configuration Options](#) for details.

The fragment reassembly behavior in TNSR can be fine-tuned globally using the commands `ip reassembly <type> <address-family> <name> <value>`.

Type Sets options for either `full` or `virtual` reassembly types. See [IP Reassembly Types](#) for details on how these modes operate.

full Options used when performing full reassembly of packet fragments.

virtual Options used when performing Shallow Virtual Reassembly for retaining fragment information.

Address Family Sets whether these parameters refer to `ipv4` or `ipv6`.

Name and Value

expire-walk-interval <expire-walk-interval-ms> The interval, in milliseconds, at which TNSR will check for fragments to expire. Decreasing this will consume more CPU time but will allow TNSR to be more proactive in cleaning up expired fragments. Increasing this will allow expired fragments to be held longer, but may be more likely to overrun the value of `max-reassemblies`. Default value is 10000 (10 seconds).

max-reassemblies <max> The maximum number of active reassemblies TNSR will maintain at any given time. Increasing this value will consume more resources, but it will also allow TNSR to reassemble a greater number of fragments at a time. Default value is 1024.

Note: In `full` reassembly mode, new fragment reassembly sessions are not created by TNSR when the limit has been reached, which enforces this limit.

In `virtual` mode, this limit is enforced by removing older sessions. If a fragment for a new packet arrives on an interface and a new reassembly session cannot be created because this limit has been reached, the last created session will be erased and a new one will be created for the current fragment.

timeout <timeout-ms> The timeout value, in milliseconds, after which TNSR will consider a reassembly attempt expired. Increasing this value will cause fragments to be held longer waiting on the remaining pieces, which means they are more likely to be successfully reassembled on slower networks, at the cost of consuming more resources. Default value is 100 milliseconds. When this value is increased, the `max-reassemblies` value may also need increased to accommodate the higher volume of fragments that TNSR will need to hold.

28.3.3 IP Reassembly Status

To view the current IP reassembly status, use the `show ip reassembly [(full|virtual) [(ipv4|ipv6)]]` command.

On its own, `show ip reassembly` will display the status of all types for both IPv4 and IPv6.

```
tnsr# show ip reassembly
full ipv4 reassembly:
  Timeout:                100 ms
  Expire walk interval:    10000 ms
  Maximum reassemblies:    1024
full ipv6 reassembly:
  Timeout:                100 ms
  Expire walk interval:    10000 ms
  Maximum reassemblies:    1024
virtual ipv4 reassembly:
  Timeout:                100 ms
  Expire walk interval:    10000 ms
  Maximum reassemblies:    1024
virtual ipv6 reassembly:
  Timeout:                100 ms
  Expire walk interval:    10000 ms
  Maximum reassemblies:    1024
```

To limit the output by type, add the `full` or `virtual` keyword, optionally followed by `ipv4` or `ipv6` to view the status only for a specific address family.

TROUBLESHOOTING

This section contains commonly encountered issues with TNSR and methods to resolve them.

29.1 Ping and traceroute do not function without host OS default route

Utilities such as `ping` and `traceroute` will send traffic using the host OS routing table by default unless a specific source address is passed to the command. See *Diagnostic Routing Behavior* for details.

29.2 Unrecognized routes in a routing table

TNSR automatically populates routing tables with necessary entries that may not appear to directly correspond with manually configured addresses. See *Common Routes* for details.

29.3 Services do not receive traffic on an interface with NAT enabled

When NAT is enabled, by default TNSR will drop traffic that doesn't match an existing NAT session or static NAT rule. This includes traffic for services on TNSR such as IPsec and BGP. To allow this traffic, see *NAT Forwarding*.

29.4 NAT session limits / “Create NAT session failed” error

The default limit for NAT sessions per IP address in the dataplane is 10240. If the number of sessions from a client IP address, including TNSR itself, exceeds that value, then new connections will fail. This value can be changed by using the `dataplane nat max-translations-per-user` command as described in *Advanced Dataplane Configuration: NAT*.

29.5 ACL rules do not match NAT traffic as expected

When NAT is active, ACL rules are always processed before NAT on interfaces where NAT is applied, in any direction. This behavior is different from some other products, such as pfSense. See *ACL and NAT Interaction* for details.

29.6 Some Traffic to the host OS management interface is dropped

TNSR includes a default set of Netfilter rules which secure the management interface. Only certain ports are allowed by default. See *Default Allowed Traffic* for details. To allow more traffic, create host ACLs as described in *Host ACLs*.

29.7 Locked out by NACM Rules

If TNSR access is lost due to the NACM configuration, access can be regained by following the directions in *Regaining Access if Locked Out by NACM*.

29.8 How to gain access to the root account

By default, the `root` account has interactive login disabled, which is the best practice. This can be changed by resetting the root password using `sudo` from another administrator account, or in the ISO installer. See *Default Accounts and Passwords* for details.

29.9 Console Messages Obscure Prompts

When connected to the console of a TNSR device, such as the serial console, the kernel may output messages to the terminal which obscure prompts or other areas of the screen. This is normal and an expected effect when using the console directly.

To work around this intended behavior, use one of the following methods:

- Press `Ctrl-L` to clear or redraw the screen without the messages.
- Press `Enter` to receive a new prompt.
- Run `sudo dmesg -D` from a shell prompt or with the TNSR `shell` command, which will disable kernel output to all consoles.
- Connect to the TNSR device using SSH instead of the console.

29.10 OSPF Neighbors Stuck in ExStart State

When attempting to form an adjacency between two OSPF (*Open Shortest Path First v2 (OSPF)*) neighbors, if the neighbor status appears to be stuck in the `ExStart` state, the most likely cause is an MTU mismatch between the routers.

To solve this problem, adjust the MTU values of the interfaces actively participating in OSPF on all routers to match. If this is not possible, try using the `mtu-ignore` option on active OSPF interfaces.

29.11 Large packets fail to pass over IPsec

Encapsulated packets which are larger than the `default-data-size` buffer will be dropped by the dataplane. On older versions of TNSR this may even result in a dataplane crash. The size of this buffer is 2048 by default, which will pass packets up to approximately 2000 bytes. Since the default MTU is 1500, this issue is not apparent in many

cases. However, when attempting to pass jumbo frames over IPsec, this becomes a problem. To pass 9000 byte frames over IPsec, increase the buffer size to 16384 and restart the dataplane.

```
tnsr(config)# dataplane buffers default-data-size 16384
tnsr(config)# service dataplane restart
```

29.12 IPsec packets are dropped or fail to pass with QAT enabled

There is a known incompatibility between QAT and VT-d on some platforms which can prevent IPsec traffic from passing when QAT acceleration is enabled. See [Disable VT-d in the BIOS](#) for details.

29.13 Console DMA / PTE Read access Error Messages

Errors similar to the following may appear on the console:

```
[110772.063766] DMAR: [DMA Read] Request device [04:01.0] fault addr_
↪406482000
   [fault reason 06] PTE Read access is not set
[110773.059440] DMAR: DRHD: handling fault status reg 102
```

The cause is likely an incompatibility between an enabled QAT device and VT-d in the BIOS. See [Disable VT-d in the BIOS](#) for details.

29.14 Diagnosing Service Issues

If a service will not stay running and the logs indicate that it is crashing, additional debugging information can be obtained from core dumps.

By default, core dumps are disabled for services. These can be individually enabled as needed by the following command:

```
tnsr(config)# service (backend|bgp|dataplane|dhcp|http|ike|ntp|restconf|unbound)
                coredump (enable|disable)
```

The resulting core files will be written under `/var/lib/systemd/coredump/`.

29.15 Debugging TNSR

The following commands enable debugging information in various aspects of TNSR. These should only be used under direction of Netgate.

debug cli [level <n>] Enable debugging in `clixon` and `cligen` at the given level.

debug tnsr (clear|set|value) <flags> Enable debugging in TNSR. The `set` or `clear` command may be repeated multiple times to add or remove individual flag values. The `value` command may be used to directly set the value. The `<flags>` value is the logical `or` of all desired debugging flags.

The following flag values are available:

Flag	Value
TDBG_NONE	0x00000000
TDBG_FRR	0x00000001
TDBG_HOST	0x00000002
TDBG_KEA	0x00000004
TDBG_VPP	0x00000008
TDBG_NTP	0x00000010
TDBG_STRONGSWAN	0x00000020
TDBG_UNBOUND	0x00000040
TDBG_HTTP	0x00000080
TDBG_DELAYED_NODE	0x00001000
TDBG_DEP_GRAPH	0x00002000
TDBG_TRANSACTION	0x00004000
TDBG_ACL	0x00010000
TDBG_BGP	0x00020000
TDBG_BRIDGE	0x00040000
TDBG_INTF	0x00080000
TDBG_NEIGHBOR	0x00100000
TDBG_SUBIF	0x00200000
TDBG_SYSCTL	0x00400000
TDBG_GRE	0x00800000
TDBG_LOOPBACK	0x01000000
TDBG_ROUTE	0x02000000
TDBG_SPAN	0x04000000
TDBG_MAP	0x08000000

debug vmgmt (clear|set|value) <flags> Enable VPP Mgmt library debug. The `set` or `clear` command may be repeated multiple times to add or remove individual flag values. The `value` command may be used to directly set the value. The `<flags>` value is the logical `or` of all desired debugging flags.

The following flag values are available:

Flag	Value
VDBG_NONE	0x0000
VDBG_API_SETUP	0x0001
VDBG_API_MSG	0x0002
VDBG_ACL	0x0004
VDBG_BRIDGE	0x0008
VDBG_INTF	0x0010
VDBG_NAT	0x0020
VDBG_TAP	0x0040
VDBG_MEMIF	0x0080
VDBG_LLDP	0x0100
VDBG_GRE	0x0200
VDBG_MAP	0x0400
VDBG_ROUTE	0x0800

no debug (cli|tnsr|vmgmt) Removes all debugging.

29.16 Diagnostic Information for Support

When contacting Netgate for support at <https://go.netgate.com>, support representatives may require information about the configuration and state of a TNSR installation. TNSR includes a utility, `tnsr-diag`, which gathers diagnostic information automatically. The archive it produces can then be included in communications with Netgate support.

The command to collect this data may be run from a shell using `sudo tnsr-diag`, or from within the TNSR CLI using `shell sudo tnsr-diag`.

For example:

```
tnsr# shell sudo tnsr-diag
Collecting...
Diagnostic data stored at /tmp/tnsr-diag-2020-02-13-161340.zip
```

As shown in the example output above, the data is stored in a compressed archive in `/tmp` named `tnsr-diag-<timestamp>.zip`. This file may be copied from the TNSR device using `scp`.

Warning: The archive generated by this utility may contain sensitive information about the router configuration and environment. Review its contents before transmitting the information, and always use a secure communications method.

COMMANDS

30.1 Mode List

Internal Name	Prompt	Mode Description
access_list	config-access-list	Dynamic Routing Accesss List
acl	config-acl	Access Control List
acl_rule	config-acl-rule	ACL Rule
aspath	config-aspath	AS Path ordered rule
auth	config-user	User Authentication
bfd	config-bfd	Bidirectional Forwarding Detection
bfd_key	config-bfd-key	BFD key
bgp	config-bgp	BGP server
bgp_ip4multi	config-bgp-ip4multi	BGP IPv4 Multicast Address Family
bgp_ip4multi_nbr	config-bgp-ip4multi-nbr	BGP IPv4 Multicast Address Family Neighbor
bgp_ip4uni	config-bgp-ip4uni	BGP IPv4 Unicast Address Family
bgp_ip4uni_nbr	config-bgp-ip4uni-nbr	BGP IPv4 Unicast Address Family Neighbor
bgp_ip6multi	config-bgp-ip6multi	BGP IPv6 Multicast Address Family
bgp_ip6multi_nbr	config-bgp-ip6multi-nbr	BGP IPv6 Multicast Address Family Neighbor
bgp_ip6uni	config-bgp-ip6uni	BGP IPv6 Unicast Address Family
bgp_ip6uni_nbr	config-bgp-ip6uni-nbr	BGP IPv6 Unicast Address Family Neighbor
bgp_neighbor	config-bgp-neighbor	BGP Neighbor
bond	config-bond	Interface bonding
bridge	config-bridge	Bridge
community_list	config-community	BGP community list
config	config	Configuration
frr_bgp	config-frr-bgp	Dynamic Routing BGP
frr_ospf	config-frr-ospf	Dynamic Routing OSPF
frr_ospf6	config-frr-ospf6	Dynamic Routing OSPF6
gre	config-gre	Generic Route Encapsulation
host_acl	config-host-acl	Host Access List
host_acl_rule	config-host-acl-rule	Host Access List Rule
host_if	config-host-if	Host interface
http	config-http	HTTP server
ike_authentication	config-ike-auth	IKE peer authentication
ike_authentication_round	config-ike-auth-round	IKE peer authentication round
ike_child	config-ike-child	IKE child SA
ike_child_proposal	config-ike-child-proposal	IKE child SA proposal
ike_identity	config-ike-identity	IKE peer identity

Continued on next page

Table 1 – continued from previous page

Internal Name	Prompt	Mode Description
ike_proposal	config-ike-proposal	IKE proposal
interface	config-interface	Interface
ipsec_crypto_ike	config-ipsec-crypto-ike	IKE
ipsec_crypto_manual	config-crypto-manual	IPsec static keying
ipsec_tunnel	config-ipsec-tun	IPsec tunnel
kea_dhcp4	config-kea-dhcp4	DHCP4 Server
kea_dhcp4_log	config-kea-dhcp4-log	DHCP4 Log
kea_dhcp4_log_out	config-kea-dhcp4-log-out	DHCP4 Log output
kea_dhcp4_opt	config-kea-dhcp4-opt	DHCP4 option
kea_dhcp4_optdef	config-kea-dhcp4-optdef	DHCP4 option definition
kea_subnet4	config-kea-dhcp4-subnet4	DHCP4 subnet4
kea_subnet4_opt	config-kea-subnet4-opt	DHCP4 subnet4 option
kea_subnet4_pool	config-kea-subnet4-pool	DHCP4 subnet4 pool
kea_subnet4_pool_opt	config-kea-subnet4-pool-opt	DHCP4 subnet4 pool option
kea_subnet4_reservation	config-kea-subnet4-reservation	DHCP4 subnet4 host reservation
kea_subnet4_reservation_opt	config-kea-subnet4-reservation-opt	DHCP4 subnet4 host res option
loopback	config-loopback	Loopback interface
macip	config-macip	MAC/IP access control list
macip_rule	config-macip-rule	MACIP Rule
map	config-map	MAP-E/MAP-T
map_param	config-map-param	MAP-E/MAP-T global parameter
master		Initial, privileged
memif	config-memif	Memif interface
nacm_group	config-nacm-group	NACM group
nacm_rule	config-nacm-rule	NACM rule
nacm_rule_list	config-nacm-rule-list	NACM rule list
ntp	config-ntp	NTP
ntp_restrict	config-ntp-restrict	NTP restriction
ntp_server	config-ntp-server	NTP server
ospf	config-ospf	Dynamic routing OSPF Server
ospf_if	config-ospf-if	Dynamic routing OSPF Interface
ospf6	config-ospf6	Dynamic routing OSPF Server
ospf6_if	config-ospf6-if	Dynamic routing OSPF Interface
prefix_list	config-pref-list	Dynamic routing prefix list
route_dynamic_manager	config-route-dynamic-manager	Dynamic routing manager
route_map	config-rt-map	Route Map
route_table_v4	config-route-table-v4	IPv4 Static Route Table
route_table_v6	config-route-table-v6	IPv6 Static Route Table
rttbl4_next_hop	config-rttbl4-next-hop	IPv4 Next Hop
rttbl6_next_hop	config-rttbl6-next-hop	IPv6 Next Hop
span	config-span	SPAN
subif	config-subif	Sub-interface VLAN
tap	config-tap	Tap
unbound	config-unbound	Unbound DNS Server
unbound_fwd_zone	config-unbound-fwd-zone	Unbound forward-zone
unbound_local_host	config-unbound-local-host	Unbound local host override
unbound_local_zone	config-unbound-local-zone	Unbound local zone override
vrrp4	config-vrrp4	IPv4 VRRP
vrrp6	config-vrrp6	IPv6 VRRP

Continued on next page

Table 1 – continued from previous page

Internal Name	Prompt	Mode Description
vxlان	config-vxlان	VXLان

30.2 Master Mode Commands

```
tnsr# configure [terminal]
tnsr# debug cli [level <n>]
tnsr# debug tnsr (clear|set|value) <flags>
tnsr# debug vmgmt (clear|set|value) <flags>
tnsr# no debug (cli|tnsr|vmgmt)
tnsr# exit
tnsr# ls [-l]
tnsr# ping (<dest-host>|<dest-ip>) [ipv4|ipv6] [interface <if-name>]
        [source <src-addr>] [count <count>] [packet-size <bytes>]
        [ttl <ttl-hops>] [timeout <wait-sec>]
tnsr# pwd
tnsr# shell [<command>]
tnsr# traceroute (<dest-host>|<dest-ip>) [ipv4|ipv6] [interface <if-name>]
        [source <src-addr>] [packet-size <bytes>] [no-dns] [timeout <seconds>]
        [ttl <ttl-hops>] [waittime <wait-sec>]
tnsr# whoami
```

30.2.1 Package Management Commands

```
tnsr# package (info|list) [available|installed|updates] [<pkg-name>]
tnsr# package install <pkg-glob>
tnsr# package remove <pkg-glob>
tnsr# package search <term>
tnsr# package upgrade <pkg-glob>
```

30.2.2 Public Key Infrastructure Commands

```
tnsr# pki ca list
tnsr# pki ca <name> (append <source-name>|delete|enter|get|import <file>)
tnsr# pki certificate list
tnsr# pki certificate <name> (delete|enter|get|import <file>)
tnsr# pki private-key list
tnsr# pki private-key <name> (delete|enter|get|import <file>)
tnsr# pki private-key <name> generate [key-length (2048|3072|4096)]
tnsr# pki signing-request list
tnsr# pki signing-request <name> (delete|generate|get|sign (ca-name <ca>|self))
tnsr# pki signing-request set (city|common-name|country|org|org-unit|state) <text>
tnsr# pki signing-request set digest (md5|sha1|sha224|sha256|sha384|sha512)
tnsr# pki signing-request settings (clear|show)
```

30.3 Config Mode Commands

```

tnsr(config)# [no] acl <acl-name>
tnsr(config)# [no] auth user <user-name>
tnsr(config)# bfd conf-key-id <conf-key-id>
tnsr(config)# bfd session <bfd-session>
tnsr(config)# [no] cli history-config (enable|disable)
tnsr(config)# [no] cli history-config lines [<count>]
tnsr(config)# [no] cli option (auto-discard|check-delete-thresholds)
tnsr(config)# configuration candidate clear
tnsr(config)# configuration candidate commit
tnsr(config)# configuration candidate discard
tnsr(config)# configuration candidate load <filename> [(replace|merge)]
tnsr(config)# configuration candidate validate
tnsr(config)# configuration copy candidate startup
tnsr(config)# configuration copy running (candidate|startup)
tnsr(config)# configuration copy startup candidate
tnsr(config)# configuration save (candidate|running) <filename>
tnsr(config)# [no] dataplane buffers buffers-per-numa [<buffers-per-numa>]
tnsr(config)# [no] dataplane buffers default-data-size [<default-data-size>]
tnsr(config)# [no] dataplane cpu corelist-workers [<core-first> [- <core-last>]]
tnsr(config)# [no] dataplane cpu coremask-workers <coremask-workers>
tnsr(config)# [no] dataplane cpu main-core <main-core>
tnsr(config)# [no] dataplane cpu scheduler-policy (batch|fifo|idle|other|rr)
tnsr(config)# [no] dataplane cpu scheduler-priority <scheduler-priority>
tnsr(config)# [no] dataplane cpu skip-cores <skip-cores>
tnsr(config)# [no] dataplane cpu workers <workers>
tnsr(config)# dataplane dpdk dev <pci-id> (crypto|crypto-vf|network)
                        [num-rx-queues [<num-rx-queues>]] [num-tx-queues [<num-tx-queues>]]
                        [num-rx-desc [<num-rx-desc>]] [num-tx-desc [<num-tx-desc>]]
                        [vlan-strip-offload (off|on)]
tnsr(config)# dataplane dpdk dev <pci-id> network name <name>
tnsr(config)# no dataplane dpdk dev <pci-id> [name] [num-rx-queues] [num-tx-queues]
                        [num-rx-desc] [num-tx-desc] [vlan-strip-offload]
tnsr(config)# dataplane dpdk iova-mode (pa|va)
tnsr(config)# [no] dataplane dpdk iova-mode
tnsr(config)# [no] dataplane dpdk no-multi-seg
tnsr(config)# [no] dataplane dpdk no-tx-checksum-offload
tnsr(config)# [no] dataplane dpdk num-crypto-mbufs [<num>]
tnsr(config)# [no] dataplane dpdk uio-driver [<uio-driver>]
tnsr(config)# [no] dataplane dpdk vdev <sw-dev-type>
tnsr(config)# [no] dataplane ip heap-size [<size>]
tnsr(config)# [no] dataplane ip6 heap-size [<size>]
tnsr(config)# [no] dataplane ip6 hash-buckets [<size>]
tnsr(config)# [no] dataplane nat dslite-ce
tnsr(config)# [no] dataplane nat max-translations-per-user <n>
tnsr(config)# [no] dataplane nat mode (deterministic|endpoint-dependent|simple)
tnsr(config)# [no] dataplane nat mode-options simple (out2in-dpo|static-mapping-only)
tnsr(config)# [no] dataplane nat translation hash buckets <size>
tnsr(config)# [no] dataplane nat translation hash memory <size>
tnsr(config)# [no] dataplane nat user hash buckets <size>
tnsr(config)# [no] dataplane nat user hash memory <size>
tnsr(config)# [no] dataplane statseg heap-size <heap-size>[kKmMgG]
tnsr(config)# [no] dataplane statseg per-node-counters enable
tnsr(config)# [no] dataplane statseg socket-name <socket-name>
tnsr(config)# debug cli [level <n>]
tnsr(config)# debug tnsr (clear|set|value) <flags>
tnsr(config)# debug vmgmt (clear|set|value) <flags>
tnsr(config)# no debug (cli|tnsr|vmgmt)

```

(continues on next page)

(continued from previous page)

```

tnsr(config)# dhcp4 (enable|disable)
tnsr(config)# [no] dhcp4 server
tnsr(config)# dslite aftr endpoint <ip6-address>
tnsr(config)# dslite b4 endpoint <ip6-address>
tnsr(config)# dslite pool address <ipv4-addr-first> [- <ipv4-addr-last>]
tnsr(config)# no dslite [pool address]
tnsr(config)# exit
tnsr(config)# [no] gre <gre-name>
tnsr(config)# [no] host acl <acl-name>
tnsr(config)# [no] host interface <host-if-name>
tnsr(config)# http (enable|disable)
tnsr(config)# [no] http server
tnsr(config)# [no] interface <if-name>
tnsr(config)# interface clear counters [<interface>]
tnsr(config)# [no] interface bond <instance>
tnsr(config)# [no] interface bridge domain <domain-id>
tnsr(config)# [no] interface loopback <name>
tnsr(config)# [no] interface memif interface <id>
tnsr(config)# [no] interface memif socket id <id> filename <file>
tnsr(config)# [no] interface subif <interface> <subid>
tnsr(config)# [no] interface tap <host-name>
tnsr(config)# [no] ip reassembly (full|virtual) (ipv4|ipv6) expire-walk-interval [
-><interval-ms>]
tnsr(config)# [no] ip reassembly (full|virtual) (ipv4|ipv6) max-reassemblies [<max>]
tnsr(config)# [no] ip reassembly (full|virtual) (ipv4|ipv6) timeout [<timeout-ms>]
tnsr(config)# [no] ipsec tunnel <tunnel-num>
tnsr(config)# [no] lldp system-name <system-name>
tnsr(config)# [no] lldp tx-hold <transmit-hold>
tnsr(config)# [no] lldp tx-interval <transmit-interval>
tnsr(config)# [no] macip <macip-name>
tnsr(config)# nacm (enable|disable)
tnsr(config)# no nacm enable
tnsr(config)# [no] nacm exec-default (deny|permit)
tnsr(config)# [no] nacm group <group-name>
tnsr(config)# [no] nacm read-default (deny|permit)
tnsr(config)# [no] nacm rule-list <rule-list-name>
tnsr(config)# [no] nacm write-default (deny|permit)
tnsr(config)# [no] nat deterministic mapping inside <inside-prefix> outside <outside-
->prefix>
tnsr(config)# [no] nat global-options nat44 forwarding (true|false)
tnsr(config)# [no] nat global-options timeouts (icmp|tcp_established|tcp_
->transitory|udp) <seconds>
tnsr(config)# [no] nat ipfix logging [domain <domain-id>] [src-port <src-port>]
tnsr(config)# [no] nat nat64 map <domain-name>
tnsr(config)# [no] nat nat64 map parameters
tnsr(config)# [no] nat pool (addresses <ip-first> [- <ip-last>]|interface <if-name>)
[<twice-nat>] [route-table <rt-tbl-name>]
tnsr(config)# [no] nat static mapping (icmp|udp|tcp) local <ip-local> [<port-local>]
external (<ip-external>|<if-name>) [<port-external>]
[<twice-nat>] [out-to-in-only] [route-table <rt-tbl-name>]
tnsr(config)# [no] neighbor <interface> <ip-address> <mac-address> [no-adj-route-
->table-entry]
tnsr(config)# ntp (enable|disable)
tnsr(config)# no ntp enable
tnsr(config)# [no] ntp server
tnsr(config)# [no] route dynamic access-list <access-list-name>
tnsr(config)# route dynamic bgp

```

(continues on next page)

(continued from previous page)

```

tnsr(config)# route dynamic manager
tnsr(config)# route dynamic ospf
tnsr(config)# route dynamic ospf6
tnsr(config)# [no] route dynamic prefix-list <prefix-list-name>
tnsr(config)# [no] route dynamic route-map <route-map-name> (permit|deny) sequence
↳<sequence>
tnsr(config)# no route dynamic route-map [<route-map-name> [(permit|deny) sequence
↳<sequence>]]
tnsr(config)# [no] route (ipv4|ipv6) table <route-table-name>
tnsr(config)# server snmp (disable|enable)
tnsr(config)# service (backend|bgp|dataplane|dhcp|http|ike|ntp|restconf|unbound)
                coredump (enable|disable)
tnsr(config)# service bgp (start|stop|restart|status)
tnsr(config)# service dataplane (start|stop|restart|status)
tnsr(config)# service dhcp (start|stop|reload|status) [dhcp4|dhcp6|dhcp_ddns]
tnsr(config)# service http (start|stop|restart|status)
tnsr(config)# service ntp (start|stop|restart|status)
tnsr(config)# service snmp (start|stop|restart|status)
tnsr(config)# service unbound (start|stop|status|restart|reload)
tnsr(config)# [no] snmp access group-name <group-name>
                        prefix (exact|prefix)
                        model (any|v1|v2c)
                        level (noauth|auth|priv)
                        read <read-view>
                        write <write-view>
tnsr(config)# snmp community community-name <community-name>
                        source (<src-prefix>|default)
                        security-name <security-name>
tnsr(config)# snmp group group-name <group-name>
                        security-name <security-name>
                        security-model (any|v1|v2c)
tnsr(config)# snmp view view-name <view-name>
                        view-type (included|excluded)
                        oid <oid>
tnsr(config)# [no] span <if-name-src>
tnsr(config)# [no] sysctl vm nr_hugepages <u64>
tnsr(config)# [no] sysctl vm max_map_count <u64>
tnsr(config)# [no] sysctl kernel shmmem <u64>
tnsr(config)# [no] system contact <text>
tnsr(config)# [no] system description <text>
tnsr(config)# [no] system location <text>
tnsr(config)# [no] system name <text>
tnsr(config)# [no] unbound server
tnsr(config)# unbound (enable|disable)
tnsr(config)# no unbound enable
tnsr(config)# [no] vxlan <vxlan-name>

```

30.4 Show Commands in Both Master and Config Modes

```

tnsr# show acl [<acl-name>]
tnsr# show bfd
tnsr# show bfd keys [conf-key-id <conf-key-id>]
tnsr# show bfd sessions [conf-key-id <conf-key-id> | peer-ip-addr <peer-addr>]
tnsr# show cli

```

(continues on next page)

(continued from previous page)

```

tnsr# show clock
tnsr# show configuration (candidate|running|startup) [xml|json]
tnsr# show dataplane cpu threads
tnsr# show dslite
tnsr# show gre [<tunnel-name>]
tnsr# show host interface (<name>|ipv4|ipv6|link)
tnsr# show http [<config-file>]
tnsr# show interface [<if-name>]
      [(access-list|bonding|counters|ip [(nat|vrrp-virtual-router)]|
      ipv6 [vrrp-virtual-router]|link|mac-address|rx-queues|vlan tag-rewrite)]
tnsr# show interface bridge domain [<bdi>]
tnsr# show interface loopback [<loopback-name>]
tnsr# show interface memif [<id>]
tnsr# show interface bond [<id>]
tnsr# show interface lacp [<if-name>]
tnsr# show interface tap
tnsr# show ip reassembly [(full|virtual) [(ipv4|ipv6)]]
tnsr# show ipsec tunnel [<tunnel_number> [child|ike|verbose]]
tnsr# show kea [(keactrl|dhcp4) [config-file]]
tnsr# show kea dhcp4 leases
tnsr# show macip [<macip-name>]
tnsr# show map [<map-domain-name>]
tnsr# show nacm [group [<group-name>] | rule-list [<rule-list-name>]]
tnsr# show nat [config|deterministic-mappings|interface-sides|static-mappings]
tnsr# show nat dynamic (addresses|interfaces)
tnsr# show nat sessions [verbose]
tnsr# show neighbor [interface <if-name>]
tnsr# show ntp [(associations|peers) [associd <id>]]
tnsr# show ntp config-file
tnsr# show packet-counters
tnsr# show route dynamic access-list [<access-list-name>]
tnsr# show route dynamic bgp as-path [<as-path-name>]
tnsr# show route dynamic bgp community-list [<community-list-name>]
tnsr# show route dynamic bgp config [<as-number>]
tnsr# show route dynamic bgp (ipv4|ipv6) neighbors [[<peer>] [advertised-routes|
      dampened-routes|flap-statistics|prefix-counts|received|received-
      ↪routes|routes]]
tnsr# show route dynamic bgp (ipv4|ipv6) network <prefix>
tnsr# show route dynamic bgp (ipv4|ipv6) summary
tnsr# show route dynamic bgp nexthop [detail]
tnsr# show route dynamic bgp peer-group <peer-group-name>
tnsr# show route dynamic manager
tnsr# show route dynamic ospf [(border-routers|config|interface|neighbor [detail]|
      route|router-info)]
tnsr# show route dynamic ospf database [(asbr-summary|external|max-age|network|nssa-
      ↪external|
      opaque-area|opaque-as|opaque-link|router|self-originate|summary)]
tnsr# show route dynamic ospf6 area [<area-id>]
tnsr# show route dynamic ospf6 [(area|border-
      ↪routers|config|database|interface|linkstate|
      neighbor [detail]|route-table|spf)]
tnsr# show route dynamic prefix-list [<prefix-list-name>]
tnsr# show route dynamic route-map [<route-map-name>]
tnsr# show route [table <route-table-name>]
tnsr# show span
tnsr# show sysctl
tnsr# show system

```

(continues on next page)

(continued from previous page)

```
tnsr# show unbound [config-file]
tnsr# show version
tnsr# show vxlan [<vxlan-name>]
```

30.5 Access Control List Modes

30.5.1 Enter Access Control List Mode

```
tnsr(config)# acl <acl-name>
tnsr(config-acl)#
```

30.5.2 Access Control List Mode Commands

```
tnsr(config-acl)# rule <seq-number>
```

30.5.3 Remove Access Control List

```
tnsr(config)# no acl <acl-name>
```

30.5.4 Enter ACL Rule Mode

```
tnsr(config-acl)# rule <seq-number>
tnsr(config-acl-rule)#
```

30.5.5 ACL Rule Mode Commands

```
tnsr(config-acl-rule)# action (deny|permit|reflect)
tnsr(config-acl-rule)# ip-version (ipv4|ipv6)
tnsr(config-acl-rule)# no action [deny|permit|reflect]
tnsr(config-acl-rule)# destination address <ip-prefix>
tnsr(config-acl-rule)# no destination address [<ip-prefix>]
tnsr(config-acl-rule)# [no] destination port (any|<first> [- <last>])
tnsr(config-acl-rule)# [no] icmp type (any|<type-first> [- <type-last>])
tnsr(config-acl-rule)# [no] icmp code (any|<code-first> [- <code-last>])
tnsr(config-acl-rule)# protocol (any|icmp|icmpv6|tcp|udp|<proto-number>)
tnsr(config-acl-rule)# no protocol
tnsr(config-acl-rule)# source address <ip-prefix>
tnsr(config-acl-rule)# no source address [<ip-prefix>]
tnsr(config-acl-rule)# [no] source port (any|<first> [- <last>])
tnsr(config-acl-rule)# [no] tcp flags mask <mask> value <value>
tnsr(config-acl-rule)# [no] tcp flags value <value> mask <mask>
```

30.5.6 Remove ACL Rule

```
tnsr(config-acl)# no rule <seq>
```

30.6 MACIP ACL Mode

30.6.1 Enter MACIP ACL Mode

```
tnsr(config)# macip <macip-name>  
tnsr(config-macip)#
```

30.6.2 MACIP ACL Mode Commands

```
tnsr(config-macip)# rule <seq>
```

30.6.3 Remove MACIP ACL

```
tnsr(config-macip)# no macip <macip-name>
```

30.6.4 Enter MACIP ACL Rule Mode

```
tnsr(config-macip)# rule <seq-number>  
tnsr(config-macip-rule)#
```

30.6.5 MACIP Rule Mode Commands

```
tnsr(config-macip-rule)# action (deny|permit)  
tnsr(config-macip-rule)# no action [deny|permit]  
tnsr(config-macip-rule)# ip-version (ipv4|ipv6)  
tnsr(config-macip-rule)# address <ip-prefix>  
tnsr(config-macip-rule)# no address [<ip-prefix>]  
tnsr(config-macip-rule)# mac address <mac-address> [mask <mac-mask>]  
tnsr(config-macip-rule)# mac mask <mac-mask> [address <mac-address>]  
tnsr(config-macip-rule)# no mac  
tnsr(config-macip-rule)# no mac address [<mac-address>] [mask [<mac-mask>]]  
tnsr(config-macip-rule)# no mac mask [<mac-mask>] [address [<mac-address>]]
```

30.6.6 Remove MACIP ACL Rule

```
tnsr(config-macip)# no rule <seq-number>
```


30.7 GRE Mode

30.7.1 Enter GRE Mode

```
tnsr(config)# gre <gre-name>
tnsr(config-gre)#
```

30.7.2 GRE Mode Commands

```
tnsr(config-gre)# encapsulation route-table <rt-table-name>
tnsr(config-gre)# instance <id>
tnsr(config-gre)# destination <ip-address>
tnsr(config-gre)# source <ip-address>
tnsr(config-gre)# tunnel-type erspan session-id <session-id>
tnsr(config-gre)# tunnel-type (l3|teb)
```

30.7.3 Remove GRE Instance

```
tnsr(config)# no gre <gre-name>
```

30.8 HTTP mode

30.8.1 Enter HTTP mode

```
tnsr(config)# http server
tnsr(config-http)#
```

30.8.2 HTTP Mode Commands

```
tnsr(config-http)# authentication client-certificate-ca <cert-name>
tnsr(config-http)# authentication type (client-certificate|password|none)
tnsr(config-http)# enable restconf
tnsr(config-http)# disable restconf
tnsr(config-http)# server certificate <cert-name>
```

30.8.3 Remove http Configuration

```
tnsr(config)# no http server
```

30.9 Interface Mode

30.9.1 Enter Interface mode

```
tnsr(config)# interface <if-name>
tnsr(config-interface)#
```

30.9.2 Interface Mode Commands

```
tnsr(config-interface)# access-list (input|output) acl <acl-name> sequence <number>
tnsr(config-interface)# access-list macip <macip-name>
tnsr(config-interface)# no access-list
tnsr(config-interface)# no access-list acl <acl-name>
tnsr(config-interface)# no access-list macip [<macip-name>]
tnsr(config-interface)# no access-list [(input|output) [acl <acl-name> [sequence
↪<number>]]]
tnsr(config-interface)# bond <instance> [long-timeout] [passive]
tnsr(config-interface)# [no] bond <instance>
tnsr(config-interface)# bridge domain <bridge-domain-id> [bvi <bvi>] [shg <shg>]
tnsr(config-interface)# description <string-description>
tnsr(config-interface)# detailed-stats (enable|disable)
tnsr(config-interface)# [no] dhcp client ipv4 [hostname <host-name>]
tnsr(config-interface)# disable
tnsr(config-interface)# [no] enable
tnsr(config-interface)# [no] ip address <ip-prefix>
tnsr(config-interface)# [no] ip nat (inside|outside|none)
tnsr(config-interface)# [no] ip reassembly enable
tnsr(config-interface)# [no] ip reassembly type (full|virtual)
tnsr(config-interface)# [no] ip route-table <route-table-name-ipv4>
tnsr(config-interface)# [no] ip vrrp-virtual-router <vrid>
tnsr(config-interface)# [no] ipv6 address <ipv6-prefix>
tnsr(config-interface)# [no] ipv6 reassembly enable
tnsr(config-interface)# [no] ipv6 reassembly type (full|virtual)
tnsr(config-interface)# [no] ipv6 route-table <route-table-name-ipv6>
tnsr(config-interface)# [no] ipv6 vrrp-virtual-router <vrid>
tnsr(config-interface)# lldp port-name <port-name>
tnsr(config-interface)# lldp management ipv4 <ip-address>
tnsr(config-interface)# lldp management ipv6 <ipv6-address>
tnsr(config-interface)# lldp management oid <oid>
tnsr(config-interface)# map (disable|enable|translate)
tnsr(config-interface)# no map (enable|translate)
tnsr(config-interface)# mac-address <mac-address>
tnsr(config-interface)# mtu <mtu>
tnsr(config-interface)# no mtu [<mtu>]
tnsr(config-interface)# rx-queue <queue_num> cpu <core-id>
tnsr(config-interface)# no rx-queue [<queue_num> [cpu <core-id>]]
tnsr(config-interface)# vlan tag-rewrite (disable|pop-1|pop-2)
tnsr(config-interface)# vlan tag-rewrite push-1 (dot1ad|dot1q) <tag1>
tnsr(config-interface)# vlan tag-rewrite push-2 (dot1ad|dot1q) <tag1> <tag2>
tnsr(config-interface)# vlan tag-rewrite (translate-1-1|translate-2-1) (dot1ad|dot1q)
↪<tag1>
tnsr(config-interface)# vlan tag-rewrite (translate-1-2|translate-2-2) (dot1ad|dot1q)
↪<tag1> <tag2>
```

30.9.3 Remove Interface

```
tnsr(config)# no interface <if-name>
```

30.10 Loopback Mode

30.10.1 Enter Loopback Mode

```
tnsr(config)# interface loopback <loopback-name>  
tnsr(config-loopback)#
```

30.10.2 Loopback Mode Commands

```
tnsr(config-loopback)# instance <ul6>  
tnsr(config-loopback)# mac-address <mac-addr>  
tnsr(config-loopback)# description <rest>
```

30.10.3 Remove Loopback interface

```
tnsr(config)# no interface <loop<n>>  
tnsr(config)# no interface loopback <loopback-name>
```

30.11 Bridge Mode

30.11.1 Enter Bridge Mode

```
tnsr(config)# interface bridge <bdi>  
tnsr(config-bridge)#
```

30.11.2 Bridge Mode commands

```
tnsr(config-bridge)# [no] arp entry ip <ip-addr> mac <mac-addr>  
tnsr(config-bridge)# [no] arp term  
tnsr(config-bridge)# [no] flood  
tnsr(config-bridge)# [no] forward  
tnsr(config-bridge)# [no] learn  
tnsr(config-bridge)# [no] mac-age <mins>  
tnsr(config-bridge)# [no] rewrite  
tnsr(config-bridge)# [no] uu-flood
```

30.11.3 Remove Bridge

```
tnsr(config)# no interface bridge <bdi>
```

30.12 NAT Commands in Configure Mode

```
tnsr(config)# [no] nat static mapping (icmp|udp|tcp)
               local <ip> [<port>]
               external (<ip>|<if-name>) [<port>]
               [twice-nat] [out-to-in-only]
               [route-table <rt-tbl-name>]
tnsr(config)# [no] nat ipfix logging [domain <domain-id>] [src-port <port>]
tnsr(config)# [no] nat pool address <ip-first> [- <ip-last>] [twice-nat]
tnsr(config)# [no] nat pool interface <if-name> [twice-nat]
```

30.13 DS-Lite Commands in Configure Mode

```
tnsr(config)# dslite aftr endpoint <ip6-address>
tnsr(config)# dslite b4 endpoint <ip6-address>
tnsr(config)# dslite pool address <ipv4-addr-first> [- <ipv4-addr-last>]
```

30.14 Tap Mode

30.14.1 Enter Tap Mode

```
tnsr(config)# interface tap <tap-name>
tnsr(config-tap)#
```

30.14.2 Tap Mode commands

```
tnsr(config-tap)# [no] host bridge <bridge-name>
tnsr(config-tap)# [no] host ipv4 gateway <ipv4-addr>
tnsr(config-tap)# [no] host ipv4 prefix <ipv4-prefix>
tnsr(config-tap)# [no] host ipv6 gateway <ipv6-addr>
tnsr(config-tap)# [no] host ipv6 prefix <ipv6-prefix>
tnsr(config-tap)# [no] host mac-address <host-mac-address>
tnsr(config-tap)# [no] host name-space <netns>
tnsr(config-tap)# [no] instance <instance>
tnsr(config-tap)# [no] mac-address <mac-address>
tnsr(config-tap)# [no] rx-ring-size <size>
tnsr(config-tap)# [no] tx-ring-size <size>
```

30.14.3 Remove Tap

```
tnsr(config)# no interface tap <tap-name>
```

30.15 BFD Key Mode

30.15.1 Enter BFD Key Mode

```
tnsr(config)# bfd conf-key-id <conf-key-id>
tnsr(config-bfd-key)#
```

30.15.2 BFD Key Mode Commands

```
tnsr(config-bfd-key)# authentication type (keyed-sha1|meticulous-keyed-sha1)
tnsr(config-bfd-key)# secret < (<hex-pair>)[1-20] >
```

30.15.3 Remove BFD Key Configuration

```
tnsr(config)# no bfd conf-key-id <conf-key-id>
```

30.16 BFD Mode

30.16.1 Enter BFD Mode

```
tnsr(config)# bfd session <bfd-session>
tnsr(config-bfd)#
```

30.16.2 BFD Mode Commands

```
tnsr(config-bfd)# [no] bfd-key-id <bfd-key-id>
tnsr(config-bfd)# [no] conf-key-id <conf-key-id>
tnsr(config-bfd)# delayed (true|false)
tnsr(config-bfd)# desired-min-tx <microseconds>
tnsr(config-bfd)# detect-multiplier <n-packets>
tnsr(config-bfd)# disable
tnsr(config-bfd)# [no] enable (true|false)
tnsr(config-bfd)# interface <if-name>
tnsr(config-bfd)# local address <ip-address>
tnsr(config-bfd)# peer address <ip-address>
tnsr(config-bfd)# remote address <ip-address>
tnsr(config-bfd)# required-min-rx <microseconds>
```

30.16.3 Remove BFD Configuration

```
tnsr(config)# no bfd session <bfd-session>
```

30.16.4 Change BFD Admin State

```
tnsr# bfd session <bfd-session>  
tnsr(config-bfd)# enable false  
tnsr(config-bfd)# enable true  
tnsr(config-bfd)#
```

30.16.5 Change BFD Authentication

```
tnsr(config)# bfd session <bfd-session>  
tnsr(config-bfd)# bfd-key-id <bfd-key-id>  
tnsr(config-bfd)# conf-key-id <conf-key-id>  
tnsr(config-bfd)# delayed (true|false)
```

30.17 Host Interface Mode

30.17.1 Enter Host Interface Mode

```
tnsr(config)# host interface <if-name>  
tnsr(config-host-if)#
```

30.17.2 Host Interface Mode Commands

```
tnsr(config-host-if)# [no] description <text>  
tnsr(config-host-if)# disable  
tnsr(config-host-if)# [no] enable  
tnsr(config-host-if)# [no] ip address <ipv4-prefix>  
tnsr(config-host-if)# [no] ipv6 address <ipv6-prefix>  
tnsr(config-host-if)# mtu <mtu-value>
```

30.17.3 Remove Host Interface

```
tnsr(config)# no host interface <if-name>
```

30.18 IPsec Tunnel Mode

30.18.1 Enter IPsec Tunnel Mode

```
tnsr(config)# ipsec tunnel <tunnel-num>
tnsr(config-ipsec-tun)#
```

30.18.2 IPsec Tunnel Mode Commands

```
tnsr(config-ipsec-tun)# crypto config-type (ike|manual)
tnsr(config-ipsec-tun)# crypto (ike|manual)
tnsr(config-ipsec-tun)# [no] local-address <ip-address>
tnsr(config-ipsec-tun)# [no] remote-address (<ip-address>|<hostname>)
```

30.18.3 Remove IPsec Tunnel

```
tnsr(config)# no ipsec tunnel <tunnel-num>
```

30.19 IKE mode

30.19.1 Enter IKE mode

```
tnsr(config-ipsec-tun)# crypto ike
tnsr(config-ipsec-crypto-ike)#
```

30.19.2 IKE Mode Commands

```
tnsr(config-ipsec-crypto-ike)# [no] authentication (local|remote)
tnsr(config-ipsec-crypto-ike)# [no] child <name>
tnsr(config-ipsec-crypto-ike)# [no] identity (local|remote)
tnsr(config-ipsec-crypto-ike)# lifetime <seconds>
tnsr(config-ipsec-crypto-ike)# no lifetime
tnsr(config-ipsec-crypto-ike)# [no] proposal <number>
tnsr(config-ipsec-crypto-ike)# version (0|1|2)
tnsr(config-ipsec-crypto-ike)# no version
```

30.19.3 Remove IKE configuration

```
tnsr(config-ipsec-tun)# no crypto ike
```

30.20 IKE Peer Authentication Mode

30.20.1 Enter IKE Peer Authentication Mode

```
tnsr(config-ipsec-crypto-ike)# authentication (local|remote)
tnsr(config-ike-auth)#
```

30.20.2 IKE Peer Authentication Mode Commands

```
tnsr(config-ike-auth)# [no] round (1|2)
```

30.20.3 Remove IKE Peer Authentication Configuration

```
tnsr(config-ipsec-crypto-ike)# no authentication (local|remote)
```

30.21 IKE Peer Authentication Round Mode

30.21.1 Enter IKE Peer Authentication Round Mode

```
tnsr(config-ike-auth)# round (1|2)
tnsr(config-ike-auth-round)#
```

30.21.2 IKE Peer Authentication Round Mode Commands

```
tnsr(config-ike-auth-round)# type psk
tnsr(config-ike-auth-round)# no type
tnsr(config-ike-auth-round)# psk <pre-shared-key>
tnsr(config-ike-auth-round)# no psk
```

30.21.3 Remove IKE Peer Authentication Round Configuration

```
tnsr(config-ike-auth)# no round (1|2)
```

30.22 IKE Child SA Mode

30.22.1 Enter IKE Child SA Mode

```
tnsr(config-ipsec-crypto-ike)# child <name>
tnsr(config-ike-child)#
```


30.22.2 IKE Child SA Mode Commands

```
tnsr(config-ike-child)# lifetime <seconds>
tnsr(config-ike-child)# no lifetime
tnsr(config-ike-child)# [no] proposal <number>
```

30.22.3 Remove IKE Child SA

```
tnsr(config-ipsec-crypto-ike)# no child <name>
```

30.23 IKE Child SA Proposal Mode

30.23.1 Enter IKE Child SA Proposal Mode

```
tnsr(config-ike-child)# proposal <number>
tnsr(config-ike-child-proposal)#
```

30.23.2 IKE Child SA Proposal Mode Commands

```
tnsr(config-ike-child-proposal)# encryption <crypto-algorithm>
tnsr(config-ike-child-proposal)# no encryption
tnsr(config-ike-child-proposal)# integrity <integrity-algorithm>
tnsr(config-ike-child-proposal)# no integrity
tnsr(config-ike-child-proposal)# group <pfs-group>
tnsr(config-ike-child-proposal)# no group
tnsr(config-ike-child-proposal)# sequence-number (esn|noesn)
tnsr(config-ike-child-proposal)# no sequence-number
```

30.23.3 Remove IKE Child SA Proposal

```
tnsr(config-ike-child)# no proposal <number>
```

30.24 IKE Peer Identity Mode

30.24.1 Enter IKE Peer Identity Mode

```
tnsr(config-ipsec-crypto-ike)# identity (local|remote)
tnsr(config-ike-identity)#
```

30.24.2 IKE Peer Identity Mode Commands

```
tnsr(config-ike-identity)# type (none|address|email|fqdn|dn|key-id)
tnsr(config-ike-identity)# no type
tnsr(config-ike-identity)# value <identity>
tnsr(config-ike-identity)# no value
```

30.24.3 Remove IKE Peer Identity Configuration

```
tnsr(config-ipsec-crypto-ike)# no identity (local|remote)
```

30.25 IKE Proposal Mode

30.25.1 Enter IKE Proposal Mode

```
tnsr(config-ipsec-crypto-ike)# proposal <number>
tnsr(config-ike-proposal)#
```

30.25.2 IKE Proposal Mode Commands

```
tnsr(config-ike-proposal)# encryption <crypto-algorithm>
tnsr(config-ike-proposal)# no encryption
tnsr(config-ike-proposal)# integrity <integrity-algorithm>
tnsr(config-ike-proposal)# no integrity
tnsr(config-ike-proposal)# prf <prf-algorithm>
tnsr(config-ike-proposal)# no prf
tnsr(config-ike-proposal)# group <diffie-hellman-group>
tnsr(config-ike-proposal)# no group
```

30.25.3 Remove IKE Proposal Configuration

```
tnsr(config-ipsec-crypto-ike)# no proposal <number>
```

30.26 Map Mode

30.26.1 Enter Map Mode

```
tnsr(config)# nat nat64 map <domain-name>
```

30.26.2 Map Mode Commands

```
tnsr(config-map)# [no] description <desc>
tnsr(config-map)# [no] embedded-address bit-length <ea-width>
tnsr(config-map)# [no] ipv4 prefix <ip4-prefix>
tnsr(config-map)# [no] ipv6 prefix <ip6-prefix>
tnsr(config-map)# [no] ipv6 source <ip6-src>
tnsr(config-map)# [no] mtu <mtu-val>
tnsr(config-map)# [no] port-set length <psid-length>
tnsr(config-map)# [no] port-set offset <psid-offset>
tnsr(config-map)# [no] rule port-set <psid> ipv6-destination <ip6-address>
```

30.26.3 Remove Map Entry

```
tnsr(config)# [no] nat nat64 map <domain-name>
```

30.27 Map Parameters Mode

30.27.1 Enter Map Parameters Mode

```
tnsr(config)# nat nat64 map parameters
```

30.27.2 Map Parameters Mode Commands

```
tnsr(config-map-param)# [no] fragment (inner|outer)
tnsr(config-map-param)# [no] fragment ignore-df
tnsr(config-map-param)# [no] icmp source-address <ipv4-address>
tnsr(config-map-param)# [no] icmp6 unreachable-msgs (disable|enable)
tnsr(config-map-param)# [no] pre-resolve (ipv4|ipv6) next-hop <ip46-address>
tnsr(config-map-param)# [no] security-check (disable|enable)
tnsr(config-map-param)# [no] security-check fragments (disable|enable)
tnsr(config-map-param)# [no] traffic-class copy (disable|enable)
tnsr(config-map-param)# [no] traffic-class tc <tc-value>
```

30.28 memif Mode

30.28.1 Enter memif Mode

```
tnsr(config)# interface memif interface <id>
tnsr(config-memif)#
```

30.28.2 memif mode Commands

```
tnsr(config-memif)# buffer-size <ul6>
tnsr(config-memif)# mac-address <mac-addr>
tnsr(config-memif)# mode (ethernet|ip|punt|inject)
```

(continues on next page)

(continued from previous page)

```
tnsr(config-memif)# ring-size <power-of-2>
tnsr(config-memif)# role master
tnsr(config-memif)# role slave [rx-queues <u8>|tx-queues <u8>]
tnsr(config-memif)# secret <string-24>
tnsr(config-memif)# socket-id <socket-id>
```

30.28.3 Remove memif Interface

```
tnsr(config)# no interface memif interface <id>
```

30.29 Dynamic Routing Access List Mode

30.29.1 Enter Dynamic Routing Access List Mode

```
tnsr(config)# route dynamic access-list <access-list-name>
tnsr(config-access-list)#
```

30.29.2 Dynamic Routing Access List Mode Commands

```
tnsr(config-access-list)# [no] remark <text>
tnsr(config-access-list)# sequence <seq> (permit|deny) <ip-prefix>
tnsr(config-access-list)# no sequence <seq> [(permit|deny) [<ip-prefix>]]
```

30.29.3 Remove Dynamic Routing Access List

```
tnsr(config)# no route dynamic access-list <access-list-name>
```

30.30 Dynamic Routing Prefix List Mode

30.30.1 Enter Dynamic Routing Prefix List Mode

```
tnsr(config)# route dynamic prefix-list <pl-name>
tnsr(config-pref-list)#
```

30.30.2 Dynamic Routing Prefix List Mode Commands

```
tnsr(config-pref-list)# [no] descripton <text>
tnsr(config-pref-list)# sequence <seq> (permit|deny) <prefix> [ge <lower-bound>] [le
↪ <upper-bound>]
tnsr(config-pref-list)# no sequence <seq> [(permit|deny) <prefix> [ge <lower-bound>]
↪ [le <upper-bound>]]
```

30.30.3 Remove Dynamic Routing Prefix List

```
tnsr(config)# no route dynamic prefix-list <pl-name>
```

30.31 Dynamic Routing Route Map Rule Mode

30.31.1 Enter Dynamic Routing Route Map Rule Mode

```
tnsr(config)# route dynamic route-map <route-map-name> (permit|deny) sequence
↪<sequence>
tnsr(config-rt-map)#
```

30.31.2 Dynamic Routing Route Map Mode Commands

```
tnsr(config-rt-map)# [no] description <string>

tnsr(config-rt-map)# [no] match as-path <as-path-name>
tnsr(config-rt-map)# [no] match community <comm-list-name> [exact-match]
tnsr(config-rt-map)# [no] match extcommunity <extcomm-list-name>
tnsr(config-rt-map)# [no] match interface <if-name>
tnsr(config-rt-map)# [no] match ip address access-list <access-list-name>
tnsr(config-rt-map)# [no] match ip address prefix-list <prefix-list-name>
tnsr(config-rt-map)# [no] match ip next-hop access-list <access-list-name>
tnsr(config-rt-map)# [no] match ip next-hop <ipv4-address>
tnsr(config-rt-map)# [no] match ip next-hop prefix-list <prefix-list-name>
tnsr(config-rt-map)# [no] match ipv6 address access-list <access-list-name>
tnsr(config-rt-map)# [no] match ipv6 address prefix-list <prefix-list-name>
tnsr(config-rt-map)# [no] match large-community <large-comm-list-name>
tnsr(config-rt-map)# [no] match local-preference <preference-uint32>
tnsr(config-rt-map)# [no] match metric <metric-uint32>
tnsr(config-rt-map)# [no] match origin (egp|igp|incomplete)
tnsr(config-rt-map)# [no] match peer <peer-ip-address>
tnsr(config-rt-map)# [no] match probability <percent>
tnsr(config-rt-map)# [no] match source-protocol <src-protocol>
tnsr(config-rt-map)# [no] match tag <value-(1-4294967295)>

tnsr(config-rt-map)# [no] set aggregator as <asn> ip address <ipv4-address>
tnsr(config-rt-map)# [no] set as-path exclude <string-of-as-numbers>
tnsr(config-rt-map)# [no] set as-path prepend <string-of-as-numbers>
tnsr(config-rt-map)# [no] set as-path prepend last-as <asn>
tnsr(config-rt-map)# [no] set atomic-aggregate
tnsr(config-rt-map)# [no] set community none
tnsr(config-rt-map)# [no] set community <community-value> [additive]
tnsr(config-rt-map)# [no] set comm-list <community-list-name> delete
tnsr(config-rt-map)# [no] set extcommunity (rt|soo) <extcommunity-list-name>
tnsr(config-rt-map)# [no] set forwarding-address <ipv6-address>
tnsr(config-rt-map)# [no] set ip next-hop <ipv4-address>|peer-address|unchanged
tnsr(config-rt-map)# [no] set ipv4 vpn next-hop (<ipv4-address>|<ipv6-address>)
tnsr(config-rt-map)# [no] set ipv6 next-hop global <ipv6-address>
tnsr(config-rt-map)# [no] set ipv6 next-hop local <ipv6-address>
tnsr(config-rt-map)# [no] set ipv6 next-hop peer-address
tnsr(config-rt-map)# [no] set ipv6 next-hop prefer-global
```

(continues on next page)

(continued from previous page)

```

tnsr(config-rt-map)# [no] set ipv6 vpn next-hop (<ipv4-address>|<ipv6-address>)
tnsr(config-rt-map)# [no] set large-community none
tnsr(config-rt-map)# [no] set large-community <large-community-value> [additive]
tnsr(config-rt-map)# [no] set large-comm-list <large-comm-list-name> delete
tnsr(config-rt-map)# [no] set local-preference <preference>
tnsr(config-rt-map)# [no] set metric [+]<metric-uint32>
tnsr(config-rt-map)# [no] set origin (egp|igp|unknown)
tnsr(config-rt-map)# [no] set originator <ipv4-addr>
tnsr(config-rt-map)# [no] set src <ip-address>
tnsr(config-rt-map)# [no] set tag <tag-(1-4294967295)>
tnsr(config-rt-map)# [no] set weight <weight>

tnsr(config-rt-map)# [no] call <rt-map-name>

tnsr(config-rt-map)# [no] on-match next
tnsr(config-rt-map)# [no] on-match goto <sequence>

```

30.31.3 Remove Dynamic Routing Route Map

```
tnsr(config-rt-map)# no route dynamic route-map <route-map-name>
```

30.31.4 Remove Dynamic Routing Route Map Rule

```
tnsr(config-rt-map)# no route dynamic route-map <route-map-name> [permit|deny] ↪
↪sequence <sequence>
```

30.32 Dynamic Routing BGP Mode

30.32.1 Enter Dynamic Routing BGP Mode

```
tnsr(config)# route dynamic bgp
tnsr(config-frr-bgp)#
```

30.32.2 Dynamic Routing BGP Mode Commands

```

tnsr(config-frr-bgp)# [no] as-path <as-path-name>
tnsr(config-frr-bgp)# [no] community-list <comm-list-name> (standard|expanded)
                                [extended|large]

tnsr(config-frr-bgp)# disable
tnsr(config-frr-bgp)# [no] enable
tnsr(config-frr-bgp)# [no] option debug (allow-martians|nht|update-groups)
tnsr(config-frr-bgp)# [no] option debug as4 [segment]
tnsr(config-frr-bgp)# [no] option debug bestpath <ipv6-prefix>
tnsr(config-frr-bgp)# [no] option debug keepalive [<peer>]
tnsr(config-frr-bgp)# [no] option debug neighbor-events [<peer>]
tnsr(config-frr-bgp)# [no] option debug updates
                                [in <peer>|out <peer>|prefix (<ipv4-prefix>|<ipv6-
↪prefix>)]

```

(continues on next page)

(continued from previous page)

```

tnsr(config-frr-bgp)# [no] option debug zebra [prefix (<ipv4-prefix>|<ipv6-prefix>)]
tnsr(config-frr-bgp)# [no] server <asn>
tnsr(config-frr-bgp)# session clear (*|<peer>|<asn>) [soft]
tnsr(config-frr-bgp)# [no] route-map delay-timer <interval-sec>
tnsr(config-frr-bgp)# neighbor <if-name> <ip-address> <mac-address>
                                [no-adj-route-table-entry]
tnsr(config-frr-bgp)# no neighbor <if-name> [<ip-address>
                                [<mac-address> [no-adj-route-table-entry]]]

```

30.33 Dynamic Routing BGP Server Mode

30.33.1 Enter Dynamic Routing BGP Server Mode

```

tnsr(config-frr-bgp)# server <asn>
tnsr(config-bgp)#

```

30.33.2 Dynamic Routing BGP Server Mode Commands

```

tnsr(config-bgp)# [no] address-family (ipv4|ipv6) (unicast|multicast)
tnsr(config-bgp)# [no] always-compare-med
tnsr(config-bgp)# [no] bestpath as-path (confed|ignore|multipath-relax|as-set|no-as-
→set)
tnsr(config-bgp)# [no] bestpath compare-routerid
tnsr(config-bgp)# [no] bestpath med [confed|missing-as-worst]
tnsr(config-bgp)# [no] client-to-client reflection
tnsr(config-bgp)# [no] cluster-id (<ipv4>|<value>)
tnsr(config-bgp)# [no] coalesce-time <value>
tnsr(config-bgp)# [no] confederation identifier <ASN>
tnsr(config-bgp)# [no] confederation peer <ASN>
tnsr(config-bgp)# [no] dampening [penalty <value> [reuse <value> [suppress <value>
                                [maximum <value>]]]]
tnsr(config-bgp)# [no] deterministic-med
tnsr(config-bgp)# [no] disable-ebgp-connected-route-check
tnsr(config-bgp)# [no] enforce-first-as
tnsr(config-bgp)# [no] listen limit <value>
tnsr(config-bgp)# [no] listen range (<ip4-prefix>|<ip6-prefix>) peer-group <peer-group-
→name>
tnsr(config-bgp)# [no] max-med administrative [<med>]
tnsr(config-bgp)# [no] max-med on-startup period <seconds> [<med>]
tnsr(config-bgp)# [no] neighbor <peer>
tnsr(config-bgp)# [no] network import-check
tnsr(config-bgp)# [no] route-reflector allow-outbound-policy
tnsr(config-bgp)# [no] router-id <A.B.C.D>
tnsr(config-bgp)# [no] timers keep-alive <interval> hold-time <hold-time>
tnsr(config-bgp)# [no] update-delay <delay>
tnsr(config-bgp)# [no] write-quanta <packets>

```

30.33.3 Remove Dynamic Routing BGP Server

```
tnsr(config-frr-bgp)# no server <asn>
```

30.34 Dynamic Routing BGP Neighbor Mode

30.34.1 Enter Dynamic Routing BGP Neighbor Mode

```
tnsr(config-bgp)# neighbor <peer>  
tnsr(config-bgp-neighbor)#
```

30.34.2 Dynamic Routing BGP Neighbor Mode Commands

```
tnsr(config-bgp-neighbor)# [no] advertisement-interval <interval>  
tnsr(config-bgp-neighbor)# [no] bfd enabled (true|false)  
tnsr(config-bgp-neighbor)# [no] capability (dynamic|extended-nexthop)  
tnsr(config-bgp-neighbor)# [no] disable-connected-check  
tnsr(config-bgp-neighbor)# [no] description <string>  
tnsr(config-bgp-neighbor)# disable  
tnsr(config-bgp-neighbor)# [no] dont-capability-negotiate  
tnsr(config-bgp-neighbor)# [no] ebgp-multihop [hop-maximum <hops>]  
tnsr(config-bgp-neighbor)# [no] enable  
tnsr(config-bgp-neighbor)# [no] local-as <asn> [no-prepend [replace-as]]  
tnsr(config-bgp-neighbor)# [no] override-capability  
tnsr(config-bgp-neighbor)# [no] passive  
tnsr(config-bgp-neighbor)# [no] password <line>  
tnsr(config-bgp-neighbor)# [no] peer-group [<peer-group-name>]  
tnsr(config-bgp-neighbor)# [no] port <port>  
tnsr(config-bgp-neighbor)# [no] remote-as <asn>  
tnsr(config-bgp-neighbor)# [no] solo  
tnsr(config-bgp-neighbor)# [no] strict-capability-match  
tnsr(config-bgp-neighbor)# [no] timers keepalive <interval> holdtime <hold>  
tnsr(config-bgp-neighbor)# [no] timers connect <seconds>  
tnsr(config-bgp-neighbor)# [no] ttl-security hops <hops>  
tnsr(config-bgp-neighbor)# [no] update-source (<if-name>|<ip-address>)
```

30.34.3 Remove Dynamic Routing BGP Neighbor

```
tnsr(config-bgp)# no neighbor <peer>
```

30.35 Dynamic Routing BGP Address Family Mode

30.35.1 Enter Dynamic Routing BGP Address Family Mode

```
tnsr(config-bgp)# address-family ipv4 unicast  
tnsr(config-bgp-ip4uni)#
```



```
tnsr(config-bgp)# address-family ipv4 multicast
tnsr(config-bgp-ip4multi)#
```

```
tnsr(config-bgp)# address-family ipv6 unicast
tnsr(config-bgp-ip6uni)#
```

```
tnsr(config-bgp)# address-family ipv6 multicast
tnsr(config-bgp-ip6multi)#
```

30.35.2 Dynamic Routing BGP IPv4 Unicast Address Family Mode Commands

```
tnsr(config-bgp-ip4uni)# [no] aggregate-address <ipv4-prefix> [as-set] [summary-only]
tnsr(config-bgp-ip4uni)# [no] distance external <extern> internal <intern> local
↪<local>
tnsr(config-bgp-ip4uni)# [no] distance administrative <dist> prefix <ipv4-prefix>
[access-list <access-list-name>]
tnsr(config-bgp-ip4uni)# [no] maximum-paths <non-ibgp-paths> [igbp <ibgp-paths>
[equal-cluster-length]]
tnsr(config-bgp-ip4uni)# [no] neighbor <existing-neighbor>
tnsr(config-bgp-ip4uni)# [no] network <ipv4-prefix> [route-map <route-map>]
tnsr(config-bgp-ip4uni)# [no] redistribute <route-source> [metric <val>|route-map
↪<route-map-name>]
tnsr(config-bgp-ip4uni)# [no] redistribute ospf [metric <val>|route-map <route-map-
↪name>]
tnsr(config-bgp-ip4uni)# [no] redistribute table id <kernel-table-id> [metric <val>|
route-map <route-map-name>]
tnsr(config-bgp-ip4uni)# [no] table-map <route-map-name>
```

30.35.3 Dynamic Routing BGP IPv4 Multicast Address Family Mode Commands

```
tnsr(config-bgp-ip4multi)# [no] aggregate-address <ipv4-prefix> [as-set] [summary-
↪only]
tnsr(config-bgp-ip4multi)# [no] distance external <extern> internal <intern> local
↪<local>
tnsr(config-bgp-ip4multi)# [no] distance administrative <dist> prefix <ipv4-prefix>
[access-list <access-list-name>]
tnsr(config-bgp-ip4multi)# [no] neighbor <existing-neighbor>
tnsr(config-bgp-ip4multi)# [no] network <ipv4-prefix> [route-map <route-map>]
tnsr(config-bgp-ip4multi)# [no] table-map <route-map-name>
```

30.35.4 Dynamic Routing BGP IPv6 Unicast Address Family Mode Commands

```
tnsr(config-bgp-ip6uni)# [no] aggregate-address <ipv6-prefix> [as-set] [summary-only]
tnsr(config-bgp-ip6uni)# [no] distance external <extern> internal <intern> local
↪<local>
tnsr(config-bgp-ip6uni)# [no] distance administrative <dist> prefix <ipv6-prefix>
[access-list <access-list-name>]
tnsr(config-bgp-ip6uni)# [no] maximum-paths <non-ibgp-paths> [igbp <ibgp-paths>
[equal-cluster-length]]
tnsr(config-bgp-ip6uni)# [no] neighbor <existing-neighbor>
```

(continues on next page)

(continued from previous page)

```
tnsr(config-bgp-ip6uni)# [no] network <ipv6-prefix> [route-map <route-map>]
tnsr(config-bgp-ip6uni)# [no] redistribute <route-source> [metric <val>|route-map
↪<route-map-name>]
tnsr(config-bgp-ip6uni)# [no] redistribute ospf [metric <val>|route-map <route-map-
↪name>]
tnsr(config-bgp-ip6uni)# [no] redistribute table id <kernel-table-id> [metric <val>|
route-map <route-map-name>]
tnsr(config-bgp-ip6uni)# [no] table-map <route-map-name>
```

30.35.5 Dynamic Routing BGP IPv6 Multicast Address Family Mode Commands

```
tnsr(config-bgp-ip6multi)# [no] distance external <extern> internal <intern> local
↪<local>
tnsr(config-bgp-ip6multi)# [no] distance administrative <dist> prefix <ipv6-prefix>
access-list <access-list-name>
tnsr(config-bgp-ip6multi)# [no] neighbor <existing-neighbor>
tnsr(config-bgp-ip6multi)# [no] network <ipv6-prefix> [route-map <route-map>]
```

30.35.6 Remove Dynamic Routing BGP Address Family

```
tnsr(config-bgp)# no address-family (ipv4|ipv6) (unicast|multicast)
```

30.36 Dynamic Routing BGP Address Family Neighbor Mode

Note: Though the samples below indicate IPv4 unicast, the same syntax is used for all address families.

30.36.1 Enter Dynamic Routing BGP Address Family Neighbor Mode

```
tnsr(config-bgp-ip4uni)# neighbor <existing-neighbor>
tnsr(config-bgp-ip4uni-nbr)#
```

30.36.2 Dynamic Routing BGP Address Family Neighbor Mode Commands

```
tnsr(config-bgp-ip4uni-nbr)# [no] activate
tnsr(config-bgp-ip4uni-nbr)# [no] addpath-tx-all-paths
tnsr(config-bgp-ip4uni-nbr)# [no] addpath-tx-bestpath-per-as
tnsr(config-bgp-ip4uni-nbr)# [no] allowas-in [<occurrence>|origin]
tnsr(config-bgp-ip4uni-nbr)# [no] as-override
tnsr(config-bgp-ip4uni-nbr)# [no] attribute-unchanged [as-path|next-hop|med]
tnsr(config-bgp-ip4uni-nbr)# [no] capability orf prefix-list (send|receive|both)
tnsr(config-bgp-ip4uni-nbr)# [no] default-originate [route-map <route-map>]
tnsr(config-bgp-ip4uni-nbr)# [no] distribute-list <access-list-name> (in|out)
tnsr(config-bgp-ip4uni-nbr)# [no] filter-list <aspath-name> (in|out)
tnsr(config-bgp-ip4uni-nbr)# [no] maximum-prefix limit <val>
```

(continues on next page)

(continued from previous page)

```
tnsr(config-bgp-ip4uni-nbr)# [no] maximum-prefix restart <val>
tnsr(config-bgp-ip4uni-nbr)# [no] maximum-prefix threshold <val>
tnsr(config-bgp-ip4uni-nbr)# [no] maximum-prefix warning-only
tnsr(config-bgp-ip4uni-nbr)# [no] next-hop-self [force]
tnsr(config-bgp-ip4uni-nbr)# [no] prefix-list <prefix-list-name> (in|out)
tnsr(config-bgp-ip4uni-nbr)# [no] remove-private-AS [all] [replace-AS]
tnsr(config-bgp-ip4uni-nbr)# [no] route-map <name> (in|out)
tnsr(config-bgp-ip4uni-nbr)# [no] route-reflector-client
tnsr(config-bgp-ip4uni-nbr)# [no] route-server-client
tnsr(config-bgp-ip4uni-nbr)# [no] send-community (standard|large|extended)
tnsr(config-bgp-ip4uni-nbr)# [no] soft-reconfiguration inbound
tnsr(config-bgp-ip4uni-nbr)# [no] unsuppress-map <route-map>
tnsr(config-bgp-ip4uni-nbr)# [no] weight <weight>
```

30.36.3 Remove Dynamic Routing BGP Address Family Neighbor

```
tnsr(config-bgp-ip4uni)# no neighbor <existing-neighbor>
```

30.37 Dynamic Routing BGP Community List Mode

30.37.1 Enter Dynamic Routing BGP Community List Mode

```
tnsr(config-frr-bgp)# community-list <name> (standard|expanded) ↵
↵ [normal|extended|large]
tnsr(config-community)#
```

30.37.2 Dynamic Routing BGP Community List Mode Commands

```
tnsr(config-community)# description <text>
tnsr(config-community)# sequence <seq> (permit|deny) <community-value>
tnsr(config-community)# no description [<text>]
tnsr(config-community)# no sequence <seq> [(permit|deny) <community-value>]
```

30.37.3 Remove Dynamic Routing BGP Community List

```
tnsr(config-frr-bgp)# no community-list <name> (standard|expanded) [extended|large]
```

30.38 Dynamic Routing BGP AS Path Mode

30.38.1 Enter Dynamic Routing BGP AS Path Mode

```
tnsr(config-frr-bgp)# as-path <as-path-name>
tnsr(config-aspath)#
```

30.38.2 Dynamic Routing BGP AS Path Mode Commands

```
tnsr(config-aspath)# [no] rule <seq> (permit|deny) <pattern>
```

30.38.3 Remove Dynamic Routing BGP AS Path

```
tnsr(config-frr-bgp)# no as-path <as-path-name>
```

30.39 Dynamic Routing OSPF Mode

30.39.1 Enter Dynamic Routing OSPF Mode

```
tnsr(config)# route dynamic ospf
tnsr(config-frr-ospf)#
```

30.39.2 OSPF Mode Commands

```
tnsr(config-frr-ospf)# [no] enable
tnsr(config-frr-ospf)# disable
tnsr(config-frr-ospf)# [no] server
```

30.40 Dynamic Routing OSPF Server Mode

30.40.1 Entering OSPF Server Mode

```
tnsr(config-frr-ospf)# server
tnsr(config-ospf)#
```

30.40.2 OSPF Server Mode Commands

```
tnsr(config-ospf)# [no] area <area-id>
tnsr(config-ospf)# [no] auto-cost reference-bandwidth <bw>
tnsr(config-ospf)# [no] capability opaque-lsa
tnsr(config-ospf)# [no] compatible rfc-1583-compatibility
tnsr(config-ospf)# [no] debug (event|nssa|sr|te)
tnsr(config-ospf)# [no] debug (ism|nsm) (events|status|timers)
tnsr(config-ospf)# [no] debug lsa (flooding|generate|install|refresh)
tnsr(config-ospf)# [no] debug packet (dd|hello|ls-acknowledgment|ls-request|ls-update)
    (send|recv) [detail]
tnsr(config-ospf)# [no] debug zebra (interface|redistribute)
tnsr(config-ospf)# [no] default-information originate
    (always|metric <val>|route-map <map>|type <type>)
tnsr(config-ospf)# [no] default-metric <val>
tnsr(config-ospf)# [no] distance [(external|inter-area|intra-area)] <dist>
```

(continues on next page)

(continued from previous page)

```

tnsr(config-ospf)# [no] distribution-list out
                        (bgp|connected|kernel|static|table)
                        access-list <name>
tnsr(config-ospf)# [no] interface <if-name>
tnsr(config-ospf)# [no] log-adjacency-changes [detail]
tnsr(config-ospf)# [no] max-metric router-lsa administrative
tnsr(config-ospf)# [no] max-metric router-lsa (on-shutdown|on-startup) <seconds>
tnsr(config-ospf)# [no] neighbor <ip4-address> [(poll-interval <interval>|priority
↪<prio>)]
tnsr(config-ospf)# [no] ospf abr-type (cisco|ibm|shortcut|standard)
tnsr(config-ospf)# [no] ospf router-id <router-id>
tnsr(config-ospf)# [no] ospf write-multiplier <write>
tnsr(config-ospf)# [no] passive-interface <if-name> [<ip4-address>]
tnsr(config-ospf)# [no] pce address
                        (<ip4-address>|domain <asn>|flags <bits>|neighbor <asn>
↪|scope <bits>)
tnsr(config-ospf)# [no] redistribute
                        (bgp|connected|kernel|ospf|static|table)
                        [(metric <val>|route-map <map>|type <type>)]
tnsr(config-ospf)# [no] refresh timer <time>
tnsr(config-ospf)# [no] router-info as
tnsr(config-ospf)# [no] timers lsa min-arrival <min>
tnsr(config-ospf)# [no] timers throttle lsa all <delay>
tnsr(config-ospf)# [no] timers throttle spf (delay|initial-hold|maximum-hold) <val>

```

30.40.3 Remove OSPF Server Configuration

```
tnsr(config-frr-ospf)# no server
```

30.41 Dynamic Routing OSPF Interface Mode

30.41.1 Enter Dynamic Routing OSPF Interface Mode

```
tnsr(config-ospf)# interface <if-name>
tnsr(config-ospf-if)#
```

30.41.2 Dynamic Routing OSPF Interface Mode Commands

```

tnsr(config-ospf-if)# [no] bfd enabled (true|false)
tnsr(config-ospf-if)# [no] ip address (*|<ip4-address>) area <area-id>
tnsr(config-ospf-if)# [no] ip address (*|<ip4-address>) authentication [message-
↪digest|null]
tnsr(config-ospf-if)# [no] ip address (*|<ip4-address>) authentication-key <key>
tnsr(config-ospf-if)# [no] ip address (*|<ip4-address>) cost <link-cost>
tnsr(config-ospf-if)# [no] ip address (*|<ip4-address>) dead-interval minimal hello
↪<multiplier>
tnsr(config-ospf-if)# [no] ip address (*|<ip4-address>) dead-interval <time>
tnsr(config-ospf-if)# [no] ip address (*|<ip4-address>) hello-interval <interval>
tnsr(config-ospf-if)# [no] ip address (*|<ip4-address>) message-digest-key key-id <id>

```

(continues on next page)

(continued from previous page)

```

                                md5-key <key>
tnsr(config-ospf-if) # [no] ip address (*|<ip4-address>) mtu-ignore
tnsr(config-ospf-if) # [no] ip address (*|<ip4-address>) retransmit-interval <interval>
tnsr(config-ospf-if) # [no] ip address (*|<ip4-address>) priority <priority>
tnsr(config-ospf-if) # [no] ip address (*|<ip4-address>) transmit-delay <delay>
tnsr(config-ospf-if) # [no] ip network
                                (broadcast|non-broadcast|point-to-multipoint|point-to-
→point)

```

30.41.3 Remove Dynamic Routing OSPF Interface

```
tnsr(config-ospf) # no interface <if-name>
```

30.42 Dynamic Routing OSPF Area Mode

30.42.1 Enter Dynamic Routing OSPF Area Mode

```
tnsr(config-ospf) # area <area-id>
tnsr(config-ospf-area) #
```

30.42.2 Dynamic Routing OSPF Area Mode Commands

```

tnsr(config-ospf-area) # authentication [message-digest]
tnsr(config-ospf-area) # default-cost <cost>
tnsr(config-ospf-area) # export-list <acl-name>
tnsr(config-ospf-area) # filter-list (in|out) prefix-list <prefix-list-name>
tnsr(config-ospf-area) # import-list <acl-name>
tnsr(config-ospf-area) # nssa [(no-summary|translate (always|candidate|never))]
tnsr(config-ospf-area) # range <prefix> [cost <val>|not-advertise|substitute <sub-
→prefix>]
tnsr(config-ospf-area) # shortcut (default|disable|enable)
tnsr(config-ospf-area) # stub [no-summary]
tnsr(config-ospf-area) # virtual-link <router-id>

```

30.42.3 Remove Dynamic Routing OSPF Area

```
tnsr(config-ospf) # no area <area-id>
```

30.43 Dynamic Routing OSPF6 Mode

30.43.1 Enter Dynamic Routing OSPF6 Mode

```
tnsr(config) # route dynamic ospf6
tnsr(config-frr-ospf6) #
```

30.43.2 OSPF6 Mode Commands

```
tnsr(config-frr-ospf6)# [no] enable
tnsr(config-frr-ospf6)# disable
tnsr(config-frr-ospf6)# [no] server
```

30.44 Dynamic Routing OSPF6 Server Mode

30.44.1 Entering OSPF6 Server Mode

```
tnsr(config-frr-ospf6)# server
tnsr(config-ospf6)#
```

30.44.2 OSPF6 Server Mode Commands

```
tnsr(config-ospf6)# [no] area <area-id>
tnsr(config-ospf6)# [no] auto-cost reference-bandwidth <bw>
tnsr(config-ospf6)# [no] debug (abr|asbr|flooding|interface)
tnsr(config-ospf6)# [no] debug border-routers (area <area-id>|router <router-id>)
tnsr(config-ospf6)# [no] debug lsa (as-external|inter-prefix|inter-router|intra-
↪prefix|link|network|router|unknown) (examine|flooding|originate)
tnsr(config-ospf6)# [no] debug message (dd|hello|ls-acknowledgment|ls-request|ls-
↪update|unknown) (recv|send)
tnsr(config-ospf6)# [no] debug neighbor [(event|state)]
tnsr(config-ospf6)# [no] debug route [(inter-area|intra-area|memory|table)]
tnsr(config-ospf6)# [no] debug spf (database|process|time)
tnsr(config-ospf6)# [no] debug zebra [(recv|send)]
tnsr(config-ospf6)# [no] distance [(external|inter-area|intra-area)] <dist>
tnsr(config-ospf6)# [no] interface <if-name>
tnsr(config-ospf6)# [no] log-adjacency-changes [detail]
tnsr(config-ospf6)# [no] ospf6 router-id <router-id>
tnsr(config-ospf6)# [no] redistribute
      (bgp|connected|kernel|static|table) [route-map <map>]
tnsr(config-ospf6)# [no] stub-router administrative
tnsr(config-ospf6)# [no] timers lsa min-arrival <min>
tnsr(config-ospf6)# [no] timers throttle spf (delay|initial-hold|maximum-hold) <val>
```

30.44.3 Remove OSPF6 Server Configuration

```
tnsr(config-frr-ospf6)# no server
```

30.45 Dynamic Routing OSPF6 Interface Mode

30.45.1 Enter Dynamic Routing OSPF6 Interface Mode

```
tnsr(config-ospf6)# interface <if-name>
tnsr(config-ospf6-if)#
```

30.45.2 Dynamic Routing OSPF6 Interface Mode Commands

```
tnsr(config-ospf6-if)# [no] advertise prefix-list <name>
tnsr(config-ospf6-if)# [no] area <area-id>
tnsr(config-ospf6-if)# [no] bfd enabled (true|false)
tnsr(config-ospf6-if)# [no] cost outgoing <outgoing-cost>
tnsr(config-ospf6-if)# [no] dead-interval <time>
tnsr(config-ospf6-if)# [no] hello-interval <interval>
tnsr(config-ospf6-if)# [no] instance-id <value>
tnsr(config-ospf6-if)# [no] mtu <value>
tnsr(config-ospf6-if)# [no] mtu-ignore
tnsr(config-ospf6-if)# [no] network (broadcast|point-to-point)
tnsr(config-ospf6-if)# [no] passive
tnsr(config-ospf6-if)# [no] priority <priority>
tnsr(config-ospf6-if)# [no] retransmit-interval <interval>
tnsr(config-ospf6-if)# [no] transmit-delay <delay>
```

30.45.3 Remove Dynamic Routing OSPF6 Interface

```
tnsr(config-ospf6)# no interface <if-name>
```

30.46 Dynamic Routing OSPF6 Area Mode

30.46.1 Enter Dynamic Routing OSPF6 Area Mode

```
tnsr(config-ospf6)# area <area-id>
tnsr(config-ospf6-area)#
```

30.46.2 Dynamic Routing OSPF6 Area Mode Commands

```
tnsr(config-ospf6-area)# range <prefix> [cost <val>|not-advertise]
tnsr(config-ospf6-area)# stub [no-summary]
```

30.46.3 Remove Dynamic Routing OSPF6 Area

```
tnsr(config-ospf6)# no area <area-id>
```

30.47 Dynamic Routing RIP Mode

30.47.1 Enter Dynamic Routing RIP Mode

```
tnsr(config)# route dynamic rip
tnsr(config-frr-rip)#
```


30.47.2 RIP Mode Commands

```
tnsr(config-frr-rip)# [no] enable
tnsr(config-frr-rip)# disable
tnsr(config-frr-rip)# [no] server
```

30.48 Dynamic Routing RIP Server Mode

30.48.1 Entering RIP Server Mode

```
tnsr(config-frr-rip)# server
tnsr(config-rip)#
```

30.48.2 RIP Server Mode Commands

```
tnsr(config-rip)# [no] allow-ecmp
tnsr(config-rip)# [no] debug (events|zebra)
tnsr(config-rip)# [no] debug packet (send|recv)
tnsr(config-rip)# [no] default-information originate
tnsr(config-rip)# [no] distance default <value>
tnsr(config-rip)# [no] distance <prefix> distance <value> [access-list <acl-name>]
tnsr(config-rip)# [no] distribution-list interface (*|<if-name>)
    (access-list|prefix-list) (in|out) <name>
tnsr(config-rip)# [no] interface <if-name>
tnsr(config-rip)# [no] key-chain <name>
tnsr(config-rip)# [no] neighbor <ip4-address>
tnsr(config-rip)# [no] network (interface <if-name>|prefix <prefix>)
tnsr(config-rip)# [no] offset-list (*|<if-name>) (in|out) <acl-name>
tnsr(config-rip)# [no] passive-interface (default|<if-name>) [<ip4-address>]
tnsr(config-rip)# [no] redistribute (bgp|connected|kernel|ospf|static)
    [(metric <value>|route-map <name>)]
tnsr(config-rip)# [no] route prefix <ip4-prefix>
tnsr(config-rip)# [no] route-map-filter interface (default|<if-name>) (in|out) route-
    map <name>
tnsr(config-rip)# [no] timers (garbage-collection|table-update|timeout) <value>
tnsr(config-rip)# [no] version (1|2)
```

30.48.3 Remove RIP Server Configuration

```
tnsr(config-frr-rip)# no server
```

30.49 Dynamic Routing RIP Interface Mode

30.49.1 Enter Dynamic Routing RIP Interface Mode

```
tnsr(config-rip)# interface <if-name>
tnsr(config-rip-if)#
```

30.49.2 Dynamic Routing RIP Interface Mode Commands

```
tnsr(config-rip-if)# [no] authentication key-chain
tnsr(config-rip-if)# [no] authentication mode (md5|text) [auth-length (old-ripd|rfc)]
tnsr(config-rip-if)# [no] authentication string <auth-string>
tnsr(config-rip-if)# [no] receive version (1|2|both)
tnsr(config-rip-if)# [no] send version (1|2|both)
tnsr(config-rip-if)# [no] split-horizon [poisoned-reverse]
tnsr(config-rip-if)# [no] v2-broadcast
```

30.49.3 Remove Dynamic Routing RIP Interface

```
tnsr(config-rip)# no interface <if-name>
```

30.50 Dynamic Routing RIP Key Chain Mode

30.50.1 Enter Dynamic Routing RIP Key Chain Mode

```
tnsr(config-rip)# key-chain <name>
tnsr(config-rip-key-chain)#
```

30.50.2 Dynamic Routing RIP Key Chain Mode Commands

```
tnsr(config-rip-key-chain)# [no] key <key-id> string <key-string>
```

30.50.3 Remove Dynamic Routing RIP Key Chain

```
tnsr(config-rip)# no key-chain <name>
```

30.51 Dynamic Routing Manager Mode

30.51.1 Enter Dynamic Routing Manager Mode

```
tnsr(config)# route dynamic manager
tnsr(config-route-dynamic-manager)#
```

30.51.2 Dynamic Routing Manager Mode Commands

```
tnsr(config-route-dynamic-manager)# [no] debug (events|fpm|nht)
tnsr(config-route-dynamic-manager)# [no] debug kernel [msgdump [send|receive]]
tnsr(config-route-dynamic-manager)# [no] debug packet [send|receive] [detail]
tnsr(config-route-dynamic-manager)# [no] debug rib [detailed]
```

(continues on next page)

(continued from previous page)

```
tnsr(config-route-dynamic-manager)# [no] log file <filename> [<level>]  
tnsr(config-route-dynamic-manager)# [no] log syslog [<level>]
```

30.52 IPv4 Route Table Mode

30.52.1 Enter IPv4 Route Table Mode

```
tnsr(config)# route ipv4 table <route-table-name>  
tnsr(config-route-table-v4)#
```

30.52.2 IPv4 Route Table Mode Commands

```
tnsr(config-route-table-v4)# description <rest-of-line>  
tnsr(config-route-table-v4)# [no] route <destination-prefix>
```

30.52.3 Remove IPv4 Route Table

```
tnsr(config-route-table-v4)# no route ipv4 table <route-table-name>
```

30.53 IPv6 Route Table Mode

30.53.1 Enter IPv6 Route Table Mode

```
tnsr(config)# route ipv6 table <route-table-name>  
tnsr(config-route-table-v6)#
```

30.53.2 IPv6 Route Table Mode Commands

```
tnsr(config-route-table-v6)# description <rest-of-line>  
tnsr(config-route-table-v6)# [no] route <destination-prefix>
```

30.53.3 Remove IPv6 Route Table

```
tnsr(config-route-table-v6)# no route ipv6 table <route-table-name>
```

30.54 IPv4 or IPv6 Next Hop Mode

30.54.1 Enter IPv4 or IPv6 Next Hop Mode

```
tnsr(config-route-table-v46)# route <destination-prefix>
tnsr(config-rttbl46-next-hop)#
```

30.54.2 IPv4 or IPv6 Next Hop Mode Commands

```
tnsr(config-rttbl46-next-hop)# [no] description <rest-of-line>
tnsr(config-rttbl46-next-hop)# [no] next-hop <hop-id> via <ip46-addr>
                                     [<if-name>|<next-hop-table <route-table-name>>]
                                     [weight <multi-path-weight>]
                                     [preference <admin-preference>]
                                     [resolve-via-host] [resolve-via-attached]

tnsr(config-rttbl46-next-hop)# [no] next-hop <hop-id> via drop
tnsr(config-rttbl46-next-hop)# [no] next-hop <hop-id> via local
tnsr(config-rttbl46-next-hop)# [no] next-hop <hop-id> via null-send-unreach
tnsr(config-rttbl46-next-hop)# [no] next-hop <hop-id> via null-send-prohibit
tnsr(config-rttbl46-next-hop)# [no] next-hop <hop-id> classify <classify-table-name>
tnsr(config-rttbl46-next-hop)# [no] next-hop <hop-id> lookup [in] route-table <route-
↪table-name>
```

30.54.3 Remove IPv4 or IPv6 Next Hop

```
tnsr(config-rttbl46-next-hop)# no next-hop <hop-id>
```

30.55 SPAN Mode

30.55.1 Enter SPAN Mode

```
tnsr(config)# span <if-name-src>
tnsr(config-span)#
```

30.55.2 SPAN Mode Commands

```
tnsr(config-span)# onto <if-name-dst> (hw|l2) (rx|tx|both|disabled)
```

30.55.3 Remove Single SPAN Destination

```
tnsr(config-span)# no onto <if-name-dst> [(hw|l2)]
```

30.55.4 Remove SPAN

```
tnsr(config)# no span <if-name-src>
```

30.56 VXLAN Mode

30.56.1 Enter VXLAN Mode

```
tnsr(config)# vxlan <tunnel-name>  
tnsr(config-vxlan)#
```

30.56.2 VXLAN Mode Commands

```
tnsr(config-vxlan)# [no] destination <ip-addr>  
tnsr(config-vxlan)# [no] encapsulation (ipv4|ipv6) route-table <rt-table-name>  
tnsr(config-vxlan)# [no] instance <id>  
tnsr(config-vxlan)# [no] multicast interface <if-name>  
tnsr(config-vxlan)# [no] source <ip-addr>  
tnsr(config-vxlan)# [no] vni <u24>
```

30.56.3 Remove VXLAN Tunnel

```
tnsr(config)# no vxlan [<tunnel-name>]
```

30.57 User Authentication Configuration Mode

30.57.1 Enter User Authentication Configuration Mode

```
tnsr(config)# auth user <user-name>  
tnsr(config-user)#
```

30.57.2 User Authentication Mode Commands

```
tnsr(config-user)# [no] password <user-password>  
tnsr(config-user)# [no] user-keys <key-name>
```

30.57.3 Remove User

```
tnsr(config)# no auth user <user-name>
```

30.58 NTP Configuration Mode

30.58.1 Enter NTP Configuration Mode

```
tnsr(config)# ntp server  
tnsr(config-ntp)#
```

30.58.2 NTP Mode Commands

```
tnsr(config-ntp)# disable monitor  
tnsr(config-ntp)# enable monitor  
tnsr(config-ntp)# driftfile <file-path>  
tnsr(config-ntp)# interface sequence <seq> (drop|ignore|listen)  
                        (all|interface <if-name>|prefix <ip-prefix>)  
tnsr(config-ntp)# logconfig sequence <seq> (add|delete|set)  
                        (all|clock|peer|sync|sys) (all|events|info|statistics|status)  
tnsr(config-ntp)# restrict (default|host <fqdn>|prefix <ip-prefix>|source)  
tnsr(config-ntp)# server (address <ip-address>|host <fqdn>)  
tnsr(config-ntp)# statsdir <directory-path>  
tnsr(config-ntp)# tinker panic <n-secs>  
tnsr(config-ntp)# tos orphan <stratum>
```

30.58.3 Remove NTP Server

```
tnsr(config)# no ntp server
```

30.59 NTP Restrict Mode

30.59.1 Enter NTP Restrict Mode

```
tnsr(config-ntp)# restrict (default|host <fqdn>|prefix <ip-prefix>|source)
```

30.59.2 NTP Restrict Mode Commands

```
tnsr(config-ntp-restrict)# kod  
tnsr(config-ntp-restrict)# limited  
tnsr(config-ntp-restrict)# nomodify  
tnsr(config-ntp-restrict)# nopeer  
tnsr(config-ntp-restrict)# noquery  
tnsr(config-ntp-restrict)# noserve  
tnsr(config-ntp-restrict)# notrap
```

30.59.3 Remove NTP Restriction

```
tnsr(config-ntp)# no restrict (default|host <fqdn>|prefix <ip-prefix>|source)
```

30.60 NTP Upstream Server Mode

30.60.1 Enter NTP Upstream Server Mode

```
tnsr(config-ntp)# server (address <ip-address>|host <fqdn>)
```

30.60.2 NTP Upstream Server Mode Commands

```
tnsr(config-ntp-server)# iburst  
tnsr(config-ntp-server)# maxpoll <power-of-2-sec>  
tnsr(config-ntp-server)# nselect  
tnsr(config-ntp-server)# operational-mode (pool|server)  
tnsr(config-ntp-server)# prefer
```

30.60.3 Remove NTP Upstream Server

```
tnsr(config-ntp)# no server (address <ip-address>|host <fqdn>)
```

30.61 NACM Group Mode

30.61.1 Enter NACM Group Mode

```
tnsr(config)# nacm group <group-name>  
tnsr(config-nacm-group)#
```

30.61.2 NACM Group Mode Commands

```
tnsr(config-nacm-group)# [no] member <user-name>
```

30.61.3 Remove NACM Group

```
tnsr(config)# no nacm group <group-name>
```

30.62 NACM Rule-list Mode

30.62.1 Enter NACM Rule-list Mode

```
tnsr(config)# nacm rule-list <rule-list-name>
tnsr(config-nacm-rule-list)#
```

30.62.2 NACM Rule-list Mode Commands

```
tnsr(config-nacm-rule-list)# [no] group (*|<group-name>)
tnsr(config-nacm-rule-list)# [no] rule <rule-name>
```

30.62.3 Remove NACM Rule-list

```
tnsr(config)# no nacm rule-list <rule-list-name>
```

30.63 NACM Rule Mode

30.63.1 Enter NACM Rule Mode

```
tnsr(config-nacm-rule-list)# rule <rule-name>
tnsr(config-nacm-rule)#
```

30.63.2 NACM Rule Mode Commands

```
tnsr(config-nacm-rule)# [no] access-operations (*|create|read|update|delete|exec)
tnsr(config-nacm-rule)# [no] action (deny|permit)
tnsr(config-nacm-rule)# [no] module (*|<module-name>)
tnsr(config-nacm-rule)# [no] comment <rest>
tnsr(config-nacm-rule)# [no] rpc (*|<rpc-name>)
tnsr(config-nacm-rule)# [no] notification (*|<notification-name>)
tnsr(config-nacm-rule)# [no] path <node-id>
```

30.63.3 Remove NACM Rule

```
tnsr(config-nacm-rule-list)# no rule <rule-name>
```


30.64 DHCP IPv4 Server Config Mode

30.64.1 Enter DHCP IPv4 Server Mode

```
tnsr(config)# [no] dhcp4 server
tnsr(config)# dhcp4 {disable|enable}
tnsr(config)# no dhcp4 enable
tnsr(config-kea-dhcp4)#
```

30.64.2 DHCP IPv4 Server Mode Commands

```
tnsr(config-kea-dhcp4)# [no] decline-probation-period <seconds>
tnsr(config-kea-dhcp4)# [no] description <desc>
tnsr(config-kea-dhcp4)# [no] echo-client-id <boolean>
tnsr(config-kea-dhcp4)# [no] interface listen <if-name>
tnsr(config-kea-dhcp4)# [no] interface listen *
tnsr(config-kea-dhcp4)# [no] interface socket (raw|udp)
tnsr(config-kea-dhcp4)# [no] lease filename <filename>
tnsr(config-kea-dhcp4)# [no] lease lfc-interval <seconds>
tnsr(config-kea-dhcp4)# [no] lease persist <boolean>
tnsr(config-kea-dhcp4)# [no] logging <logger-name>
tnsr(config-kea-dhcp4)# [no] match-client-id <boolean>
tnsr(config-kea-dhcp4)# [no] next-server <ipv4-address>
tnsr(config-kea-dhcp4)# [no] option <dhcp4-option>
tnsr(config-kea-dhcp4)# [no] option-def <name>
tnsr(config-kea-dhcp4)# [no] rebind-timer <seconds>
tnsr(config-kea-dhcp4)# [no] renew-timer <seconds>
tnsr(config-kea-dhcp4)# [no] valid-lifetime <seconds>
```

30.64.3 Remove DHCP IPv4 Server Configuration

```
tnsr(config)# no dhcp4 server
```

30.65 DHCP4 Subnet4 Mode

30.65.1 Enter DHCP4 Subnet4 Mode

```
tnsr(config-kea-dhcp4)# subnet <ipv4-prefix>
tnsr(config-kea-subnet4)#
```

30.65.2 DHCP4 Subnet4 Mode Commands

```
tnsr(config-kea-subnet4)# [no] id <uint32>
tnsr(config-kea-subnet4)# [no] option <dhcp4-option>
tnsr(config-kea-subnet4)# [no] pool <ipv4-prefix>|<ipv4-range>
tnsr(config-kea-subnet4)# [no] interface <if-name>
```

30.65.3 Remove DHCP4 IPv4 Subnet4 Configuration

```
tnsr(config-kea-dhcp4)# no subnet <ipv4-prefix>|<ipv4-range>
```

30.66 DHCP4 Subnet4 Pool Mode

30.66.1 Enter DHCP4 Subnet4 Pool Mode

```
tnsr(config-kea-subnet4)# pool <ipv4-prefix>|<ipv4-range>
tnsr(config-kea-subnet4-pool)#
```

30.66.2 DHCP4 Subnet4 Pool Mode Commands

```
tnsr(config-kea-subnet4-pool)# [no] option <dhcp4-option>
```

30.66.3 Remove DHCP4 IPv4 Subnet4 Pool

```
tnsr(config-kea-subnet4)# no pool <ipv4-prefix>|<ipv4-range>
```

30.67 DHCP4 Subnet4 Reservation Mode

30.67.1 Enter DHCP4 Subnet4 Reservation Mode

```
tnsr(config-kea-subnet4)# reservation <ipv4-address>
tnsr(config-kea-subnet4-reservation)#
```

30.67.2 DHCP4 Subnet4 Reservation Mode Commands

```
tnsr(config-kea-subnet4-reservation)# [no] hostname <hostname>
tnsr(config-kea-subnet4-reservation)# [no] mac-address <mac-address>
tnsr(config-kea-subnet4-reservation)# [no] option <dhcp4-option>
```

30.67.3 Remove DHCP4 IPv4 Subnet4 Reservation

```
tnsr(config-kea-subnet4)# no reservation <ipv4-address>
```

30.68 Kea DHCP4, Subnet4, Pool, or Reservation Option Mode

30.68.1 Enter DHCP4 Option Mode

```
tnsr(config-kea-*)# option <dhcp4-option|option-def-nam>  
tnsr(config-kea-**-opt)#
```

30.68.2 DHCP4 Option Mode Commands

```
tnsr(config-kea-**-opt)# [no] always-send <boolean>  
tnsr(config-kea-**-opt)# [no] csv-format <boolean>  
tnsr(config-kea-**-opt)# [no] data <option-data>  
tnsr(config-kea-**-opt)# [no] space <space-name>
```

30.68.3 Remove DHCP4 Option Configuration

```
tnsr(config-kea-*)# no option <dhcp4-option>
```

30.69 Kea DHCP4 Option Definition Mode

30.69.1 Enter DHCP4 Option Definition Mode

```
tnsr(config-kea-dhcp4)# option-def <name>  
tnsr(config-kea-dhcp4-optdef)#
```

30.69.2 DHCP4 Option Definition Mode Commands

```
tnsr(config-kea-dhcp4-optdef)# array <array-val>  
tnsr(config-kea-dhcp4-optdef)# code <code-val>  
tnsr(config-kea-dhcp4-optdef)# encapsulate <encap>  
tnsr(config-kea-dhcp4-optdef)# record-types <types>  
tnsr(config-kea-dhcp4-optdef)# space <space-name>  
tnsr(config-kea-dhcp4-optdef)# type <type>
```

30.69.3 Remove DHCP4 Option Definition

```
tnsr(config-kea-dhcp4)# no option-def <name>
```

30.70 DHCP4 Log Mode

30.70.1 Enter DHCP4 Log Mode

```
tnsr(config-kea-dhcp4)# logging <logger-name>
tnsr(config-kea-dhcp4-log)#
```

30.70.2 DHCP4 Log Mode Commands

```
tnsr(config-kea-dhcp4-log)# [no] debug-level <level>
tnsr(config-kea-dhcp4-log)# [no] output <location>
tnsr(config-kea-dhcp4-log)# [no] severity (debug|error|fatal|info|warn)
```

30.70.3 Remove DHCP4 Log Configuration

```
tnsr(config-kea-dhcp4)# no logging <logger-name>
```

30.71 DHCP4 Log Output Mode

30.71.1 Enter DHCP4 Log Output Mode

```
tnsr(config-kea-dhcp4-log)# output <location>
tnsr(config-kea-dhcp4-log-out)#
```

30.71.2 DHCP4 Log Output Mode Commands

```
tnsr(config-kea-dhcp4-log-out)# [no] flush (false|true)
tnsr(config-kea-dhcp4-log-out)# [no] maxsize <size>
tnsr(config-kea-dhcp4-log-out)# [no] maxver <rotate>
```

30.71.3 Remove DHCP4 Log Output Configuration

```
tnsr(config-kea-dhcp4-log)# no output [<location>]
```

30.72 Unbound Server Mode

30.72.1 Enter Unbound Server Mode

```
tnsr(config)# unbound server
tnsr(config-unbound)#
```

30.72.2 Unbound Server Mode Commands

```
tnsr(config-unbound)# disable (caps-for-id | harden (dnssec-stripped|glue) |
                             hide (version|identity) | ip4 | ip6 | message prefetch |
                             serve-expired | tcp | udp)
tnsr(config-unbound)# edns reassembly size <s>
tnsr(config-unbound)# enable (caps-for-id | harden (dnssec-stripped|glue) |
                              hide (version|identity) | ip4 | ip6 | message prefetch |
                              serve-expired | tcp | udp)
tnsr(config-unbound)# forward-zone <zone-name>
tnsr(config-unbound)# host cache (num-hosts <num> | slabs <s> | ttl <t>)
tnsr(config-unbound)# interface <ip4-address>
tnsr(config-unbound)# jostle timeout <t>
tnsr(config-unbound)# key cache slabs <s>
tnsr(config-unbound)# message cache (size <s> | slabs <s>)
tnsr(config-unbound)# outgoing-interface <ip-address>
tnsr(config-unbound)# no outgoing-interface [<ip-address>]
tnsr(config-unbound)# port outgoing range <n>
tnsr(config-unbound)# rrset cache (size <s> | slabs <s>)
tnsr(config-unbound)# rrset-message cache ttl (minimum <min> | maximum <max>)
tnsr(config-unbound)# socket receive-buffer size <s>
tnsr(config-unbound)# tcp buffers (incoming <n> | outgoing <n>)
tnsr(config-unbound)# thread (num-queries <n> | num-threads <n> |
                              unwanted-reply-threshold <threshold>)
tnsr(config-unbound)# verbosity <level-0..5>
```

30.72.3 Remove Unbound Server

```
tnsr(config)# no unbound server
```

30.73 Unbound Forward-Zone Mode

30.73.1 Enter Unbound Forward-Zone Mode

```
tnsr(config-unbound)# forward-zone <zone-name>
tnsr(config-unbound-fwd-zone)#
```

30.73.2 Unbound Forward-Zone Mode Commands

```
tnsr(config-unbound-fwd-zone)# disable (forward-first | forward-tls-upstream)
tnsr(config-unbound-fwd-zone)# enable (forward-first | forward-tls-upstream)
tnsr(config-unbound-fwd-zone)# nameserver address <ip-address> [port <port>] [auth-
↪name <name>]
tnsr(config-unbound-fwd-zone)# nameserver host <host-name>
```

30.73.3 Remove Unbound Forward-Zone Zone

```
tnsr(config-unbound)# no forward-zone <zone-name>
```

30.74 Subif Mode

30.74.1 Enter Subif Mode

```
tnsr(config)# interface subif <if-name> <subid>
tnsr(config-subif)#
```

30.74.2 Subif Mode Commands

```
tnsr(config-subif)# default
tnsr(config-subif)# dot1q (<outer-vlan-id>|any)
tnsr(config-subif)# exact-match
tnsr(config-subif)# inner-dot1q (<inner-vlan-id>|any)
tnsr(config-subif)# outer-dot1ad (<outer-vlan-id>|any)
tnsr(config-subif)# outer-dot1q (<outer-vlan-id>|any)
```

30.74.3 Remove Subif

```
tnsr(config)# no interface subif <if-name> <subid>
```

30.75 Bond Mode

30.75.1 Enter Bond Mode

```
tnsr(config)# interface bond <instance>
tnsr(config-bond)#
```

30.75.2 Bond Mode Commands

```
tnsr(config-bond)# [no] load-balance (12|123|134)
tnsr(config-bond)# [no] mode (round-robin|active-backup|xor|broadcast|lacp)
tnsr(config-bond)# [no] mac-address <mac-address>
```

30.75.3 Remove Bond

```
tnsr(config)# no interface bond <instance>
```

30.76 Host ACL Mode

30.76.1 Enter Host ACL Mode

```
tnsr(config)# host acl <acl-name>
tnsr(config-host-acl)#
```

30.76.2 Host ACL Mode Commands

```
tnsr(config-host-acl)# [no] description <text>
tnsr(config-host-acl)# [no] rule <rule-seq>
tnsr(config-host-acl)# [no] sequence <acl-seq>
```

30.76.3 Remove Host ACL

```
tnsr(config)# no host acl <acl-name>
```

30.77 Host ACL Rule Mode

30.77.1 Enter Host ACL Rule Mode

```
tnsr(config-host-acl)# rule <rule-seq>
tnsr(config-host-acl-rule)#
```

30.77.2 Host ACL Rule Mode Commands

```
tnsr(config-host-acl-rule)# [no] action (deny|permit)
tnsr(config-host-acl-rule)# [no] description <text>
tnsr(config-host-acl-rule)# [no] match input-interface <host-interface>
tnsr(config-host-acl-rule)# [no] match ip address (source|destination) <prefix>
tnsr(config-host-acl-rule)# [no] match ip icmp type
                                (address-mask-reply|address-mask-request|destination-
↪ unreachable|
                                echo-reply|echo-request|info-reply|info-
↪ request|parameter-problem|
                                redirect|router-advertisement|router-
↪ solicitation|source-quench|
                                time-exceeded|timestamp-reply|timestamp-request) [code
↪ <code>]
tnsr(config-host-acl-rule)# [no] match ip icmpv6 type
                                (destination-unreachable|echo-reply|echo-request|
                                mld-listener-query|mld-listener-reduction|mld-listener-
↪ report|
                                nd-neighbor-advert|nd-neighbor-solicit|nd-redirect|
                                nd-router-advert|nd-router-solicit|packet-too-big|
                                parameter-problem|router-renumbering|time-exceeded) ↪
↪ [code <code>]
```

(continues on next page)

(continued from previous page)

```
tnsr(config-host-acl-rule)# [no] match ip port (source|destination) <port-num>
tnsr(config-host-acl-rule)# [no] match ip port (source|destination) range start <low-
↪port-num>
                                [end <high-port-num>]
tnsr(config-host-acl-rule)# [no] match ip protocol (icmp|tcp|udp|<proto-number>)
tnsr(config-host-acl-rule)# [no] match ip tcp flag (ack|cwr|ece|fin|psh|rst|syn|urg)
tnsr(config-host-acl-rule)# [no] match ip version (4|6)
tnsr(config-host-acl-rule)# [no] match mac address (source|destination) <mac>
```

30.77.3 Remove Host ACL Rule

```
tnsr(config-host-acl)# no rule <rule-seq>
```

30.78 VRRP Mode

30.78.1 Enter VRRP Mode

IPv4:

```
tnsr(config-interface)# ip vrrp-virtual-router <vrid>
tnsr(config-vrrp4)#
```

IPv6:

```
tnsr(config-interface)# ipv6 vrrp-virtual-router <vrid>
tnsr(config-vrrp6)#
```

30.78.2 VRRP Mode Commands

```
tnsr(config-vrrp46)# [no] accept-mode (false|true)
tnsr(config-vrrp46)# [no] preempt (false|true)
tnsr(config-vrrp46)# [no] priority <priority>
tnsr(config-vrrp46)# [no] track-interface <interface> priority-decrement <value>
tnsr(config-vrrp46)# [no] v3-advertisement-interval <advertise-interval-centi-sec>
tnsr(config-vrrp46)# [no] virtual-address <ipv4-address>
```

30.78.3 Remove VRRP

IPv4:

```
tnsr(config-interface)# no ip vrrp-virtual-router [<vrid>]
```

IPv6:

```
tnsr(config-interface)# no ipv6 vrrp-virtual-router [<vrid>]
```


API ENDPOINTS

In addition to the CLI, there are a variety of ways to configure TNSR, including a RESTful API.

31.1 YANG Data Models

The sets of functions and procedures used to manipulate the TNSR configuration are generated from the [RFC 7950](#) data models defined in the [TNSR YANG models](#).

31.2 RESTCONF API

TNSR can be controlled via a [RESTCONF API](#). Reference material, code examples, and more on the RESTCONF API may be found in the [TNSR API Documentation](#).

NETGATE TNSR RELEASES

- *TNSR 20.02 Release Notes*
 - *TNSR 20.02.1 Release Notes*
 - *TNSR 20.02.2 Release Notes*
- *TNSR 19.12 Release Notes*
- *TNSR 19.08 Release Notes*
- *TNSR 19.05 Release Notes*
- *TNSR 19.02 Release Notes*
 - *TNSR 19.02.1 Release Notes*
- *TNSR 18.11 Release Notes*
- *TNSR 18.08 Release Notes*
- *TNSR 18.05 Release Notes*
- *TNSR 0.1.0 Release Notes*

32.1 TNSR 20.02.2 Release Notes

32.1.1 About This Release

This is a maintenance release for TNSR software version 20.02 with bug fixes specific to the Azure platform. This version is only published on Azure.

Warning: For more information on changes in TNSR 20.02, including important information about upgrading from versions prior to TNSR 20.02, see *TNSR 20.02 Release Notes*.

General

- Fixed potential delays in initial connectivity for new Azure instances [3729]
- Fixed issues with the default startup configuration on new Azure instances [3699]

Interfaces

- Changed interface names on Azure to be prefixed with `VirtualFunctionEthernet` instead of using the speed and type, to improve consistency and reduce the likelihood of interface names changing unexpectedly [3728]
- Fixed issues with Azure instances attaching to the wrong devices [3660, 3727]
- Fixed issues with Azure instances failing to attach to any interfaces [3668]
- Fixed incorrect default MTU on Azure [3730]

32.2 TNSR 20.02.1 Release Notes

32.2.1 About This Release

This is a maintenance release for TNSR software version 20.02 with bug fixes.

Warning: For more information on changes in TNSR 20.02, including important information about upgrading from versions prior to TNSR 20.02, see [TNSR 20.02 Release Notes](#).

Dataplane

- Fixed backend crash when setting dataplane stat segment heap size [3598]
- Fixed validation during deletion/change of custom interface names [3461]

BGP

- Fixed static route restoration after failing over to a BGP route [3543]

IPsec

- Fixed interface validation during deletion of IPsec tunnel configuration [3456]

32.3 TNSR 20.02 Release Notes

32.3.1 About This Release

This is a regularly scheduled TNSR release including new features and bug fixes.

Warning: [TNSR 20.02.1](#) contains additional fixes for problems found in TNSR 20.02 and users should upgrade to that version instead.

The [TNSR 20.02.2](#) release corrects problems specific to Azure and is only available on that platform.

General

- Updated DPDK to 19.11 [2968]
- Updated VPP to 20.01 [2970]
- Updated strongswan to 5.8.2 [2964]
- Updated clixon to 4.3.2 [2570]
- Yang module version data is now stored in the configuration database [3022]
- Added support for *Shallow Virtual Reassembly* [2954]
 - This replaces manual reassembly configuration for NAT and MAP with *global reassembly configuration parameters*.
 - The old reassembly options under NAT and MAP must be removed from the configuration database. This change can be made automatically by *the configuration database update script* [3019, 3021].
- Added a *diagnostic information utility* for use when submitting support requests [2769]

Configuration Changes

Several areas of the configuration were changed. These changes must either be made manually or see *Updating the Configuration Database* for information on how to automatically update the configuration using a script included in this update.

- IPsec interfaces in the dataplane changed from `ipsec<N>` to `ipip<N>` and all references in the configuration must be updated to follow that change [2970]

This change can be made automatically by *the configuration database update script* [2972]

ACLs

- Fixed issues with accessing very large ACLs (100K rules) repeatedly [2558]

Azure

- Fixed network connectivity issues on Azure [2952]

Dataplane

- Fixed dataplane auto pinning of worker threads to cores not following expected conventions [2846]
- Fixed dataplane reporting incorrect physical core ID for main thread [2845]
- Added QAT crypto Virtual Functions (VF) to VPP `startup.conf` when `{corelist, coremask}-workers` is set and a crypto Physical Function (PF) is white listed [3248]
- Fixed potential situations where DPDK driver sections may not have been written to the dataplane startup configuration [3160]
- Added dataplane DPDK `iova-mode` configuration options [3416]
- The default dataplane UIO driver has been changed to `igb_uio` instead of using automatic driver selection [3414]
- Fixed issues with loading the `vfio-pci` driver at boot time [2686]

DHCP

- Added methods to view the current DHCP lease database via CLI and RESTCONF [2241]
- Added the ability for the DHCP server to use new custom option definitions rather than only redefining existing options with custom values [2934]

Interfaces

- Added options to assign per-interface RX queues to specific worker threads [2018]
- Fixed issues on XG-1537 and other systems with X552 NICs where if one of the **SFP+** (not copper) interfaces did not have an active link when the dataplane restarted, the interface would remain down when the link was reconnected. [2965]
- SPAN interfaces may now utilize *VXLAN interfaces* as destinations. [1027]

IPsec

- Fixed a dataplane and clixon crash due to large packets attempting to pass over IPsec. [2902]

Though the crash has been solved, packets larger than the `default-data-size` buffer value in the dataplane will fail to pass. To pass large IPsec packets, increase this buffer size. For example:

```
tnsr(config)# dataplane buffers default-data-size 16384
tnsr(config)# service dataplane restart
```

NAT

- Fixed incompatibility with NAT outside interfaces with output feature enabled being configured as a DHCP client [2914]
- Increased the default maximum NAT translations per user from 100 to 10240 [2752]

MAP

- Improved dataplane MAP-T RFC compliance [2977]
 - Fixed MAP-T IPv4 to IPv6 echo request not being translated correctly [2978]
 - Fixed MAP-T IPv4 to IPv6 echo reply not being translated correctly [2979]
 - Fixed MAP-T IPv6 to IPv4 echo request not being translated correctly [2980]
 - Fixed MAP-T IPv4 to IPv6 MTU Exceeded, DF flag set being handled incorrectly [2982]
 - Fixed MAP-T IPv4 to IPv6 TTL Expires at BR being handled incorrectly [2983]
 - Fixed MAP-T handling of spoofed IPv4 source prefix IPv6 to IPv4 [3053]
- Fixed an issue where MAP BR encapsulated/translated only the last fragment when it received fragmented packets from an IPv4 network [1887]
- Fixed fragmentation of IPv4 packets being performed regardless of configured MAP fragmentation behavior in MAP-T mode [1826]

Neighbors

- Fixed ARP responses for VPP outside interfaces responding incorrectly from the Host OS interface when both are connected to the same layer 2 [2266, 3314]
- Fixed issues with ARP table contents not being expired over time [3200]

QAT

- Added the capability to configure QAT VF entries passed to a virtual machine from the hypervisor [3250]

RESTCONF

- Added support for PATCH method in RESTCONF for API [1109]
- RESTCONF responses for leaf nodes with a value of an empty string ("") have changed, but still may not contain the expected encoded JSON output. [3450]

Previous versions of TNSR with clixon 4.0 or earlier returned the value as `null`, while clixon 4.3 now returns `[null]`. Per [RFC 7951](#), the previous behavior was incorrect. While the new behavior is closer to that mentioned in RFC 7951 section 6.9, the behavior described there is for `empty` type nodes, not `string` type. The intended behavior for empty strings is not yet clearly defined in RFC 7951.

This behavior is likely to change in future releases as the specification is refined.

Dynamic Routing

- Removed a redundant BGP command `enforce-multihop` which is identical to `disable-connected-check`.
 - Configuration database entries for `enforce-multihop` must be removed or changed to `disable-connected-check`. This change can be made automatically by *the configuration database update script* [3004]
- Fixed configuration of distance values for BGP address families via CLI [2869]
- Added validation to prevent configuring a `route-map` with a sequence number of 0 [2876]
- Removed incorrect `route-reflector-client` BGP option for eBGP peer from CLI [2936]
- Fixed setting multiple `attribute-unchanged` values via CLI [2941]
- Fixed setting `attribute-unchanged` BGP option without specifying a value [2942]
- Fixed setting `route-map` as a value for `unsuppress-map` via CLI [2944]
- Fixed disabling `send-community` BGP option in the CLI [2945]
- Fixed disabling `client-to-client reflection` BGP option in the CLI [2946]
- Fixed issue with displaying a large amount of received or advertised BGP prefixes taking a long time [2778]

SNMP

- Fixed SNMP configuration changes requiring a service restart [2568]

32.3.2 Known Limitations

General

- TNSR instances on VMWare configured for VM Hardware Compatibility with ESX 6.7 (VM Version 14 or later) cannot initialize their VMXNET3 interfaces unless there are 2 or more RX queues due to an upstream DPDK issue [2576]
 - Workaround 1: Create the VM with VM version 13 (ESX 6.5) and do not upgrade its compatibility level until this issue is resolved.
 - Workaround 2: Configure a `num-rx-queues` value of at least 2 for each VMXNET3 interface in the DPDK settings for the device(s) (*DPDK Configuration*) and restart the dataplane.

ACLs

- ACLs used with `access-list output` do not work on traffic sent to directly connected hosts [2057]

BFD

- Unable to set `delayed` option on an existing BFD session [2709]

CLI

- CLI does not return from shell in certain situations [2651]

Dataplane

- Systems with multiple CPU sockets using NUMA may experience dataplane issues at startup or when the dataplane is restarted manually [2383]
- CLI does not prevent the user from configuring a custom interface name which uses reserved keywords which may cause the dataplane to fail (e.g. `span`) [3234]
- UIO driver changes are not reflected on interfaces which are already in use [3209]
Workaround: Reboot the TNSR device.
- Setting dataplane stat segment heap size causes backend to crash [3598]
- Deletion/change of custom interface names is not validated properly [3461]

DHCP

- Unable to delete all DHCP server options at once from CLI [2667]

GRE

- Unable to modify GRE tunnel settings [2698]

Host Interfaces

- Configuration of host OS interface clears TNSR TAP interface configuration [2640]

Workaround: Remove and reconfigure the TAP interface.

- DHCP on Host Interface stops trying DHCP if a response is not received in a timely manner (Service = Failed) [3015]

Workaround: Set `PERSISTENT_DHCLIENT=1` in `/etc/sysconfig/network-scripts/ifcfg-<name>` for the affected host interface.

- Cannot remove an IP address assigned to a host interface during the installation process from within the TNSR CLI [3013]

HTTP Server

- HTTP server retains old configuration after TNSR services restart [2453]
- SSL certificate error when the HTTP server is configured with a certificate that uses md5 digest [2403]

Installer

- TNSR Install over OOB Management GUI may appear to fail due to the screen saver activating before installation is completed.

This affects installation using a console such as iDRAC Virtual Media redirector.

Workarounds: Press `tab` when the screensaver activates. Alternately, use vFlash instead of iDRAC for better performance.

Interfaces

- Packets do not pass through a subinterface after the subinterface configuration has been modified [1612]
- Chelsio interfaces crash the dataplane [1896]
- VLAN subinterfaces may not work under KVM using virtio drivers [2189]
- An IPv6 link-local address cannot manually be configured on an interface [2394]
- IPv6 addresses on IPsec or GRE interfaces may not be displayed in `show` command output [2425]
- Bridge domain ARP entries are not displayed in the CLI [2378]
- Bridge domain ARP entries cannot be removed from the CLI [2380]
- Bridge domain MAC age cannot be removed from the CLI [2381]
- Link state always reported as “up” when using `e1000` network drivers [2831]
- `vmxnet3` RSS fails to initialize, cannot pass packets [2576]

Workaround: Set `dataplane dpdk dev <device id> network num-rx-queues 2` in the TNSR CLI and restart the dataplane.

- Cannot add a DHCP client hostname to an existing DHCP client [2557]

Workaround: Remove the dhcp client from the interface and then re-add it with the hostname.

- Re-enabling loopback interface breaks packet forwarding until the dataplane is restarted [2828]

- Subinterface settings are not applied on change without restarting dataplane [2696]
- Unable to create multiple IP QinQ subinterfaces with the same outer vlan tag [2659]
- Unable to create a subinterface with dot1q any [2652]
- Full reassembly may not disable on an interface once enabled when using `no ip reassembly enable` [3360]

Workaround: Remove both the reassembly enable and type configuration on the interface:

```
tnsr(config-interface)# no ip reassembly enable
tnsr(config-interface)# no ip reassembly type
```

IPsec

- An IPsec tunnel which was removed and then added back in may take longer than expected to establish [1313]
- An SA ordering issue may prevent IPsec traffic from passing if both endpoints attempt to establish a tunnel at the same time [2391]
- Attempting to change `IKE lifetime` for an existing tunnel to a value lower than the lifetime of a child entry results in an unintuitive error message [3243]
- Deletion of IPsec tunnel configuration is not validated properly [3456]

LACP

- If a bond interface does not have a MAC address explicitly configured, the MAC address may become out of sync between the dataplane and host tap interfaces [2126]

Workaround: The MAC address will be synchronized when the interface status changes (up or down), so disable and enable the interface or restart the dataplane.

- There may be a 10-15 second delay with ARP resolution after configuring an LACP bond [2867]

LLDP

- All LLDP interface parameters must be configured at the same time. [3462]
- When LLDP parameters change, TNSR requires a dataplane restart for the new settings to take effect. [3486]
- LLDP parameter values are not validated by the CLI or RESTCONF and invalid values are rejected by the dataplane directly [3459]

MAP

- MAP-T BR cannot translate IPv4 ICMP echo reply to IPv6 [1749]
- MAP BR does not send ICMPv6 unreachable messages when a packet fails to match a MAP domain [1869]
- Pre-resolve does not work when MAP-T mode is used [1871]
- Full ip reassembly does not work with MAP [3386]
- ICMP6 echo request packets are being dropped on MAP-T BR when MAP domain with non-zero PSID offset is used [3401]

- Initial fragment of UDP and ICMP6 packets is dropped on MAP-T border router when it receives fragments from an IPv6 network [3412]
- Ethernet padding is incorrectly copied from IPv4 to IPv6 frames when translated by MAP [3460]

NACM

- Default parameters rule for NACM node `access-operation` and `module` does not work without explicit settings [2514]

NAT

- `twice-nat` does not work [1023]
- NAT forwarding is not working for `in2out` direction [1039]
- NAT forwarding fails with more than one worker thread [2031]
Note: This also affects connectivity to services on TNSR, such as RESTCONF, when the client is not on a directly connected network.
- Router with 1:1 NAT will drop packets with `ttl=2` from input interface [2849]
- VPP service fails if `NAT concurrent-reassemblies` is set to 1 and several fragments arriving to the NAT outside interface [2739]
- ICMP fragments arriving to NAT Inside interface aren't being reassembled by NAT reassembly function [2733]
- Dataplane fails on DS-Lite AFTR router when packets from B4 are received before pool is configured [3024]
Workaround: Configure the DS-Lite pool `**before**` the ``after` endpoint.
- DS-Lite CE configuration is not fully removed when deleted via CLI, which may leave TNSR with an invalid configuration database which cannot start [3030]
- Deterministic nat option is not compatible with a pool of IP addresses [3257]
- Reassembly timeout does not work when full IP reassembly is configured with NAT [3269]
- Shallow Virtual Reassembly cannot be disabled when it is enabled implicitly by other features such as NAT and MAP [3361]
- Shallow Virtual Reassembly may fail when configured explicitly after it is implicitly enabled by other features such as NAT and MAP [3362]
- Re-enabling full IP reassembly on an interface which has implicit shallow virtual reassembly enabled breaks the packet flow [3379]
- Setting reassembly type `full` and then enabling ip reassembly on an interface which has implicit shallow virtual reassembly enabled breaks packet flow [3380]
- Second fragment of a packet is not being virtually reassembled when `max-reassemblies` counter for shallow virtual reassembly is set to 1 [3384]

Neighbor / ARP / NDP

- Packet loss during ARP transaction immediately after Dataplane restart or interface disable/enable [2868]

NTP

- NTP server default restriction list cannot be deleted in CLI [3413]

RESTCONF

- RESTCONF responses for leaf nodes with a value of an empty string (" ") may not contain the expected encoded JSON output. [3450]

See [RESTCONF](#) earlier in this document for more details.

- RESTCONF responses containing certain IETF error types such as `application` errors may contain an extra JSON key, `rpc-error`, in the `error` list. RESTCONF users should accommodate this extra key, if present, when parsing IETF error messages. [3455]
- Incorrect BGP configuration is generated when IPv6 address family is configured via REST [2915]
- Adding a user via RESTCONF requires a password even when key is provided [2875]
- Adding MACIP rule via RESTCONF fails [2844]
- Cannot rename an ACL via RESTCONF [2843]
- Deleting ACL rule via RESTCONF crashes Clixon [2841]

Static Routing

- IPv6 packet loss may be observed between TNSR instances [2382]
- TNSR drops packets when an output interface configured in the routing table is disabled, even when other usable paths are present to the same destination [3359]

Dynamic Routing

- CLI shows that only IPv4 prefix is available within `prefix-list` sequence configuration [2689]

BGP

- An IPv6 BGP session cannot be established over IPsec or GRE [2429]
- BGP `maximum-path` option for eBGP and iBGP can not be configured simultaneously [2879]
- BGP `network backdoor` feature does not work without service restart [2873]
- Unable to verify received prefix-list entries via CLI when ORF capability is used [2864]
- `extended-nexthop` capability is not being negotiated between IPv6 BGP peers [2850]
- BGP session soft reset option does not work for IPv6 peers [2833]
Workaround: Reset the connection without soft option.
- `ttl-security hops` value can be set when `ebgp-multihop` is already configured (the options are mutually exclusive) [2832]
- `clixon-backend` fails when loading BGP config with 150k advertised prefixes [2784]
- BGP updates for new prefixes are sent every 60 seconds despite configured `advertisement-interval` value [2757]

- TNSR installs additional duplicated `next-hop` entries for multipath routes received via BGP [2935]
- IPv4 BGP summary command returns results for both IPv4 and IPv6 [3270]
- BGP `next-hop` attributes are not sent unmodified to an eBGP peer when `route-server-client` option is configured [2940]
- `show route dynamic bgp ipv6 summary` command will not show any information if address family is not specified when configuring BGP for IPv6 [2967]

Workaround: Set the address family when configuring BGP. Alternately, due to [3270], IPv6 information is current visible in `show route dynamic bgp ipv4 summary`, so use that command instead.

- Unable to configure BGP IPv4/IPv6 multicast address family using CLI [3038]
Workaround: Configure this feature via RESTCONF
- BGP listen range option disappears from the active dynamic routing daemon configuration after restarting BGP service [3043]
- Unable to verify dynamic BGP peer information from TNSR CLI [3044]
- Unable to configure BGP dampening values via TNSR CLI [3057]
- Unable to configure BGP `write-quanta` value via TNSR CLI [3087]
- Unable to configure BGP debug logging via TNSR CLI [3199]
- Unable to configure BGP confederation identifier via TNSR CLI [3210]
- Static routes may not be restored correctly after failing over to a BGP route [3543]

OSPF

- OSPF `default-information originate` does not work with static route `0.0.0.0/0` as default route [2477]
- Changing redistributed kernel routes does not trigger addition/removal of corresponding OSPF Type-5 LSAs [2389]
- Routing information in the forwarding table is not updated correctly when removing a static route which overlaps a route received via OSPF [2320]
- The OSPF RIB is not updated when the ABR type changes from standard to shortcut, and vice versa [2699]
- Changing the default metric for OSPF server does not result in update on other routers [2586]

OSPF6

- IPv6 routes in the OSPF6 database may not appear in the OSPF RIB until the service is restarted [2891]
- When deleting an OSPF6 interface via RESTCONF, it may remain active in the OSPF6 daemon despite being removed from the TNSR configuration [3481]

RIP

- `key-chain` string is not applied in the routing daemon if configured after RIP is enabled [2878]
Workaround: Disable and enable RIP after making the change.
- RIP `timeout` value is not respected [2796]

SNMP

- There are no changes when using “write” community [2567]

VRRP

- VRRP does not function on an outside NAT interface with a priority of 255 [2419]

Workaround: Set the `priority` of the VR address on the primary router to a value less than 255 yet higher than that of other routers. Enable Accept Mode on the VR address if the VR address will be used by services on TNSR.

VXLAN

- Changes to a VXLAN interface do not apply until the dataplane is restarted [1778]
- VXLAN and OSPF may not work properly if OSPF is configured after VXLAN in the dataplane [2511]

32.3.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

32.4 TNSR 19.12 Release Notes

32.4.1 About This Release

General

- Updated to CentOS 7.7 [2638]

ACL

- Fixed a backend crash when requesting a non-existent ACL via RESTCONF [2613]
- Fixed a backend crash when displaying an ACL with a description in the CLI [2606]

BFD

- Integrated BFD implementation with dynamic routing protocol daemons [2106, 2131]
- Removed redundant BFD configuration parameters from routing daemon configuration, configure options directly in BFD instead [2578]

Counters

- Fixed an issue with invalid interface counter data at first boot. [2572]
- Fixed an issue with multicast counter output containing unicast counter data [2526]

Dataplane

- Fixed error message displayed when attempting to assign more than the available number of CPU cores [2625]
- Enhanced the CPU corelist-workers command to accept ranges of cores [1943]
- Fixed an issue where the value of `ip reassembly max-reassemblies` was ignored if `ip reassembly expire-walk-interval` was also set [2561]
- Added commands to configure dataplane network device receive and transmit descriptors [2020]

DHCP

- Added commands to define custom DHCP options [2774]
- Fixed an error when running `service dhcp reload` [2666]

Host ACLs

- Changed default host ACL ruleset to allow IPv6 traceroute [2627]

Interfaces

- Fixed display of tag rewriting configuration in `show interface` output [2807]
- Fixed IPv6 addresses not being reapplied to an interface when it was disabled and later re-enabled [2648]
- Fixed use of renamed interfaces with bonding [2740]
- Fixed adding interfaces to a bond when they previously had been configured with an IP address [2654]
- Fixed an issue where data may fail to pass through a bond interface after changing its settings [1603]

IPsec

- Fixed an issue with RESTCONF IPsec status data returning every value as a string type [2642]
- Improved IPsec to be thread-safe with multiple workers [1334, 2084]

MAP

- Fixed an issue where IPv6 packets were not translated to IPv4 for MAP domain rules where PSID offset and length are specified [2808]
- Fixed an issue where changing MAP behavior from translate to encapsulate required restarting the dataplane [1779]
- Fixed TCP MSS value not being applied to encapsulated packets in MAP-E mode [1816]

NAT

- Fixed an issue with `show nat deterministic-mappings` returning IPv6 data instead of IPv4 [2887]
- Fixed issues with `show nat sessions` not returning results via RESTCONF or the CLI [2746, 2251]
- Added commands to adjust values of NAT hash buckets and memory [1762, 2611]
- Increased the maximum value of `max-translations-per-user` to 262144 [2612]
- Fixed NAT and ACL permit+reflect rules not working when configured together [2262]

Routing

- Fixed an issue with adding routes to the same destination via different next-hop routers [2407]

Dynamic Routing

- Fixed an issue preventing OS-level interface events/status from being recognized by FRR daemons [2755]
- Fixed an issue with creating `access-list` entries for IPv6 prefixes using the CLI [2624]
- Fixed an issue with creating route map `match peer` entries for IPv6 addresses using the CLI [2623]

BGP

- Fixed setting the `solo` option for BGP neighbors [2826]
- Fixed setting the `maximum-paths` BGP option via CLI [2822]
- Fixed setting the `table-map filter` BGP option via CLI [2821]
- Fixed setting the `route-map` option for BGP network entries via CLI [2820]
- Fixed setting the `backdoor` option for BGP network entries via CLI [2819]
- Fixed the `show route dynamic bgp ipv4 network` command so it does not require a full prefix with mask length [2773]
- Fixed an issue where setting a new BGP `update-delay` timer did not override the previous `peer-wait` value [2772]
- Fixed input validation of the BGP `update-delay` value so it cannot be set larger than `peer-wait` [2771]
- Fixed an issue where BGP would fail to install a received IPv6 route into the routing table [2650]

OSPF

- Added `detail` modifier to `show route dynamic ospf neighbor` which displays more detailed OSPF neighbor information [2742]
- Fixed an issue where an OSPF LSA was not added to the LSDB if there was a dead LSA for same route present [2626]
- Fixed an issue where OSPF did not send LSA-5 messages to a backbone area if an NSSA area session was already established [2559]
- Fixed setting the `timer throttle lsa` value for OSPF in the CLI [2555]

OSPF6

- Added support for OSPFv3 (Also known as OSPF6) to handle OSPF for IPv6 [2517]
 - OSPF6 is now also allowed in the default host ACL ruleset [2668]

RIP

- Added support for RIP (v2 and v1) [2498]
 - RIP is now also allowed in the default host ACL ruleset (UDP port 520) [2657]

SNMP

- Fixed `ifOutUcastPkts` returning value of `rx-bytes` instead of `tx-bytes` [2584]

VRRP

- Added commands to configure interface tracking for VRRP and display its status [2521]
- Fixed an issue where multiple VRs with the same VR ID on a hardware interface (via subinterfaces) could interfere with each other [2865]
- Fixed an issue where a VRRP VR only removes the virtual MAC from an interface when transitioning from master to backup [2842]
- Fixed an issue with using VRRP on bond interfaces [2829]
- Fixed an issue with incorrect VRRP VR behavior with priority 255 and accept mode enabled [2816]
- Added input validation to prevent conflicting VRRP and NAT configurations [2799]
- Fixed an issue where VRRP may fail to add a virtual IP address [2706]

Configuration Changes

Several areas of the configuration were changed. These changes must either be made manually or see [Updating the Configuration Database](#) for information on how to automatically update the configuration using a script included in this update.

- `netgate-bgp`
 - Configuration under `/route-config/dynamic/bgp/routers/router:`
 - * `update-delay-peer-wait` had a constraint added. Its value must be less than or equal to `.. / update-delay-updates`
 - * `address-families/ipv4/unicast/mutliple-path-maximums` was renamed to `multiple-path-maximums` to correct a spelling error
 - * `address-families/ipv6/unicast/mutliple-path-maximums` was renamed to `multiple-path-maximums` to correct a spelling error
 - * `neighbors/neighbor/bidirectional-forwarding-detection` did not have any effect on BGP so it was removed.
- `netgate-ospf`
 - Type definitions

- * Enumerated type `ospf-route-out` had several values removed which are not supported. This type was used in `/route-config/dynamic/ospf/routers/router/distribute-list/out/route-out`
- `netgate-snmp`
 - Type definitions
 - * Enumerated type `snmp-security-level` had several values removed which are not supported. This type is used in `/snmp-config/snmp-access-control/access/access-entry/security-level`
 - * Enumerated type `snmp-security-model` had several values removed which are not supported. This type is used in `/snmp-config/snmp-access-control/access/access-entry/security-model` and `/snmp-config/snmp-access-control/group/group-entry/security-model`
 - * Enumerated type `snmp-context-match` had several values removed which are not supported. This type is used in `/snmp-config/snmp-access-control/access/access-entry/prefix`
- `netgate-ip`
 - Renamed `/ip` to `ip-config` – This only contains IP reassembly settings.

32.4.2 Known Limitations

Upgrade Issues

Warning: Due to a build dependency issue with `librtnl` in TNSR 19.12, installations of TNSR 19.08 upgraded to TNSR 19.12 will not end up with a functional copy of `librtnl`. This library must be linked against the current version of VPP. Since VPP had a version change between 19.08 and 19.12, but the version number of `librtnl` did not change, it is not reinstalled on upgrade with an appropriately relinked copy.

To resolve this problem, manually reinstall the `librtnl` package using a shell prompt:

```
$ sudo yum reinstall librtnl
```

This may also be run from within TNSR by using the `shell` command, for example:

```
tnsr# shell sudo yum reinstall librtnl
```

This problem has been fixed so it will not recur for TNSR 20.02 or later releases which will carry the TNSR version on these packages to ensure they match appropriately. Installations of TNSR versions prior to 19.08 can safely upgrade to 19.12 without encountering this issue as there was a version change in `librtnl` after that time.

Symptoms of this problem include:

- Sporadic VPP and configuration backend crashes.
- VPP failing to forward packets as expected.
- Configured services (e.g. BGP, IPsec, DNS) not functioning correctly due to host stack connectivity being impaired.

Azure

Warning: The TNSR 19.12 release is not compatible with Azure. Instances of TNSR 19.08 running on Azure should not be upgraded until the next release (TNSR 20.02).

ACLs

- ACLs used with `access-list output` do not work on traffic sent to directly connected hosts [2057]
- Accessing very large (100K rules) ACLs repeatedly results in a Clixon crash [2558]

BFD

- Unable to set `delayed` option on an existing BFD session [2709]

CLI

- CLI does not return from shell in certain situations [2651]

Dataplane

- Dataplane auto pinning of worker threads to cores does not follow expected convention [2846]
- Dataplane reports incorrect physical core ID for main thread [2845]
- Systems with multiple CPU sockets using NUMA may experience dataplane issues at startup or when the dataplane is restarted manually [2383]

DHCP

- Unable to delete all DHCP server options at once from CLI [2667]

GRE

- Unable to modify GRE tunnel settings [2698]

HTTP Server / RESTCONF

- HTTP server retains old configuration after TNSR services restart [2453]
- SSL certificate error when the HTTP server is configured with a certificate that uses md5 digest [2403]

Interfaces

- Packets do not pass through a subinterface after the subinterface configuration has been modified [1612]
- Chelsio interfaces crash the dataplane [1896]
- VLAN subinterfaces may not work under KVM using virtio drivers [2189]
- An IPv6 link-local address cannot manually be configured on an interface [2394]

- IPv6 addresses on IPsec or GRE interfaces may not be displayed in `show` command output [2425]
- Bridge domain ARP entries are not displayed in the CLI [2378]
- Bridge domain ARP entries cannot be removed from the CLI [2380]
- Bridge domain MAC age cannot be removed from the CLI [2381]
- Link state always reported as “up” when using `e1000` network drivers [2831]
- `vmxnet3` RSS fails to initialize, cannot pass packets [2576]

Workaround: Set `dataplane dpdk dev <device id> network num-rx-queues 2` in the TNSR CLI and restart the dataplane.

- Cannot add a DHCP client hostname to an existing DHCP client [2557]
Workaround: Remove the dhcp client from the interface and then re-add it with the hostname.
- Re-enabling loopback interface breaks packet forwarding until the dataplane is restarted [2828]
- Subinterface settings are not applied on change without restarting dataplane [2696]
- Unable to create multiple IP QinQ subinterfaces with the same outer vlan tag [2659]
- Configuration of host OS interface clears TNSR TAP interface configuration [2640]

Workaround: Remove and reconfigure the TAP interface.

- On the XG-1537 and other systems with X552 NICs, if one of the **SFP+** (not copper) interfaces does not have an active link when the dataplane is restarted, and presumably during startup, the interface remains down when the link is reconnected. The link lights come on as though the interface is working and the opposing interface shows the correct link state and speed. This has been confirmed with LR and SR SFP+ modules.

If an affected interface has an active link when the dataplane is started, the link can later change to be down/up or removed/reconnected without issue.

Workaround: Restart the dataplane once the links are active.

IPsec

- An IPsec tunnel which was removed and then added back in may take longer than expected to establish [1313]
- An SA ordering issue may prevent IPsec traffic from passing if both endpoints attempt to establish a tunnel at the same time [2391]
- Large packets over IPsec crash VPP and `clixon-backend` [2902]

Workaround: Increase the `default-data-size` buffer size to 16384 and restart the dataplane.

```
tnsr(config)# dataplane buffers default-data-size 16384
tnsr(config)# service dataplane restart
```

MAP

- MAP-T BR cannot translate IPv4 ICMP echo reply to IPv6 [1749]
- Fragmentation of IPv4 packets is performed regardless of configured MAP fragmentation behavior when MAT-T mode is used [1826]
- MAP BR does not send ICMPv6 unreachable messages when a packet fails to match a MAP domain [1869]
- Pre-resolve does not work when MAP-T mode is used [1871]

- MAP BR encapsulates/translate only last fragment when receiving fragmented packets from IPv4 network [1887]

NACM

- Default parameters rule for NACM node `access-operation` and module does not work without explicit settings [2514]

NAT

- `twice-nat` does not work [1023]
- NAT forwarding is not working for `in2out` direction [1039]
- DS-Lite is not functional; B4 router sends encapsulated IPv4-in-IPv6 packets, but AFTR replies with an error [1626]
- NAT forwarding fails with more than one worker thread [2031]

Note: This also affects connectivity to services on TNSR, such as RESTCONF, when the client is not on a directly connected network.

- Router with 1:1 NAT will drop packets with `ttl=2` from input interface [2849]
- VPP service fails if NAT `concurrent-reassemblies` is set to 1 and several fragments arriving to the NAT outside interface [2739]
- ICMP fragments arriving to NAT Inside interface aren't being reassembled by NAT reassembly function [2733]

Neighbor / ARP / NDP

- Packet loss during ARP transaction immediately after Dataplane restart or interface disable/enable [2868]

RESTCONF

- Incorrect BGP configuration is generated when IPv6 address family is configured via REST [2915]
- Adding a user via RESTCONF requires a password even when key is provided [2875]
- Adding MACIP rule via RESTCONF fails [2844]
- Cannot rename an ACL via RESTCONF [2843]
- Deleting ACL rule via RESTCONF crashes Clixon [2841]

Routing

- IPv6 packet loss may be observed between TNSR instances [2382]

Dynamic Routing

- CLI shows that only IPv4 prefix is available within `prefix-list` sequence configuration [2689]
- `route-map` with sequence number 0 can be configured in the CLI but cannot be used [2876]

BGP

- An IPv6 BGP session cannot be established over IPsec or GRE [2429]
- BGP `maximum-path` option for eBGP and iBGP can not be configured simultaneously [2879]
- BGP `network backdoor` feature does not work without service restart [2873]
- Unable to configure BGP distance values via CLI [2869]
- Unable to verify received prefix-list entries via CLI when ORF capability is used [2864]
- `extended-nexthop` capability is not being negotiated between IPv6 BGP peers [2850]
- BGP session soft reset option does not work for IPv6 peers [2833]
Workaround: Reset the connection without soft option.
- `ttl-security hops` value can be set when `ebgp-multihop` is already configured (the options are mutually exclusive) [2832]
- `clixon-backend` fails when loading BGP config with 150k advertised prefixes [2784]
- Displaying a large amount of received or advertised BGP prefixes takes a long time [2778]
- BGP updates for new prefixes are sent every 60 seconds despite configured `advertisement-interval` value [2757]
- TNSR installs additional duplicated `next-hop` entries for multipath routes received via BGP [2935]

OSPF

- OSPF `default-information originate` does not work with static route `0.0.0.0/0` as default route [2477]
- Changing redistributed kernel routes does not trigger addition/removal of corresponding OSPF Type-5 LSAs [2389]
- Routing information in the forwarding table is not updated correctly when removing a static route which overlaps a route received via OSPF [2320]
- The OSPF RIB is not updated when the ABR type changes from standard to shortcut, and vice versa [2699]
- Changing the default metric for OSPF server does not result in update on other routers [2586]

OSPF6

- IPv6 routes in the OSPF6 database may not appear in the OSPF RIB until the service is restarted [2891]

RIP

- `key-chain` string is not applied in the routing daemon if configured after RIP is enabled [2878]
Workaround: Disable and enable RIP after making the change.

SNMP

- SNMP configuration change requires a service restart [2568]
- There are no changes when using “write” community [2567]

VRRP

- VRRP does not function on an outside NAT interface with a priority of 255 [2419]

Workaround: Set the `priority` of the VR address on the primary router to a value less than 255 yet higher than that of other routers. Enable Accept Mode on the VR address if the VR address will be used by services on TNSR.

VXLAN

- Changes to a VXLAN interface do not apply until the dataplane is restarted [1778]
- VXLAN and OSPF may not work properly if OSPF is configured after VXLAN in the dataplane [2511]

32.4.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

32.5 TNSR 19.08 Release Notes

32.5.1 About This Release

Note: TNSR 19.08.1 installation images are identical to 19.08 except that they have the most recent (as of the time it was built) set of updates from CentOS applied instead of the base release version of CentOS 7.6.1810.

There is no need to reinstall 19.08 to reach 19.08.1 using these images as running an update from 19.08 will result in the same, or even newer, CentOS packages.

General

- Fixed removal of SSH authorized-keys entries from user entries in the OS when they are removed from TNSR users [1162]
- Cleaned up extraneous logging messages from the configuration backend [2230]

ACL

- Fixed manual selection of ACL protocol value 0, and renamed it to `any` [2134]
- Fixed setting type and code values for ICMP ACLs [2325, 2426]
- Fixed issues with removing the protocol value from an ACL rule [2252, 2307]
- Expanded TNSR ACL rule protocol choices to any protocol, specified by number [2224]
- Improved performance and display of large ACL rulesets (e.g. 10,000+ ACLs) [2139]

BFD

- Fixed editing unused BFD keys [1891]
- Fixed the BFD `delayed` option [1885]
- Added validation to prevent changing the BFD interface, local address, or peer address since this is not allowed by the dataplane. [1549]
- Fixed administratively disabling BFD via CLI [1883]

CLI

- Improved handling of resizing terminal dimensions [2214]
- Added options to enable and disable command history as well as to set the history size to a given value [2011]

Counters

- Added verbose counter information to `show interface [<if-name>] counters output` [2413]
- Removed redundant `show counters` command [2377]

Dataplane

- Improved memory handling with large ACL rulesets [2442]
- Added dataplane configuration option for `num-crypto-mbufs` [2160]
- Added dataplane configuration options for buffer parameters [2399]
- Fixed `service dataplane restart` potentially causing `clixon_backend` to lose its configuration [1383]

DHCP

- Removed invalid * DHCP logging category [1307]
- Fixed DHCP reservation required value validation so entries cannot be created without a MAC address [1530]

DNS

- Removed invalid `allow_setrd` value from `Unbound access-control` command [1747]
- Fixed handling of local zone hostname and domain when forming A/AAAA and PTR entries [1384]
- Added `outgoing-interface` command to `config-unbound` mode to control how TNSR will originate DNS requests to upstream DNS servers [1884]

GRE

- Fixed routing IPv6 inner traffic over IPv4 outer GRE tunnel [2424]

Host ACLs

- Expanded Host ACL rule protocol choices to any protocol, specified by number [2227]
- Fixed host ACL ICMP rule matching [2217, 2226]
- Fixed duplication of rules in the nftables ruleset when the dataplane restarts [2207]

HTTP Server / RESTCONF

- Fixed handling of the HTTP daemon configuration file when the service is not enabled in TNSR [1153]
- Added new default index and error pages to the HTTP daemon [1531]

Interfaces

- Fixed loopback interfaces responding to ICMP echo requests when in the down state [850]
- Added commands to enable and configure IP reassembly [1302, 1277]
- Changed `show interface` subcommands to be more consistent with other areas of the CLI [2376]

Note: Only one output-limiting keyword may now be specified, and several keywords were renamed to match their corresponding configuration parameters.

- Added the ability to remove a MAC address from an interface, which will return the MAC address back to the native address after a dataplane restart [2310]
- Fixed a clixon crash while executing `show interface lacp` [2438]
- Fixed MAC address change propagation from dataplane to host tap interfaces [1502]
- Fixed QinQ VLAN termination [1550]
- Added `no mtu` interface command to remove the MTU setting and revert to the default value [2021]

IPsec

- Fixed IPv6 traffic traversing an IPv4 IKEv2 IPsec tunnel [2422]
- Fixed IPsec Child SA failures with AES-GCM combined with DPDK cryptodevs (QAT or aesni vdev) [2309]
- Fixed IPsec tunnels with a Child SA using MD5 integrity failing to establish [2505]

- Fixed IPsec tunnels with a Child SA using 3DES encryption failing to establish [2476]
- Added elliptic curve DH group 31 (curve25519, 256 bit) to IPsec proposal choices [2179]

MAP

- Added input validation to enforce MAP `ip6-src-prefix` values [2087]

NACM

- Added improved error messages showing failed paths when access is denied by NACM [2443]
- Changes to interface-related validation now require that users with access to configure interface-related items must also be able to get `/interfaces-state/interface` to read the interface list [2443]

NAT

- Added commands to manage NAT session timeout values [2232]
- Fixed issues with static NAT mappings with defined ports occasionally leading to a clixon-backend crash when restarting [1103]
- Added input validation to prevent deterministic NAT crashes in the dataplane due to incorrect user configuration [1856]

NTP

- Fixed NTP configuration generated for `restrict` lists [1705]

RESTCONF

- Improved information returned in queries for `netgate-system:system-state` [2324]
- Fixed malformed requests causing the API to return unexpected errors for a few seconds while it restarts [2079]

Routing

- Improved handling of route table display with large route tables [506]
- Improved output of `show route table` [2229]
- Fixed handling and display of IPv6 static neighbors [2005]
- Fixed FIB lookup option for static routes [1280]
- Fixed creating static routes with the same next-hop ID in multiple routing tables [2510]

Dynamic Routing

Warning: Commands for BGP and related dynamic routing functionality have been restructured so everything is under `route dynamic`. Changes are extensive and the documentation has been updated to reflect the new commands.

- Added support for OSPF [1895]
- Length of BGP neighbor passwords is now limited to 63 characters [1454]
- Fixed removal of IPv6 next-hop peer address from a route map [2304]
- Fixed BGP advertisement of connected routes after interface status changes [746, 2409]
- Changed BGP status commands for `summary`, `neighbors`, and `network` to require an address family [2367]
- Fixed handling of BGP debug commands [2385]
- Fixed handling of BGP `maximum-prefix` configuration parameter [859]
- Fixed session handling when `maximum-prefix-limit` is exceeded [858]
- Fixed handling of IPv6 static routes in the dynamic routing manager (`zebra`) [2279]
- Cleaned up commands for unsupported dynamic routing features [2312]
- Fixed handling of BGP `import-check` [781]
- Fixed handling of routes from `aggregate-address` via `next-hop 0.0.0.0` [832]
- Eliminated unnecessary restarts of the dynamic routing daemons when making changes [1758]
- Fixed positive relative metric adjustments in route-maps [2493]
- Fixed displaying specific IPv6 BGP networks by address [2479]
- Fixed configuring a BGP IPv6 aggregate address with `summary-only` option [2509]

SNMP

- Support for SNMP monitoring has been added, see *Simple Network Management Protocol* for implementation details [2286]

Updates

- Fixed handling of `igb_uio` module during an upgrade which also updates the kernel [2216]

VRRP

- Support for VRRP has been added, see *Virtual Router Redundancy Protocol* for implementation details and limitations [1894]

VXLAN

- Fixed configuration of alternate VXLAN encapsulation routing tables [1872]

32.5.2 Known Limitations

Updates

- The UIO drivers may not be present in the correct directory after a kernel upgrade. Since the UIO drivers are kernel-specific, they must be rebuilt after any change in the kernel [2216].

To work around this issue, force a reinstall of the DPDK package which will rebuild the UIO drivers and place them in the appropriate location for the updated kernel:

```
$ sudo yum -y reinstall dpdk
```

This procedure will not be necessary when upgrading to future releases from 19.08.

ACLs

- ACLs used with `access-list output` do not work on traffic sent to directly connected hosts [2057]

BFD

- BFD does not integrate with BGP [2106]

BGP

- Changing `update-source` from an IP address to `loop1` allows a session to establish but remote prefixes do not appear in the FIB until reboot [1104]

Counters

- At first boot, interface counter data may be invalid. [2572]

Workaround: Restart the dataplane to correct this problem until next reboot.

Hardware

- Systems with multiple CPU sockets using NUMA may experience dataplane issues at startup or when the dataplane is restarted manually [2383]

HTTP Server / RESTCONF

- HTTP server retains old configuration after TNSR services restart [2453]
- SSL certificate error when the HTTP server is configured with a certificate that uses md5 digest [2403]

Interfaces

- Non-LACP bond interfaces may experience packet drops when a bond member interface is down [1603]
- Packets do not pass through a subinterface after the subinterface configuration has been modified [1612]
- Chelsio interfaces crash the dataplane [1896]

- VLAN subinterfaces may not work under KVM using virtio drivers [2189]
- An IPv6 link-local address cannot manually be configured on an interface [2394]
- IPv6 addresses on IPsec or GRE interfaces may not be displayed in `show` command output [2425]
- Bridge domain ARP entries are not displayed in the CLI [2378]
- Bridge domain ARP entries cannot be removed from the CLI [2380]
- Bridge domain MAC age cannot be removed from the CLI [2381]

IPsec

- An IPsec tunnel which was removed and then added back in may take longer than expected to establish [1313]
- An SA ordering issue may prevent IPsec traffic from passing if both endpoints attempt to establish a tunnel at the same time [2391]

MAP

- MAP-T BR cannot translate IPv4 ICMP echo reply to IPv6 [1749]
- MAP behavior cannot be changed from translate to encapsulate without restarting the dataplane [1779]
- TCP MSS value is not applied to encapsulated packets when MAP-E mode is used [1816]
- Fragmentation of IPv4 packets is performed regardless of configured MAP fragmentation behavior when MAT-T mode is used [1826]
- MAP BR does not send ICMPv6 unreachable messages when a packet fails to match a MAP domain [1869]
- Pre-resolve does not work when MAP-T mode is used [1871]
- MAP BR encapsulates/translate only last fragment when receiving fragmented packets from IPv4 network [1887]

NACM

- Default parameters rule for NACM node `access-operation` and `module` does not work without explicit settings [2514]

NAT

- `twice-nat` does not work [1023]
- NAT forwarding is not working for `in2out` direction [1039]
- DS-Lite is not functional; B4 router sends encapsulated IPv4-in-IPv6 packets, but AFTR replies with an error [1626]
- NAT forwarding fails with more than one worker thread [2031]

Note: This also affects connectivity to services on TNSR, such as RESTCONF, when the client is not on a directly connected network.

- Connections to and from the TNSR host are included in NAT sessions when connecting through an interface with `ip nat outside` [1892] [1979]
- NAT and ACL permit+reflect rules do not work together [2262]

Routing

- Deleting a non-empty route table fails with an error and the table remains in the configuration, but it cannot be changed afterward [1241]

Workaround: Remove all routes from the table before deleting. Alternately, copy the running configuration to startup and restart TNSR, which will make the route table appear again so the routes and then the table can be removed.

- Cannot add multiple routes to the same destination using different next hops [2407]

Dynamic Routing

- An IPv6 BGP session cannot be established over IPsec or GRE [2429]
- iBGP router advertises redistributed static IPv6 routes with next-hop value set to link-local address [2478]
- OSPF `default-information originate` does not work with static route `0.0.0.0/0` as default route [2477]
- Changing redistributed kernel routes does not trigger addition/removal of corresponding OSPF Type-5 LSAs [2389]
- Routing information in the forwarding table is not updated correctly when removing a static route which overlaps a route received via OSPF [2320]

VRRP

- VRRP does not function on an outside NAT interface [2419]

VXLAN

- Changes to a VXLAN interface do not apply until the dataplane is restarted [1778]
- VXLAN and OSPF may not work properly if OSPF is configured after VXLAN in the dataplane [2511]

32.5.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

32.6 TNSR 19.05 Release Notes

32.6.1 About This Release

General

- Added support for QAT C62x crypto devices [1718]
- Added service management RPCs to data model [1715]

ACL

- Fixed creating an ACL using only a description [1558]
- Fixed creating an empty ACL [1735]
- Fixed creating an ACL rule with a destination port [1796]

BGP

- IPv6 BGP neighbors get entered as `peer-groups` only in `bgpd.conf` [1190]
- Removed deprecated `neighbor <peer> interface <if> BGP` command [2113]
- Restructured BGP address family configuration to accommodate IPv4 and IPv6 [2049]
- Removed option to create a new neighbor inside address family mode [2194]
- Removed `route-map set metric` options for `+/- rtt` and `+/- metric` as they were not supported as users expected in FRR [2191]

CLI

- `[no] shutdown` style syntax has been removed. Use `enable` and `disable`, or `no enable` [1652]
- Fixed paging issues in output that could lead to incorrect or missing output after certain actions taken with multi-page output (e.g. pressing `q` or `Enter` at a `More` prompt) [1774, 1773]
- The CLI now stores command history between sessions (*Command History*) [514, 1949]
- Standardized commands to enabled coredumps for services, and added support for coredumps from `ike`, `unbound`, `http`, and `ntp` (*Diagnosing Service Issues*) [1831]
- Fixed `ping` so it can work with IPv6 source addresses [2004]
- Improved CLI performance when working with large lists [2127]
- Increased timeout for `package` commands to allow longer processes to finish completely, such as upgrades [1768]

Dataplane

- Fixed writing default values to the dataplane configuration when no dataplane options are set in the configuration [1982]
- Fixed dataplane crashes when using NAT with forwarding enabled with certain packet combinations when the protocol is not ICMP, TCP, or UDP [1998]
- Mellanox support: Added option to disable multi-segment buffers in the dataplane [2022]
- Fixed an error when configuring a dataplane crypto device without first configuring the UIO driver [1812]
- Added worker thread and core affinity options [1675]
- Added an option to set custom interface names for dataplane interfaces [2062]
- Added commands to configure dataplane statistics segment options [2199]

DHCP

- The DHCP server can now function when an interface is configured as a DHCP client [1801]
- DHCP server no longer uses link-local interface IP addresses (169.254.0.x) as a source address for DHCP packets or as a DHCP Server Identifier [1222]
- Removed incorrect references to the `netgate-interface` module from the DHCP server CLI specification API paths [1810]
- Removed redundant `ipv4` forms of DHCP-related commands [1557]

Host ACLs

- Added support for Host ACLs to control traffic to host OS interfaces using `nftables` [1651]

HTTP Server / RESTCONF

- `nginx` now behaves as expected with `authentication type none` and TLS [1086]

Warning: This mode is intended only for testing, not production use.

- Fixed RESTCONF `get of /restconf/data/` so it properly returns state data [1534]

Installer

- Improved consistency in post-install login procedures across all TNSR platforms [2013]
- Fixed installation issues on hardware that has an eMMC device, such as the SG-5100 [2048]
- Fixed the default NACM configuration when installing from ISO [2133]
- Added Infiniband/rdma packages to the default installation [2201]

Interfaces

- An interface can now be deleted if has had an ACL or MACIP applied [1177, 1178]
- MACIP ACLs no longer remain in the interface configuration after being removed [1179]
- Bond interfaces in LACP mode no longer send LACPDUs when configured for passive mode [1614]
- VLAN tag rewrite settings have been relocated to interfaces, as they do not require a subinterface [1344]
- VXLAN validation now properly reflects that a VXLAN entry requires a VNI [1821]
- GRE and VXLAN now create interfaces on the host [1999]
- Fixed display of link speeds for 40G and 100G interfaces [1867]
- Removed unused “Admin status” field from state information for host interfaces [1864]
- Fixed interface counters for Mellanox interfaces [2039]
- Fixed interface counters for IPsec interfaces [2075]
- VLAN tag-rewrite attributes are now included in `show interface output` [1654]

- Changed `show interfaces` to output interfaces in a consistent order [2046]
- Fixed a problem with neighbor location (ARP/NA) when VLAN tags are present [1326]
- Fixed default handling of VMXNET3 interfaces [1703]

IPsec

- Added support for the 3DES encryption algorithm in IPsec proposals [1444]

NACM

- NACM now supports all access operations and module restrictions (*NACM Rule Lists*) [1809]
- The method to manually disable NACM has changed. *Regaining Access if Locked Out by NACM* has been updated to reflect the new method [1750, 1752]

NAT

- DS-Lite B4 endpoint is now shown in the output of `show dslite` [1625]
- NAT sessions may now be queried with `show nat sessions [verbose]` (*View NAT Sessions*) [975, 1456]
- Fixed issues with NAT and multiple worker threads [1844]
- NAT mode deletion is now properly respected in VPP startup configuration after TNSR services restart [1017]
- Fixed incorrect NAT static mappings being added when a new rule differed from an existing rule only by the `port-local` value [1100]

32.6.2 Known Limitations

Updates

- The UIO drivers may not be present in the correct directory after a kernel upgrade. Since the UIO drivers are kernel-specific, they must be rebuilt after any change in the kernel [2216]

To work around this issue, force a reinstall of the DPDK package which will rebuild the UIO drivers and place them in the appropriate location for the updated kernel:

```
$ sudo yum -y reinstall dpdk
```

ACLs

- ACLs used with `access-list output` do not work on traffic sent to directly connected hosts [2057]

BFD

- Attempting to change a BFD local/peer address fails [1549]
- BFD cannot be administratively disabled via CLI [1883]
- The BFD `delayed` option does not work [1885]

- An unused BFD `conf-key` cannot be modified [1891]
- BFD does not integrate with BGP [2106]

BGP

- TNSR does not send BGP updates without restarting service with `redistribute from connected` option [746]
- Route with `aggregate-address` via `next-hop 0.0.0.0` does not appear in TNSR route table [832]
- BGP sessions may fail to establish or rapidly reconnect when receiving more prefixes than defined by `maximum-prefix limit` [858]
- The `maximum-prefix restart` command does not work [859]
- TNSR installs multiple paths for received routes even though support for multiple paths is not enabled [885]
Workaround: Run `systemctl reset-failed frr` from the shell to clear the error which will allow the BGP service to start again.
- Changing `update-source` from an IP address to `loop1` allows a session to establish but remote prefixes do not appear in the FIB until reboot [1104]
- BGP `import-check` feature does not work [781]
- Logs may include spurious BGP message `binary API client 'route_daemon' died` which do not affect BGP routing [1714]

CLI

- `show route table` causes the backend to die with large numbers of routes in the table [506]
For example, this crash happens with a full BGP feed.
- Using `service dataplane restart` can cause `clixon_backend` to lose its configuration [1383]
- Large lists (e.g. 10,000+ ACLs) can cause significant delays in related CLI operations [2139]

DHCP

- Adding a DHCP reservation without a MAC address causes Kea to fail and the entry cannot be removed [1530]
Workaround: A MAC address is required for DHCP reservations, so always enter a MAC address when creating an entry.
- Configuring Kea to log all names with `*` does not work [1307]
Workaround: Configure each name separately instead of using a wildcard.

DNS

- Local zone FQDN handling for forward (A) and reverse (PTR) data is inconsistent, only allowing one or the other to work as expected for a given FQDN [1384]
- Using the `allow_setrd` attribute for `access-control` entries causes `unbound` to fail [1747]
- Unbound requires a default route in the host OS to resolve [1884]

Host ACLs

- Host ACL entries are duplicated after a dataplane restart [2207]

HTTP Server / RESTCONF

- HTTP server runs even though it's not configured to run after TNSR services restart [1153]
Workaround: Manually stop the `nginx` service using `systemctl`.
- RESTCONF query replies may contain CDATA tags in JSON [1463]
- Adding an ACL rule entry via RESTCONF may appear to add a duplicate ACL [1238]

Interfaces

- Loopback interface responds to ICMP echo from an outside host even when in a *Down* state [850]
- Non-LACP bond interfaces may experience packet drops when a bond member interface is down [1603]
- MAC address changes on dataplane interfaces are not reflected on the host tap interface until the dataplane is restarted [1502] Workaround: Restart the dataplane after changing an interface MAC address.
- Bond interface MAC addresses do not match their host tap interface unless a MAC address is explicitly set at creation [1502]
Workaround: Set the MAC address when creating the bond interface.
- Packets do not pass through a subinterface after the subinterface configuration has been modified [1612]
- QinQ VLAN termination is not working [1550]
- Chelsio interfaces crash the dataplane [1896]
- VLAN subinterfaces may not work under KVM using virtio drivers [2189]

IPsec

- An IPsec tunnel which was removed and then added back in may take longer than expected to establish [1313]

MAP

- MAP-T BR cannot translate IPv4 ICMP echo reply to IPv6 [1749]
- MAP security check configuration differs between the dataplane and CLI [1777]
- MAP behavior cannot be changed from translate to encapsulate without restarting the dataplane [1779]
- TCP MSS value is not applied to encapsulated packets when MAP-E mode is used [1816]
- Fragmentation of IPv4 packets is performed regardless of configured MAP fragmentation behavior when MAT-T mode is used [1826]
- MAP BR does not send ICMPv6 unreachable messages when a packet fails to match a MAP domain [1869]
- Pre-resolve does not work when MAP-T mode is used [1871]
- MAP BR encapsulates/translate only last fragment when receiving fragmented packets from IPv4 network [1887]

NACM

- Permitted default read and write operations cannot be executed if default exec policy is set to `deny` [1158]

NAT

- `twice-nat` does not work [1023]
- NAT forwarding is not working for `in2out` direction [1039]
- NAT static mapping with defined ports leads to `clixon-backend` crash after restart [1103]
- DS-Lite is not functional; B4 router sends encapsulated IPv4-in-IPv6 packets, but AFTR replies with an error [1626]
- NAT forwarding fails with more than one worker thread [2031]
Note: This also affects connectivity to services on TNSR, such as RESTCONF, when the client is not on a directly connected network.
- Deterministic NAT crashes the dataplane [1856]
- Connections to and from the TNSR host are included in NAT sessions when connecting through an interface with `ip nat outside` [1892] [1979]

Neighbors

- IPv6 static neighbors entries do not work [2005]

NTP

- NTP restrictions for prefixes do not work [1705]

RESTCONF

- A malformed request may cause the API to return unexpected errors for a few seconds while it restarts [2079]

Routing

- Deleting a non-empty route table fails with an error and the table remains in the configuration, but it cannot be changed afterward [1241]
Workaround: Remove all routes from the table before deleting. Alternately, copy the running configuration to startup and restart TNSR, which will make the route table appear again so the routes and then the table can be removed.

User Management

- When deleting a user key from the running configuration it is not removed from the user's `authorized_keys` file [1162]
Workaround: Manually edit the `authorized_keys` file for the user and remove the key.

VXLAN

- Changes to a VXLAN interface do not apply until the dataplane is restarted [1778]
- Alternate VXLAN encapsulation routing tables cannot be configured [1872]

32.6.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

32.7 TNSR 19.02.1 Release Notes

32.7.1 About This Release

This is a maintenance release for TNSR software version 19.02 with bug fixes and Azure support.

See also:

For more information on changes in TNSR version 19.02, see *[TNSR 19.02 Release Notes](#)*.

General

- TNSR is now supported on Azure [974]

NAT

- Fixed a problem with removing MAP entries after restarting TNSR [1653]

32.7.2 Known Limitations

ACL

- Attempting to create an ACL containing only a description fails [1558]
Workaround: Define one or more rules on the ACL.

BFD

- Attempting to change a BFD local/peer address fails [1549]

BGP

- TNSR does not send BGP updates without restarting service with `redistribute from connected` option [746]
- Route with `aggregate-address` via `next-hop 0.0.0.0` does not appear in TNSR route table [832]
- BGP sessions may fail to establish or rapidly reconnect when receiving more prefixes than defined by `maximum-prefix limit` [858]
- The `maximum-prefix restart` command does not work [859]
- TNSR installs multiple paths for received routes even though support for multiple paths is not enabled [885]
Workaround: Run `systemctl reset-failed frr` from the shell to clear the error which will allow the BGP service to start again.
- Changing `update-source` from an IP address to `loop1` allows a session to establish but remote prefixes do not appear in the FIB until reboot [1104]
- IPv6 BGP neighbors get entered as `peer-groups` only in `bgpd.conf` [1190]
- BGP `import-check` feature does not work [781]

CLI

- `show route table` causes the backend to die with large numbers of routes in the table [506]
For example, this crash happens with a full BGP feed.
- Using `service dataplane restart` can cause `clixon_backend` to lose its configuration [1383]

DHCP

- The DHCP server does not function if an interface is configured as a DHCP client [1801]
Corrected in the next release under development (19.05).
- DHCP server uses default VPP interface IP address (169.254.0.x) as a source address for DHCP packets and as a DHCP Server Identifier [1222]
- Adding a DHCP reservation without a MAC address causes Kea to fail and the entry cannot be removed [1530]
Workaround: A MAC address is required for DHCP reservations, so always enter a MAC address when creating an entry.
- Configuring Kea to log all names with `*` does not work [1307]
Workaround: Configure each name separately instead of using a wildcard.

DNS

- Local zone FQDN handling for forward (A) and reverse (PTR) data is inconsistent, only allowing one or the other to work as expected for a given FQDN [1384]

HTTP Server / RESTCONF

- `nginx` does not behave as expected with `authentication type none` and TLS [1086]

This mode is primarily for testing and not production use.

Workaround: Use password or certificate-based authentication for RESTCONF.

- HTTP server runs even though it's not configured to run after TNSR services restart [1153]

Workaround: Manually stop the `nginx` service using `systemctl`.

- RESTCONF get of `/restconf/data/` does not properly return state data [1534]
- RESTCONF query replies may contain CDATA tags in JSON [1463]
- Adding an ACL rule entry via RESTCONF may appear to add a duplicate ACL [1238]

Interfaces

- Loopback interface responds to ICMP echo from an outside host even when in a *Down* state [850]
- Unable to delete an interface if has had an ACL or MACIP applied [1177, 1178]
Workaround: Remove the entire ACL or MACIP entry. Then, the interface may be removed.
- MACIP ACL remains in the interface configuration after being removed [1179]
- Bond interfaces in LACP mode will send LACPDUs even when configured for passive mode [1614]
- Non-LACP bond interfaces may experience packet drops when a bond member interface is down [1603]
- MAC address change on tap interfaces may not be reflected in the dataplane until the dataplane is restarted [1502]

Workaround: Restart the dataplane after changing an interface MAC address.

- MAC address change on bond interfaces may not be reflected in the dataplane until the dataplane is restarted [1502]
Workaround: Set the MAC address when creating the bond interface.
- VLAN tag rewrite settings are only available in subinterfaces [1344]
- Packets do not pass through a subinterface after the subinterface configuration has been modified [1612]
- QinQ VLAN termination is not working [1550]
- ARP replies received from another host on a VLAN subinterface are not processed correctly [1326]

IPsec

- An IPsec tunnel which was removed and then added back in may take longer than expected to establish [1313]

NACM

- Permitted default read and write operations cannot be executed if default exec policy is set to `deny` [1158]

NAT

- `twice-nat` does not work [1023]
- NAT mode is not deleted from VPP startup configuration after TNSR services restart [1017]
- NAT forwarding is not working for `in2out` direction [1039]
- NAT static mappings are not added as expected when only the `port-local` value differs [1100]
- NAT static mapping with defined ports leads to `clixon-backend` crash after restart [1103]
- DS-Lite is not functional; B4 router sends encapsulated IPv4-in-IPv6 packets, but AFTR replies with an error [1626]
- DS-Lite B4 endpoint is not shown by `show dslite` command [1625]
- Unable to view a list of NAT sessions [975, 1456]

Routing

- Deleting a non-empty route table fails with an error and the table remains in the configuration, but it cannot be changed afterward [1241]

Workaround: Remove all routes from the table before deleting. Alternately, copy the running configuration to startup and restart TNSR, which will make the route table appear again so the routes and then the table can be removed.

User Management

- When deleting a user key from the running configuration it is not removed from the user's `authorized_keys` file [1162]

Workaround: Manually edit the `authorized_keys` file for the user and remove the key.

32.7.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

32.8 TNSR 19.02 Release Notes

32.8.1 About This Release

Warning: A number of commands were reorganized with this release, more information will be noted below in individual sections. If a command that worked in a previous release is no longer present, it has most likely been changed to a more logical and consistent location.

Warning: RESTCONF queries now require a namespace in the format of `module:name` where only the `name` was required in previous versions. To locate the correct `module:name` combination, see [API Endpoints](#).

General

- The data models have been updated with more consistent naming and locations
- Introduced a YANG `id` type for `name` fields [1318]
- Miscellaneous code cleanup and refactoring for stability and performance improvements [1516] [1571]
- Updated to CentOS 7.6 [1335]
- Updated build to use gcc 7 [1147]
- Fixed a potential crash when listing packages [1312]
- Improved handling of package versions to better handle situations where a dependency update requires reinstalling related packages [950]

BGP

- BGP commands reorganized under `route dynamic` for configuration and `show route dynamic` for status. See [Commands](#) and [Border Gateway Protocol](#). [1369]
- FRR updated to 6.0.x

CLI

- The configuration database commands have been reorganized under `configuration` for making changes, such as `copy`, and under `show configuration` for viewing the contents of a configuration. See [Commands](#) and [Configuration Database](#). [1347]
- Fixed `system location` text handling when the value contains whitespace [1584]

Dataplane

- Updated DPDK `igb_uio` module to v19.02 [842]

DHCP Server

- Updated Kea to 1.4.0-P1 [1239]

DNS

- Fixed removal of `access-control` entries in the CLI [1417]

Host

- Fixed inconsistent behavior of `host interface` commands [1611]
- Added a default set of `nftables` rules to limit inbound traffic to the host [476]

Interfaces

- Several interface-related configuration commands have been moved under the `interface` command for better consistency. These include: `bridge`, `loopback`, `memif`, `subif`, and `tap`. See [Commands](#) and [Types of Interfaces](#) [1336]
- Added support for [Bonding Interfaces](#) for link aggregation and redundancy, including support for LACP [1025]
- Fixed display of a single TAP interface [1554]
- Fixed state data returned from a GET request for `/netgate-interface:interfaces-state/interface` [1553]
- Corrected validation of `memif` socket ID to exclude 0 which is reserved, and enforce a maximum of 4294967294 [1527]
- Corrected validation of bridge domain ID to exclude 0 which is reserved, and enforce a maximum of 16777215 [1526]
- Fixed handling of non-default routing tables assigned to interfaces at startup [1518]
- Removed unused container `/interfaces-config/interface/tunnel` from data model [1427]
- Fixed `subif` commands `outer-dot1q any` and `outer-dot1ad any` [1552] [1352]
- Fixed subinterfaces failing after changing configuration [1346]
- Removed the `untagged` command from `subif` as it was non-functional and unnecessary (use the parent interface for untagged traffic) [1345]

NAT

- Added support for [MAP-T and MAP-E BR](#) [1399]

RESTCONF

Warning: RESTCONF queries now require a namespace in the format of `module:name` where only the `name` was required in previous versions. To locate the correct `module:name` combination, see [API Endpoints](#).

- Fixed RESTCONF calls for RPCs returning error 400 despite succeeding [1511]

Routing

- Fixed removing a route table reporting failure when the operation succeeded [1515]

32.8.2 Known Limitations

ACL

- Attempting to create an ACL containing only a description fails [1558]

Workaround: Define one or more rules on the ACL.

BFD

- Attempting to change a BFD local/peer address fails [1549]

BGP

- TNSR does not send BGP updates without restarting service with `redistribute from connected` option [746]
- Route with `aggregate-address` via `next-hop 0.0.0.0` does not appear in TNSR route table [832]
- BGP sessions may fail to establish or rapidly reconnect when receiving more prefixes than defined by `maximum-prefix limit` [858]
- The `maximum-prefix restart` command does not work [859]
- TNSR installs multiple paths for received routes even though support for multiple paths is not enabled [885]
Workaround: Run `systemctl reset-failed frr` from the shell to clear the error which will allow the BGP service to start again.
- Changing `update-source` from an IP address to `loop1` allows a session to establish but remote prefixes do not appear in the FIB until reboot [1104]
- IPv6 BGP neighbors get entered as `peer-groups` only in `bgpd.conf` [1190]
- BGP `import-check` feature does not work [781]

CLI

- `show route table` causes the backend to die with large numbers of routes in the table [506]
For example, this crash happens with a full BGP feed.
- Using `service dataplane restart` can cause `clixon_backend` to lose its configuration [1383]

DHCP

- DHCP server uses default VPP interface IP address (169.254.0.x) as a source address for DHCP packets and as a DHCP Server Identifier [1222]
- Adding a DHCP reservation without a MAC address causes Kea to fail and the entry cannot be removed [1530]
Workaround: A MAC address is required for DHCP reservations, so always enter a MAC address when creating an entry.
- Configuring Kea to log all names with `*` does not work [1307]
Workaround: Configure each name separately instead of using a wildcard.

DNS

- Local zone FQDN handling for forward (A) and reverse (PTR) data is inconsistent, only allowing one or the other to work as expected for a given FQDN [1384]

HTTP Server / RESTCONF

- `nginx` does not behave as expected with `authentication type none` and TLS [1086]

This mode is primarily for testing and not production use.

Workaround: Use password or certificate-based authentication for RESTCONF.

- HTTP server runs even though it's not configured to run after TNSR services restart [1153]

Workaround: Manually stop the `nginx` service using `systemctl`.

- RESTCONF `get of /restconf/data/` does not properly return state data [1534]
- RESTCONF query replies may contain CDATA tags in JSON [1463]
- Adding an ACL rule entry via RESTCONF may appear to add a duplicate ACL [1238]

Interfaces

- Loopback interface responds to ICMP echo from an outside host even when in a *Down* state [850]
- Unable to delete an interface if has had an ACL or MACIP applied [1177, 1178]
Workaround: Remove the entire ACL or MACIP entry. Then, the interface may be removed.
- MACIP ACL remains in the interface configuration after being removed [1179]
- Bond interfaces in LACP mode will send LACPDUs even when configured for passive mode [1614]
- Non-LACP bond interfaces may experience packet drops when a bond member interface is down [1603]
- MAC address change on tap interfaces may not be reflected in the dataplane until the dataplane is restarted [1502]

Workaround: Restart the dataplane after changing an interface MAC address.

- MAC address change on bond interfaces may not be reflected in the dataplane until the dataplane is restarted [1502]
Workaround: Set the MAC address when creating the bond interface.
- VLAN tag rewrite settings are only available in subinterfaces [1344]
- Packets do not pass through a subinterface after the subinterface configuration has been modified [1612]
- QinQ VLAN termination is not working [1550]
- ARP replies received from another host on a VLAN subinterface are not processed correctly [1326]

IPsec

- An IPsec tunnel which was removed and then added back in may take longer than expected to establish [1313]

NACM

- Permitted default read and write operations cannot be executed if default exec policy is set to `deny` [1158]

NAT

- `twice-nat` does not work [1023]
- NAT mode is not deleted from VPP startup configuration after TNSR services restart [1017]
- NAT forwarding is not working for `in2out` direction [1039]
- NAT static mappings are not added as expected when only the `port-local` value differs [1100]
- NAT static mapping with defined ports leads to `clixon-backend` crash after restart [1103]
- DS-Lite is not functional; B4 router sends encapsulated IPv4-in-IPv6 packets, but AFTR replies with an error [1626]
- DS-Lite B4 endpoint is not shown by `show dslite` command [1625]
- Unable to view a list of NAT sessions [975, 1456]

Routing

- Deleting a non-empty route table fails with an error and the table remains in the configuration, but it cannot be changed afterward [1241]

Workaround: Remove all routes from the table before deleting. Alternately, copy the running configuration to startup and restart TNSR, which will make the route table appear again so the routes and then the table can be removed.

User Management

- When deleting a user key from the running configuration it is not removed from the user's `authorized_keys` file [1162]

Workaround: Manually edit the `authorized_keys` file for the user and remove the key.

32.8.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

32.9 TNSR 18.11 Release Notes

32.9.1 About This Release

Access Lists (ACLs)

- Added a description field to ACL rule entries [1195]
- Fixed issues with numerical sorting of ACL entries in `show` output [1255]
- Fixed issues with order of installed ACL rules in the dataplane with large sequence numbers [1270]

Authentication & Access Control

- Removed users from the TNSR configuration so they are stored/managed directly in the host operating system, which eliminates any chance to be out of sync [1067]
- Fixed issues with deleting NACM rule lists [1137]

BGP

- Fixed an issue where the BGP service could not restart more than three times in a row [902]
- Added `bgp clear` command to clear active BGP sessions [923]

Bridge

- Fixed a problem where the TNSR CLI incorrectly allowed multiple bridge interfaces to have `bvi` set [984]

CLI

- Fixed a problem where applied `dataplane` commands were not immediately present in the running configuration database until another change was made [1099]
- Fixed a problem where the candidate configuration database could not be emptied with the `clear` command [1066]

Hardware & Installation

- Added an ISO image to install TNSR on supported hardware [1364]
- Added support for VMware installations [1026]
- Added support for Mellanox network adapters [1268]

Interfaces

- Fixed interface link speed displaying incorrectly in CLI and RESTCONF [672]
- Fixed issues with duplicate entries being generated in the dataplane interface configuration [1243]

Host

- Added the ability to configure host OS management interfaces in the CLI [260, 261, 262]
- Fixed issues with `ping` command parameter parsing [1133]
- Fixed issues specifying a source address with `ping` [1134]

IPsec

- Fixed issues with IPsec tunnels failing to establish after a dataplane restart [1138]

NAT

- Changed the default NAT mode to `endpoint-dependent` [1079]
- Fixed creating a `twice-nat` pool [972]
- Fixed creating `out-to-in-only` static mappings [976]
- Fixed NAT reassembly for ICMP packets [990]
- Fixed fragment limitations for NAT reassembly [1065]
- Added support for deterministic NAT [360]

NTP

- Fixed issues with the `ntp restrict` command [1163]

RESTCONF

- Fixed validation when submitting invalid MAC addresses via RESTCONF [1197]
- Fixed validation when submitting invalid IP addresses via RESTCONF [1199]

VLAN/Subinterfaces

- Fixed issues where daemons such as Kea and `ntpd` did not correctly form configuration file references to subinterface names [1150]
- Fixed issues with clients on subinterface networks from receiving return traffic that passes through TNSR [1152]
The upstream VPP issue causing this has been fixed, but an additional source of problems in this area is that the `dot1q` setting for a subinterface must use `exact-match` to communicate properly with hosts on the VLAN. Ensure subinterfaces are configured to use this property.

32.9.2 Known Limitations

Authentication & Access Control

BGP

- TNSR does not send BGP updates without restarting service with `redistribute from connected` option [746]
- Route with `aggregate-address` via `next-hop 0.0.0.0` does not appear in TNSR route table [832]
- BGP sessions may fail to establish or rapidly reconnect when receiving more prefixes than defined by `maximum-prefix limit` [858]
- The `maximum-prefix restart` command does not work [859]
- TNSR installs multiple paths for received routes even though support for multiple paths is not enabled [885]
Workaround: Run `systemctl reset-failed frr` from the shell to clear the error which will allow the BGP service to start again.
- Changing `update-source` from an IP address to `loop1` allows a session to establish but remote prefixes do not appear in the FIB until reboot [1104]

- IPv6 BGP neighbors get entered as `peer-groups` only in `bgpd.conf` [1190]
- `peer-group` attribute `remote-as` does not get into FRR `bgpd.conf` [1272]

CLI

- `show route table` causes the backend to die with large numbers of routes in the table [506]
For example, this crash happens with a full BGP feed.

DHCP

- A single IP address can be set in a pool range, but the DHCP daemon requires a start/end IP address or a prefix [1208]
Workaround: Configure a pool with a start and end address or prefix.
- DHCP server uses default VPP interface IP address (169.254.0.x) as a source address for DHCP packets and as a DHCP Server Identifier [1222]
- Unable to delete DHCPv4 options specified within the pool configuration [1267]

HTTP Server / RESTCONF

- `nginx` does not behave as expected with `authentication type none` and TLS [1086]
This mode is primarily for testing and not production use.
Workaround: Use password or certificate-based authentication for RESTCONF.
- HTTP server runs even though it's not configured to run after TNSR services restart [1153]
Workaround: Manually stop the `nginx` service using `systemctl`.

Interfaces

- Loopback interface responds to ICMP echo from an outside host even when in a *Down* state [850]
- Unable to delete an interface if has had an ACL or MACIP applied [1177, 1178]
Workaround: Remove the entire ACL or MACIP entry. Then, the interface may be removed.
- MACIP ACL remains in the interface configuration after being removed [1179]

NAT

- `twice-nat` does not work [1023]
- NAT mode is not deleted from VPP startup configuration after TNSR services restart [1017]
- NAT forwarding is not working for `in2out` direction [1039]
- NAT static mappings are not added as expected when only the `port-local` value differs [1100]
- NAT static mapping with defined ports leads to `clixon-backend` crash after restart [1103]
- PAT dynamic sessions limited to 100 entries per address [1303]
This is the default limit per user in VPP and will be configurable in the next release.

Routing

- Deleting a non-empty route table fails with an error and the table remains in the configuration, but it cannot be changed afterward [1241]

Workaround: Remove all routes from the table before deleting. Alternately, copy the running configuration to startup and restart TNSR, which will make the route table appear again so the routes and then the table can be removed.

User Management

- When deleting a user key from the running configuration it is not removed from the user's `authorized_keys` file [1162]

Workaround: Manually edit the `authorized_keys` file for the user and remove the key.

32.9.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

32.10 TNSR 18.08 Release Notes

32.10.1 About This Release

Authentication & Access Control

- Added support for NETCONF Access Control Model (NACM) management.

NACM provides group-based controls to selectively allow command access for users. Users are authenticated by other means (e.g. RESTCONF certificates or users, CLI user) and then mapped to groups based on username.

- Added default configurations for NACM for different platforms [891]

These default rules allow members of group `admin` to have unlimited access and sets the default values to `deny`. It includes the users `tnsr` and `root` in the group `admin`.

Warning: TNSR Does not prevent a user from changing the rules in a way that would cut off all access.

- Changed password management to allow changing passwords for users in the host OS as well as for TNSR users [1091]

BGP

- Added explicit sequence numbering to BGP AS Path statements to support multiple patterns in a single AS Path [898]
- Added `show bgp network A.B.C.D` command to display detailed information about BGP routes [922]

CLI

- Added `enable` and `disable` commands to be used in favor of `no shutdown/shutdown` [938]
- Fixed CLI issues with data encoding that could lead to XML Parsing errors [887]

DHCP

- Improved support and control for DHCP server ([Kea](#)) management [490, 738, 1037, 1045]
- Added explicit `enable/disable` for DHCP Server daemon [1053]
- Added logging support to the DHCP Server [907]

DNS Resolver

- Added support for management of a DNS Resolver ([Unbound](#)) [492, 1072, 1093, 1094]

Hardware & Installation

- Added support for installation on Xeon D, C3000 SoCs [961]
- Added configuration packages for Netgate hardware that can run TNSR [1056]
- Fixed a Layer 2 connectivity issue with certain Intel 10G fiber configurations due to a timeout waiting for link [509]

IPsec

- Added QAT cryptographic acceleration enabled for IPsec [912, 940]
This acceleration works with QAT CPIC cards as well as C62X, C3XXX, and D15XX QAT devices.
- Fixed an issue where an IPsec Child SA would disappear after an IKEv1 Security Association re-authenticates [628]

NAT

- Fixed creating a NAT pool for custom route tables in the CLI [1055]
- Fixed handling of the NAT reassembly timeout value [1000]
- Added support for `output` feature NAT [867, 897]
- Fixed an error when changing static NAT command boolean properties [703]
- Addressed NAT issues which prevent the TNSR host OS network services from working on `nat outside` interfaces [616]

This can only work in `endpoint-dependent` NAT mode, which can be enabled as follows:

```
dataplane nat endpoint-dependent
service dataplane restart
```

This may become the default NAT mode in future TNSR releases [1079]

NTP

- Added support for NTP server (ntp.org) management [847, 939, 948, 952]

PKI (Certificates)

- Added support to the PKI CLI for managing certificate authority (CA) entries as well as certificate signing [930]

RESTCONF

- Added commands for **RESTCONF** management and authentication (HTTP server, [nginx](http://nginx.org)) [933]
- Added support to RESTCONF for certificate-based authentication [937]
When using certificates to authenticate, the common name (CN) part of the subject is used as the username.
- Added PAM support for HTTP authentication to the HTTP server [934]

32.10.2 Known Limitations

Authentication & Access Control

- Unable to delete a user from the CLI after TNSR services restart [1067]

BGP

- TNSR does not send BGP updates without restarting service with `redistribute from connected` option [746]
- Route with `aggregate-address` via `next-hop 0.0.0.0` does not appear in TNSR route table [832]
- BGP sessions may fail to establish or rapidly reconnect when receiving more prefixes than defined by `maximum-prefix limit` [858]
- The `maximum-prefix restart` command does not work [859]
- TNSR installs multiple paths for received routes even though support for multiple paths is not enabled [885]
- Unable to restart BGP service more than three times in a row [902]
Workaround: Run `systemctl reset-failed frr` from the shell to clear the error which will allow the BGP service to start again.
- Changing `update-source` from an IP address to `loop1` allows a session to establish but remote prefixes do not appear in the FIB until reboot [1104]

Bridge

- TNSR CLI allows multiple bridge interfaces to have `bvi` set [984]
Only the first interface set with `bvi` will work properly.
Workaround: Only set `bvi` on a single interface.

CLI

- Applied `dataplane` commands are not immediately present in the running configuration database until another change is made [1099]
- The candidate configuration database cannot be emptied with the `clear` command [1066]
- `show route table` causes the backend to die with large numbers of routes in the table [506]

For example, this crash happens with a full BGP feed.

RESTCONF

- `nginx` does not behave as expected with `authentication type none` [1086]

This mode is primarily for testing and not production use.

Workaround: Use password or certificate-based authentication for RESTCONF.

Interfaces

- Interface link speed displayed incorrectly in CLI and RESTCONF [672]
- Loopback interface responds to ICMP echo from an outside host even when in a *Down* state [850]

NAT

- Unable to create a `twice-nat` pool [972] or `twice-nat` not working [1023]

`twice-nat` can only work in endpoint-dependent NAT mode, which can be enabled as follows:

```
dataplane nat endpoint-dependent
service dataplane restart
```

- Unable to create `out-to-in-only` static mapping [976]

`out-to-in-only` can only work in endpoint-dependent NAT mode, which can be enabled as follows:

```
dataplane nat endpoint-dependent
service dataplane restart
```

- NAT Reassembly is not working for ICMP packets [990]
- Fragment limitation for NAT reassembly is not working [1065]
- NAT mode is not deleted from VPP startup configuration after TNSR services restart [1017]
- NAT forwarding is not working for `in2out` direction [1039]
- NAT static mappings are not added as expected when only the `port-local` value differs [1100]
- NAT static mapping with defined ports leads to `clixon-backend` crash after restart [1103]

VLAN/Subinterfaces

- Daemons such as Kea and `ntpd` do not correctly form configuration file references to subinterface names [1150]
- A VPP issue is preventing clients on subinterface networks from receiving return traffic that passes through TNSR [1152]

- These clients can communicate to TNSR, but not to hosts on other interfaces or subinterfaces.
- Other interface types work properly

32.10.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

32.11 TNSR 18.05 Release Notes

32.11.1 About This Release

This is the first public release of the TNSR product.

32.11.2 Known Limitations

- Loopback with IPv6 address will not respond to IPv6 pings [295]
 - **Workaround:** none.
- Linux route rules for the router-plugin/tap-inject are not cleaned up [477]

If the dataplane crashes, route rules added to the host system network stack are not cleaned up when it restarts.

 - **Workaround:** none.
- Deleting in-use prefix-list fails [483]

If you attempt to delete an in-use prefix list, the command will fail, but the configuration is left in an inconsistent state.

 - **Workaround:** remove the use of the prefix list prior to deleting it.
- DHCP Server Issues [490][739]

There are multiple issues with the DHCP Server, it's use is not recommended at this time.

 - **Workaround:** none.
- The command “show route table” causes backend crash [506]

A large route table (> 50k routes) can cause the “show route table” command to crash the backend process.

 - **Workaround:** Use “vppctl show ip fib” from a shell or vtysh to view route tables when a large number of routes have been added.
- RPC error when input includes “<” character [612]

Using the “<” character as input to the CLI can cause an RPC error. The error is properly detected, reported, and handled in the known cases. This affects all cases where there is free-form input.

 - **Workaround:** Do not use the “<” character.
- Enabling NAT on an outside interface disables services on that interface [616]

If you configure NAT on an outside interface, then that interface cannot provide services (like DHCP, ssh, etc.).

- **Workaround:** none
- SLAAC is not supported in dataplane, but host stack interfaces have it enabled [618]
 - **Workaround:** none.
- Child SAs can disappear after an IKEv1 SA reauth [628]
 - **Workaround:** none.
- Interface speed and duplex show as unknown [672]

The link speed and duplex indicators (visible with the “show interface” command) can display as “unknown”.

 - **Workaround:** Use the “vppctl show interface” command from an OS shell.
- Unable to change DHCP client hostname option [706]

The DHCP Client hostname can not be changed.

 - **Workaround:** none.
- Data plane restart breaks RESTCONF [741]

If you restart the data plane, the RESTCONF service loses its connection and does not reestablish it.

 - **Workaround:** Restart the data plane via the CLI, which does not have the same issue.
- RESTCONF RPC output is invalid JSON [745]

Some RPCs return multiple line output and the new line characters are not handled properly resulting in the inability of a JSON parser to process the output.

 - **Workaround:** none.
- BGP updates not being sent when “redistribute from connected” option specified [746]

Routes from connected routers are not propagated when the redistribute from connected option is set.

 - **Workaround:** none. You can temporarily resolve the problem by resetting the BGP service.
- BGP import-check feature does not work [781]

If the import-check option is set and then BGP is configured to advertise an unreachable network then the network is still advertised.

 - **Workaround:** none.
- Unable to create a default route when more than one loopback interface exists [824]
 - **Workaround:** none.
- Unable to create a second static NAT translation on a loopback interface [831]
 - **Workaround:** none.
- Route with aggregate-address via next-hop 0.0.0.0 doesn't appear in routing table [832]
 - **Workaround:** none.
- Loopback interface can be ping from an outside host even when marked down [850]
 - **Workaround:** none.
- BGP session constantly flapping when receiving more prefixes than defined in maximum-prefix limit command [858]
 - **Workaround:** none.
- BGP maximum-prefix restart option doesn't work [859]

- **Workaround:** none.
- No warning message in CLI when BGP `maximum-prefix` option is configured [860]
If the maximum number of prefixes is exceeded, there is no indication to a user that this has occurred.
 - **Workaround:** none.
- Unable to set BGP warning-only option for maximum-prefix option [861]
 - **Workaround:** none.

32.11.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

32.12 TNSR 0.1.0 Release Notes

32.12.1 About This Release

The TNSR 0.1.0 Release is the first release of the Netgate TNSR product. As there is no previous release of the TNSR products, there can be no changes relative to a previous version. Everything is new!

This release constitutes an early, evaluation version of the product.

32.12.2 Known Limitations

BGP Routes

While BGP may be configured, started, and run, reports of it not recording and displaying the learned BGP routes using the TNSR command “show routes” have been reported.

A possible work-around appears to be to stop, and then restart the BGP daemon using:

```
tnsr# service bgp stop
tnsr# service bgp start
```

BGP route-map and prefix-list Entries

TNSR route-maps and prefix-lists may be configured, and subsequently passed along to the underlying FRR configuration. TNSR will also allow removal of route-maps or prefix-lists from its configuration. However, they are not removed from the underlying FRR configuration.

A possible work-around is to manually remove them from the underlying FRR configuration using `vttysh` directly.

DHCP Server

The DHCP server does not support any form of Options yet.

The “server dhcp stop dhcp4” will not effectively terminate the Kea IPv4 DHCP server. A work-around is to run some form of “sudo killall kea-dhcp4” from a shell prompt.

32.12.3 Reporting Issues

For issues, please contact the Netgate Support staff.

- Send email to support@netgate.com
- Phone: 512.646.4100 (Support is Option 2)

LICENSING

The Netgate TNSR product uses a combination of Open Source and proprietary software subject to several different licenses.

The following list shows each Open Source component along with its license.

Table 1: Table of Open Source Licenses Used

Software	License
CentOS 7	<i>CentOS EULA</i>
Linux kernel and modules	<i>GPLv2</i>
cligen	<i>Apache 2.0</i>
clixon	<i>Apache 2.0</i>
curl	<i>MIT</i>
davici	<i>LGPLv2.1</i>
frr	<i>GPLv2</i>
kea	<i>MPL 2.0</i>
libnl	<i>LGPLv2.1</i>
net-snmp	<i>Net SNMP</i>
nginx	<i>BSD 2-clause</i>
ntp	<i>NTP License</i>
openssl	<i>OpenSSL/SSLeay</i>
strongswan	<i>GPLv2</i>
unbound	<i>BSD 3-clause</i>
VPP	<i>Apache 2.0</i>

GPL-licensed code modified for use in TNSR is available in source form:

Table 2: Table of Modified Open Source Repositories

Package	Repository Location
frr	http://github.com/netgate/frr
strongswan	http://github.com/netgate/strongswan
Hyper-V Linux kernel modules	https://github.com/netgate/uio_hv_generic

33.1 Apache 2.0 License

A copy of the Apache 2.0 License is found at <https://www.apache.org/licenses>.

The full text of the Apache 2.0 license is included below.

Apache License
Version 2.0, January 2004
<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the

(continues on next page)

(continued from previous page)

Licensors for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
 - (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
 - (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and

(continues on next page)

(continued from previous page)

wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify,

(continues on next page)

(continued from previous page)

defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

33.2 BSD 2-Clause License

A copy of the BSD 2-Clause License template is found at <https://opensource.org/licenses/BSD-2-Clause>.

The full text of the license as used by nginx is included below.

```
/*
 * Copyright (C) 2002-2019 Igor Sysoev
 * Copyright (C) 2011-2019 Nginx, Inc.
 * All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 * 1. Redistributions of source code must retain the above copyright
 *    notice, this list of conditions and the following disclaimer.
 * 2. Redistributions in binary form must reproduce the above copyright
 *    notice, this list of conditions and the following disclaimer in the
 *    documentation and/or other materials provided with the distribution.
 *
 * THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS ``AS IS'' AND
 * ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
```

(continues on next page)

(continued from previous page)

```
* ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
* SUCH DAMAGE.
*/
```

33.3 BSD 3-Clause License

A copy of the BSD 3-Clause License template is found at <https://opensource.org/licenses/BSD-3-Clause>.

The full text of the license as used by unbound is included below.

```
Copyright (c) 2007, NLnet Labs. All rights reserved.

This software is open source.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions
are met:

Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice,
this list of conditions and the following disclaimer in the documentation
and/or other materials provided with the distribution.

Neither the name of the NLNET LABS nor the names of its contributors may
be used to endorse or promote products derived from this software without
specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT
HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED
TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR
PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

33.4 CentOS EULA License

The TNSR project is built on a foundation of CentOS which is governed by the CentOS EULA License. It is found at http://mirror.centos.org/centos/7/os/x86_64/EULA.

Its full text is included below.

CentOS Linux 7 EULA

CentOS Linux 7 comes with no guarantees or warranties of any sorts, either written or implied.

The Distribution is released as GPLv2. Individual packages in the distribution come with their own licences.

33.5 GPLv2.0 License

A copy of the GPLv2 License is found at <https://www.gnu.org/licenses/old-licenses/gpl-2.0.txt>.

The full text of the GPLv2 license is included below.

GNU GENERAL PUBLIC LICENSE
Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

(continues on next page)

(continued from previous page)

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in

(continues on next page)

(continued from previous page)

whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for

(continues on next page)

(continued from previous page)

making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other

(continues on next page)

(continued from previous page)

circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

(continues on next page)

(continued from previous page)

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>
Copyright (C) <year> <name of author>
```

```
This program is free software; you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation; either version 2 of the License, or
(at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License along
with this program; if not, write to the Free Software Foundation, Inc.,
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if

(continues on next page)

(continued from previous page)

necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program
'Gnomovision' (which makes passes at compilers) written by James Hacker.

<signature of Ty Coon>, 1 April 1989
Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License.

33.6 LGPLv2.1 License

A copy of the LGPLv2.1 License is found at <https://www.gnu.org/licenses/lgpl-2.1.txt>.

The full text of the LGPLv2.1 license is included below.

GNU LESSER GENERAL PUBLIC LICENSE Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts
as the successor of the GNU Library Public License, version 2, hence
the version number 2.1.]

Preamble

The licenses for most software are designed to take away your
freedom to share and change it. By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it. You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use,
not price. Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

(continues on next page)

(continued from previous page)

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be

(continues on next page)

(continued from previous page)

allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

(continues on next page)

(continued from previous page)

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

(continues on next page)

(continued from previous page)

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative

(continues on next page)

(continued from previous page)

work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for

(continues on next page)

(continued from previous page)

reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or

(continues on next page)

(continued from previous page)

otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

(continues on next page)

(continued from previous page)

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the library's name and a brief idea of what it does.>
Copyright (C) <year> <name of author>
```

```
This library is free software; you can redistribute it and/or
modify it under the terms of the GNU Lesser General Public
License as published by the Free Software Foundation; either
version 2.1 of the License, or (at your option) any later version.
```

```
This library is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU
Lesser General Public License for more details.
```

```
You should have received a copy of the GNU Lesser General Public
License along with this library; if not, write to the Free Software
Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
```

Also add information on how to contact you by electronic and paper mail.

(continues on next page)

(continued from previous page)

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the
library 'Frob' (a library for tweaking knobs) written by James Random Hacker.
```

```
<signature of Ty Coon>, 1 April 1990
Ty Coon, President of Vice
```

That's all there is to it!

33.7 MIT License

A copy of the MIT License template is found at <https://opensource.org/licenses/MIT>.

The full text of the license as used by CURL is included below.

```
Copyright (c) 1996 - 2018, Daniel Stenberg, <daniel@haxx.se>, and many
contributors, see the THANKS file.
```

```
All rights reserved.
```

```
Permission to use, copy, modify, and distribute this software for any
purpose with or without fee is hereby granted, provided that the above
copyright notice and this permission notice appear in all copies.
```

```
THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT
OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT
HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER
IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR
IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE
SOFTWARE.
```

```
Except as contained in this notice, the name of a copyright holder
shall not be used in advertising or otherwise to promote the sale, use
or other dealings in this Software without prior written authorization
of the copyright holder.
```

33.8 Mozilla Public License 2.0 (MPL-2.0)

A copy of the Mozilla Public License 2.0 (MPL-2.0) template is found at <https://opensource.org/licenses/MPL-2.0>.

```
Mozilla Public License Version 2.0
=====
```

```
1. Definitions
-----
```

(continues on next page)

(continued from previous page)

- 1.1. "Contributor"
means each individual or legal entity that creates, contributes to the creation of, or owns Covered Software.
- 1.2. "Contributor Version"
means the combination of the Contributions of others (if any) used by a Contributor and that particular Contributor's Contribution.
- 1.3. "Contribution"
means Covered Software of a particular Contributor.
- 1.4. "Covered Software"
means Source Code Form to which the initial Contributor has attached the notice in Exhibit A, the Executable Form of such Source Code Form, and Modifications of such Source Code Form, in each case including portions thereof.
- 1.5. "Incompatible With Secondary Licenses"
means
- (a) that the initial Contributor has attached the notice described in Exhibit B to the Covered Software; or
 - (b) that the Covered Software was made available under the terms of version 1.1 or earlier of the License, but not also under the terms of a Secondary License.
- 1.6. "Executable Form"
means any form of the work other than Source Code Form.
- 1.7. "Larger Work"
means a work that combines Covered Software with other material, in a separate file or files, that is not Covered Software.
- 1.8. "License"
means this document.
- 1.9. "Licensable"
means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently, any and all of the rights conveyed by this License.
- 1.10. "Modifications"
means any of the following:
- (a) any file in Source Code Form that results from an addition to, deletion from, or modification of the contents of Covered Software; or
 - (b) any new file in Source Code Form that contains any Covered Software.
- 1.11. "Patent Claims" of a Contributor
means any patent claim(s), including without limitation, method, process, and apparatus claims, in any patent Licensable by such Contributor that would be infringed, but for the grant of the License, by the making, using, selling, offering for sale, having

(continues on next page)

(continued from previous page)

made, import, or transfer of either its Contributions or its Contributor Version.

1.12. "Secondary License"

means either the GNU General Public License, Version 2.0, the GNU Lesser General Public License, Version 2.1, the GNU Affero General Public License, Version 3.0, or any later versions of those licenses.

1.13. "Source Code Form"

means the form of the work preferred for making modifications.

1.14. "You" (or "Your")

means an individual or a legal entity exercising rights under this License. For legal entities, "You" includes any entity that controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. License Grants and Conditions

2.1. Grants

Each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license:

- (a) under intellectual property rights (other than patent or trademark) Licensable by such Contributor to use, reproduce, make available, modify, display, perform, distribute, and otherwise exploit its Contributions, either on an unmodified basis, with Modifications, or as part of a Larger Work; and
- (b) under Patent Claims of such Contributor to make, use, sell, offer for sale, have made, import, and otherwise transfer either its Contributions or its Contributor Version.

2.2. Effective Date

The licenses granted in Section 2.1 with respect to any Contribution become effective for each Contribution on the date the Contributor first distributes such Contribution.

2.3. Limitations on Grant Scope

The licenses granted in this Section 2 are the only rights granted under this License. No additional rights or licenses will be implied from the distribution or licensing of Covered Software under this License. Notwithstanding Section 2.1(b) above, no patent license is granted by a Contributor:

- (a) for any code that a Contributor has removed from Covered Software;
or

(continues on next page)

(continued from previous page)

- (b) for infringements caused by: (i) Your and any other third party's modifications of Covered Software, or (ii) the combination of its Contributions with other software (except as part of its Contributor Version); or
- (c) under Patent Claims infringed by Covered Software in the absence of its Contributions.

This License does not grant any rights in the trademarks, service marks, or logos of any Contributor (except as may be necessary to comply with the notice requirements in Section 3.4).

2.4. Subsequent Licenses

No Contributor makes additional grants as a result of Your choice to distribute the Covered Software under a subsequent version of this License (see Section 10.2) or under the terms of a Secondary License (if permitted under the terms of Section 3.3).

2.5. Representation

Each Contributor represents that the Contributor believes its Contributions are its original creation(s) or it has sufficient rights to grant the rights to its Contributions conveyed by this License.

2.6. Fair Use

This License is not intended to limit any rights You have under applicable copyright doctrines of fair use, fair dealing, or other equivalents.

2.7. Conditions

Sections 3.1, 3.2, 3.3, and 3.4 are conditions of the licenses granted in Section 2.1.

3. Responsibilities

3.1. Distribution of Source Form

All distribution of Covered Software in Source Code Form, including any Modifications that You create or to which You contribute, must be under the terms of this License. You must inform recipients that the Source Code Form of the Covered Software is governed by the terms of this License, and how they can obtain a copy of this License. You may not attempt to alter or restrict the recipients' rights in the Source Code Form.

3.2. Distribution of Executable Form

If You distribute Covered Software in Executable Form then:

- (a) such Covered Software must also be made available in Source Code Form, as described in Section 3.1, and You must inform recipients of the Executable Form how they can obtain a copy of such Source Code Form by reasonable means in a timely manner, at a charge no more

(continues on next page)

(continued from previous page)

than the cost of distribution to the recipient; and

- (b) You may distribute such Executable Form under the terms of this License, or sublicense it under different terms, provided that the license for the Executable Form does not attempt to limit or alter the recipients' rights in the Source Code Form under this License.

3.3. Distribution of a Larger Work

You may create and distribute a Larger Work under terms of Your choice, provided that You also comply with the requirements of this License for the Covered Software. If the Larger Work is a combination of Covered Software with a work governed by one or more Secondary Licenses, and the Covered Software is not Incompatible With Secondary Licenses, this License permits You to additionally distribute such Covered Software under the terms of such Secondary License(s), so that the recipient of the Larger Work may, at their option, further distribute the Covered Software under the terms of either this License or such Secondary License(s).

3.4. Notices

You may not remove or alter the substance of any license notices (including copyright notices, patent notices, disclaimers of warranty, or limitations of liability) contained within the Source Code Form of the Covered Software, except that You may alter any license notices to the extent required to remedy known factual inaccuracies.

3.5. Application of Additional Terms

You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, You may do so only on Your own behalf, and not on behalf of any Contributor. You must make it absolutely clear that any such warranty, support, indemnity, or liability obligation is offered by You alone, and You hereby agree to indemnify every Contributor for any liability incurred by such Contributor as a result of warranty, support, indemnity or liability terms You offer. You may include additional disclaimers of warranty and limitations of liability specific to any jurisdiction.

4. Inability to Comply Due to Statute or Regulation

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Software due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be placed in a text file included with all distributions of the Covered Software under this License. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

5. Termination

(continues on next page)

(continued from previous page)

5.1. The rights granted under this License will terminate automatically if You fail to comply with any of its terms. However, if You become compliant, then the rights granted under this License from a particular Contributor are reinstated (a) provisionally, unless and until such Contributor explicitly and finally terminates Your grants, and (b) on an ongoing basis, if such Contributor fails to notify You of the non-compliance by some reasonable means prior to 60 days after You have come back into compliance. Moreover, Your grants from a particular Contributor are reinstated on an ongoing basis if such Contributor notifies You of the non-compliance by some reasonable means, this is the first time You have received notice of non-compliance with this License from such Contributor, and You become compliant prior to 30 days after Your receipt of the notice.

5.2. If You initiate litigation against any entity by asserting a patent infringement claim (excluding declaratory judgment actions, counter-claims, and cross-claims) alleging that a Contributor Version directly or indirectly infringes any patent, then the rights granted to You by any and all Contributors for the Covered Software under Section 2.1 of this License shall terminate.

5.3. In the event of termination under Sections 5.1 or 5.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or Your distributors under this License prior to termination shall survive termination.

```
*****
*
* 6. Disclaimer of Warranty
* -----
*
* Covered Software is provided under this License on an "as is"
* basis, without warranty of any kind, either expressed, implied, or
* statutory, including, without limitation, warranties that the
* Covered Software is free of defects, merchantable, fit for a
* particular purpose or non-infringing. The entire risk as to the
* quality and performance of the Covered Software is with You.
* Should any Covered Software prove defective in any respect, You
* (not any Contributor) assume the cost of any necessary servicing,
* repair, or correction. This disclaimer of warranty constitutes an
* essential part of this License. No use of any Covered Software is
* authorized under this License except under this disclaimer.
*
*****
```

```
*****
*
* 7. Limitation of Liability
* -----
*
* Under no circumstances and under no legal theory, whether tort
* (including negligence), contract, or otherwise, shall any
* Contributor, or anyone who distributes Covered Software as
* permitted above, be liable to You for any direct, indirect,
* special, incidental, or consequential damages of any character
* including, without limitation, damages for lost profits, loss of
* goodwill, work stoppage, computer failure or malfunction, or any
*
*****
```

(continues on next page)

(continued from previous page)

* and all other commercial damages or losses, even if such party *
 * shall have been informed of the possibility of such damages. This *
 * limitation of liability shall not apply to liability for death or *
 * personal injury resulting from such party's negligence to the *
 * extent applicable law prohibits such limitation. Some *
 * jurisdictions do not allow the exclusion or limitation of *
 * incidental or consequential damages, so this exclusion and *
 * limitation may not apply to You. *
 * *

8. Litigation

Any litigation relating to this License may be brought only in the courts of a jurisdiction where the defendant maintains its principal place of business and such litigation shall be governed by laws of that jurisdiction, without reference to its conflict-of-law provisions. Nothing in this Section shall prevent a party's ability to bring cross-claims or counter-claims.

9. Miscellaneous

This License represents the complete agreement concerning the subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not be used to construe this License against a Contributor.

10. Versions of the License

10.1. New Versions

Mozilla Foundation is the license steward. Except as provided in Section 10.3, no one other than the license steward has the right to modify or publish new versions of this License. Each version will be given a distinguishing version number.

10.2. Effect of New Versions

You may distribute the Covered Software under the terms of the version of the License under which You originally received the Covered Software, or under the terms of any subsequent version published by the license steward.

10.3. Modified Versions

If you create software not governed by this License, and you want to create a new license for such software, you may create and use a modified version of this License if you rename the license and remove any references to the name of the license steward (except to note that such modified license differs from this License).

10.4. Distributing Source Code Form that is Incompatible With Secondary

(continues on next page)

(continued from previous page)

Licenses

If You choose to distribute Source Code Form that is Incompatible With Secondary Licenses under the terms of this version of the License, the notice described in Exhibit B of this License must be attached.

Exhibit A - Source Code Form License Notice

This Source Code Form is subject to the terms of the Mozilla Public License, v. 2.0. If a copy of the MPL was not distributed with this file, You can obtain one at <http://mozilla.org/MPL/2.0/>.

If it is not possible or desirable to put the notice in a particular file, then You may include the notice in a location (such as a LICENSE file in a relevant directory) where a recipient would be likely to look for such a notice.

You may add additional accurate notices of copyright ownership.

Exhibit B - "Incompatible With Secondary Licenses" Notice

This Source Code Form is "Incompatible With Secondary Licenses", as defined by the Mozilla Public License, v. 2.0.

33.9 Net SNMP License

The net-snmp repository is used governed by several licenses collectively listed as the Net SNMP License. It is found at <http://www.net-snmp.org/about/license.html>.

Its full text is included below.

Various copyrights apply to this package, listed in various separate parts below. Please make sure that you read all the parts.

---- Part 1: CMU/UCD copyright notice: (BSD like) ----

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

(continues on next page)

(continued from previous page)

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---- Part 2: Networks Associates Technology, Inc copyright notice (BSD) ----

Copyright (c) 2001-2003, Networks Associates Technology, Inc
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 3: Cambridge Broadband Ltd. copyright notice (BSD) ----

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

(continues on next page)

(continued from previous page)

- * The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

----- Part 4: Sun Microsystems, Inc. copyright notice (BSD) -----

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Use is subject to license terms below.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

(continues on next page)

(continued from previous page)

----- Part 5: Sparta, Inc copyright notice (BSD) -----

Copyright (c) 2003-2009, Sparta, Inc
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright
notice, this list of conditions and the following disclaimer in the
documentation and/or other materials provided with the distribution.
- * Neither the name of Sparta, Inc nor the names of its contributors may
be used to endorse or promote products derived from this software
without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS
IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR
CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,
EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS;
OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR
OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF
ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

----- Part 6: Cisco/BUPTNIC copyright notice (BSD) -----

Copyright (c) 2004, Cisco, Inc and Information Network
Center of Beijing University of Posts and Telecommunications.
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright
notice, this list of conditions and the following disclaimer in the
documentation and/or other materials provided with the distribution.
- * Neither the name of Cisco, Inc, Beijing University of Posts and
Telecommunications, nor the names of their contributors may
be used to endorse or promote products derived from this software
without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS
IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR
CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,

(continues on next page)

(continued from previous page)

EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 7: Fabasoft R&D Software GmbH & Co KG copyright notice (BSD) ----

Copyright (c) Fabasoft R&D Software GmbH & Co KG, 2003
oss@fabasoft.com
Author: Bernhard Penz

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * The name of Fabasoft R&D Software GmbH & Co KG or any of its subsidiaries, brand or product names may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 8: Apple Inc. copyright notice (BSD) ----

Copyright (c) 2007 Apple Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of Apple Inc. ("Apple") nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

(continues on next page)

(continued from previous page)

```
THIS SOFTWARE IS PROVIDED BY APPLE AND ITS CONTRIBUTORS "AS IS" AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL APPLE OR ITS
CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF
USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND
ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT
OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
SUCH DAMAGE.
```

```
---- Part 9: ScienceLogic, LLC copyright notice (BSD) ----
```

```
Copyright (c) 2009, ScienceLogic, LLC
All rights reserved.
```

```
Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are
met:
```

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of ScienceLogic, LLC nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

```
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
`AS IS' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT
HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,
BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS
OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND
ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR
TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE
USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH
DAMAGE.
```

33.10 NTP License

The ntp.org NTP daemon uses a custom open source license known as the NTP License.

A copy of the NTP License template is found at <https://opensource.org/licenses/NTP>.

The full text of the license as used by NTP is included below.

The following copyright notice applies to all files collectively called the Network Time Protocol Version 4 Distribution. Unless specifically declared otherwise in an individual file, this entire notice applies as if the text was explicitly included in the file.

```
*****
*
* Copyright (c) University of Delaware 1992-2015
*
* Permission to use, copy, modify, and distribute this software and
* its documentation for any purpose with or without fee is hereby
* granted, provided that the above copyright notice appears in all
* copies and that both the copyright notice and this permission
* notice appear in supporting documentation, and that the name
* University of Delaware not be used in advertising or publicity
* pertaining to distribution of the software without specific,
* written prior permission. The University of Delaware makes no
* representations about the suitability this software for any
* purpose. It is provided "as is" without express or implied
* warranty.
*
*****
```

Content starting in 2011 from Harlan Stenn, Danny Mayer, and Martin Burnicki is:

```
*****
*
* Copyright (c) Network Time Foundation 2011-2015
*
* All Rights Reserved
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
*
* 1. Redistributions of source code must retain the above copyright
*    notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above
*    copyright notice, this list of conditions and the following
*    disclaimer in the documentation and/or other materials provided
*    with the distribution.
*
*
* THIS SOFTWARE IS PROVIDED BY THE AUTHORS ``AS IS'' AND ANY EXPRESS
* OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
* WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
* ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE
* LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
* CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT
* OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR
* BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
* LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
* (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE
* USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH
* DAMAGE.
*****
```

33.11 Joint OpenSSL and SSLeay License

OpenSSL 1.1.1, 1.1.0, 1.0.2, and all prior releases are covered by the joint OpenSSL and SSLeay License as noted at <https://www.openssl.org/source/license.html>

The full text of the joint OpenSSL and SSLeay license is included below, and is also available at <https://www.openssl.org/source/license-openssl-ssleay.txt>

LICENSE ISSUES
=====

The OpenSSL toolkit stays under a double license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. See below for the actual license texts.

OpenSSL License

```
/* =====
 * Copyright (c) 1998-2018 The OpenSSL Project. All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 * 1. Redistributions of source code must retain the above copyright
 *    notice, this list of conditions and the following disclaimer.
 *
 * 2. Redistributions in binary form must reproduce the above copyright
 *    notice, this list of conditions and the following disclaimer in
 *    the documentation and/or other materials provided with the
 *    distribution.
 *
 * 3. All advertising materials mentioning features or use of this
 *    software must display the following acknowledgment:
 *    "This product includes software developed by the OpenSSL Project
 *    for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
 *
 * 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
 *    endorse or promote products derived from this software without
 *    prior written permission. For written permission, please contact
 *    openssl-core@openssl.org.
 *
 * 5. Products derived from this software may not be called "OpenSSL"
 *    nor may "OpenSSL" appear in their names without prior written
 *    permission of the OpenSSL Project.
 *
 * 6. Redistributions of any form whatsoever must retain the following
 *    acknowledgment:
 *    "This product includes software developed by the OpenSSL Project
 *    for use in the OpenSSL Toolkit (http://www.openssl.org/)"
 *
 * THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY
 * EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
 * PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR
 * ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
```

(continues on next page)

(continued from previous page)

```

* SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
* LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
* STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
* ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
* OF THE POSSIBILITY OF SUCH DAMAGE.
* =====
*
* This product includes cryptographic software written by Eric Young
* (eay@cryptsoft.com). This product includes software written by Tim
* Hudson (tjh@cryptsoft.com).
*
*/

```

Original SSLeay License

```

/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
* All rights reserved.
*
* This package is an SSL implementation written
* by Eric Young (eay@cryptsoft.com).
* The implementation was written so as to conform with Netscapes SSL.
*
* This library is free for commercial and non-commercial use as long as
* the following conditions are aheared to. The following conditions
* apply to all code found in this distribution, be it the RC4, RSA,
* lhash, DES, etc., code; not just the SSL code. The SSL documentation
* included with this distribution is covered by the same copyright terms
* except that the holder is Tim Hudson (tjh@cryptsoft.com).
*
* Copyright remains Eric Young's, and as such any Copyright notices in
* the code are not to be removed.
* If this package is used in a product, Eric Young should be given attribution
* as the author of the parts of the library used.
* This can be in the form of a textual message at program startup or
* in documentation (online or textual) provided with the package.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the copyright
* notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in the
* documentation and/or other materials provided with the distribution.
* 3. All advertising materials mentioning features or use of this software
* must display the following acknowledgement:
* "This product includes cryptographic software written by
* Eric Young (eay@cryptsoft.com)"
* The word 'cryptographic' can be left out if the rouines from the library
* being used are not cryptographic related :-).
* 4. If you include any Windows specific code (or a derivative thereof) from
* the apps directory (application code) you must include an acknowledgement:
* "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"
*

```

(continues on next page)

(continued from previous page)

```
* THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND
* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
* ARE DISCLAIMED.  IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
* SUCH DAMAGE.
*
* The licence and distribution terms for any publically available version or
* derivative of this code cannot be changed.  i.e. this code cannot simply be
* copied and put under another distribution licence
* [including the GNU Public Licence.]
*/
```