



Static IP Routing (SIR)

```
SIR>endpoint statistics 1
Routing Endpoint L3EPL statistics:
Rx Packets:          140715  Tx Packets:          920
Rx Octets:           14277950  Tx Octets:           256822
Rx Unicast:          1539  Tx Unicast:          906
Rx Multicast:        7516  Tx Multicast:        11
Rx Broadcast:        131660  Tx Broadcast:        3
Rx Fomw Datagrams:   1  Tx Fomw D
Tx Reque
Rx Discards:         0  Tx Discar
Rx No Routes:        0  Tx No Rou
Rx Unknown Protos:
Rx Addr Errors:
Rx Header Errors:
Rx ICMP Packets:     54  Tx ICMP P
Rx ICMP Discards:    0
```

The screenshot shows the configuration page for a 33280 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with categories like Configuration, System, External Protection, and Security. The main content area is titled 'L3 Routing Table Configuration' and features a table with columns for 'Delete', 'Destination', 'Gateway', and 'Egress L3 Endpoint'. The table lists several routes, including a default route. Below the table are buttons for 'Add New Route', 'Add Default Gateway', 'Save', and 'Reset'.

Delete	Destination	Gateway	Egress L3 Endpoint
<input type="checkbox"/>	16.100.0.0/24		L3EPL 100 / default
<input type="checkbox"/>	16.100.1.0/24		L3EPL 101 / default
<input type="checkbox"/>	16.100.2.0/24	16.100.0.100	L3EPL 100 / default
<input type="checkbox"/>	Default	16.100.10.1	default
<input type="checkbox"/>	Default	16.100.16.100	uplink

User Guide

33542 Rev. A

Trademarks

All trademarks and registered trademarks are the property of their respective owners.

Copyright Notice/Restrictions

Copyright © 2012, 2013 Transition Networks

All rights reserved.

No part of this work may be reproduced or used in any form or by any means (graphic, electronic or mechanical) without written permission from Transition Networks.

The information contained herein is confidential property of Transition Networks, Inc. The use, copying, transfer or disclosure of such information is prohibited except by express written agreement with Transition Networks, Inc.

Printed in the U.S.A.

Static IP Routing (SIR) User Guide, 33542 Rev. A

Contact Information

Transition Networks

10900 Red Circle Drive

Minnetonka, MN 55343 USA

Tel: 952-941-7600 or 1-800-526-9267

Fax: 952-941-2322

Revision History

Rev	Date	Description
A	06/12/13	Initial release for SIR version v1.0 at S3280 SW v 1.5.

Contents

Trademarks	2
Copyright Notice/Restrictions	2
Contact Information	2
Revision History	2
Contents	3
1. Introduction	5
Supported Models	5
Static IP Specs	5
Related Manuals and Online Help.....	6
Static IP Routing Process.....	7
System Functions	8
Static IP Routing / CIDR Example.....	9
SIR Details.....	9
High-Level Functional View	10
Component Descriptions	11
The Routing Manager Component.....	11
The Fast Path Component.....	11
The ARP Component.....	12
The ICMP Component	12
Hardware Layer Interface.....	12
General Information	12
Modifications to the CPU Port.....	13
The Packet Module	13
2. CLI (Command Line Interface)	15
Changes to Existing CLI Commands	15
Routing Endpoint Commands	16
Routing Table Commands.....	22
ARP Table Commands.....	24
Network Diagnostic Commands	26
SIR Configuration Example (via CLI)	29
Background Operations	30
Detailed Configuration (CLI)	31
CLI Automation Scripts.....	38
Start Background Script.....	38
Script Files used for Background Operations	38
3. Web Interface	40
Configuration > Static Routing > Endpoints Configuration.....	41
L3 Static Routing Endpoints Configuration	41
Adding a New L3 Endpoint	42
Configuration > Static Routing > Routing Table.....	44
L3 Routing Table Configuration	44
Configuration > Static Routing > ARP Table.....	49
L3 Static Routing ARP Config.....	49
Flushing of ARP Table	49
Monitor > Static Routing > Statistics	51
L3 Static Routing Statistics	51
Monitor > Static Routing > ARP Table	53
L3 Static Routing ARP Table	53
Monitor > Static Routing > Alarms.....	55
Diagnostics > Static Routing Utilities.....	56
Diagnostics > Static Routing > Ping.....	56
Ping Output	57
Diagnostics > Static Routing > Trace.....	58

Send Trace Parameters58

Trace Output59

SIR Configuration Example (via Web GUI)60

4. Messages and Troubleshooting.....61

Static IP Routing Messages61

CLI Messages and Web GUI Messages.....61

For More Error Message Information63

Static IP Routing (SIR) Troubleshooting64

5. SNMP MIBs and Traps.....65

Supported MIBs.....65

Standard MIBs.....65

Proprietary MIBs.....65

SIP-MIB header.....66

IP-MIB extension.....66

FORWADR-MIB extension67

ARP-MIB extension.....67

For Additional MIB Information69

Appendix A: Service, Warranty & Compliance Information70

Service.....70

Warranty70

Compliance Information70

Cautions and Warnings70

Glossary of SIR Terms71

Figures

Figure 1: IP Routing Feature Architectural Positioning7

Figure 2. Static IP Routing Module Functional Diagram10

Figure 3. Software Data Path14

Figure 4. Test Setup29

1. Introduction

The Static IP Routing (SIR) feature provides routing capabilities for up to 16 L3 endpoints. The routing is based on statically configured routes. Up to 128 static routes can be stored in the routing tables. The SIR feature lets you define L3 routing endpoints over the VLANs configured in the NID. Through such endpoints, the traffic is routed between the VLANs, such as hosts from one VLAN can communicate through hosts located in different VLANs. The Static IP routing functionality also includes:

- ARP and proxy ARP capabilities
- Split horizon routing
- Hairpin routing
- ICMP (Internet Control Message protocol)
- Diagnostic tools: traceroute and ping

More details about the Static IP Routing functionality are provided below.

Supported Models

The Static IP Routing (SIR) feature is supported on Transition Networks model S3280-TST and S4140 NIDs. SIR operation is essentially identical between models, with minor differences noted where they exist.

Static IP Specs


- Static Routes
 - Maximum 129 routes in the routing tables:
 - Up to 8 subnets per VLAN.
 - Up to 16 VLANs supported (including “default” VLAN at port level).
 - Routing table organized by subnet:
 - No common subnets across VLAN (i.e. cannot have subnet a.b.c.d on both VLAN x and VLAN y).
 - Cannot have same subnet across two different VLANs.
 - Can route across VLANs.
 - IPv4 only.
 - Classless Inter-Domain Routing (CIDR).
- Support ICMP, with mandatory options only.
- No IP fragmentation.
 - All subnets must have the same MTU for routing (1500 bytes).
- Untagged and single tagged frames only
- Gratuitous ARP:
 - Process received packets and send on IP config change/startup.
- Proxy ARP – enable/disable by configuration.
- Tools:
 - Ping.
 - Traceroute, and
 - Initiator / reflector for both.
- L3 processing enabled on a <Port, VLAN> with an IP address configured. Otherwise, pure L2 forwarding.
- No support for IGMP or multicast IP addresses.
- Treat multicast frames based on L2 rules (MAC multicast).
- Routing endpoints:
 - <VLAN, Port> pairs,
 - One IP address/subnet may be assigned to more than one Port, within the same VLAN, and
 - One <VLAN, Port> may have more than one IP subnet (hairpin routing).
- Low speed traffic

Related Manuals and Online Help

This manual is one of several S3280 family manuals which can include:

- Product Documentation Postcard, 33504
- S3280 Install Guide, 33513
- S3280-TST Install Guide, 33520
- S3280 Web User Guide, 33506
- S3280 CLI User Guide, 33507
- Converge™ EMS Administrator's Procedures (33544) and Install Guides (33543, 33548)
- RFC2544 User Guide (33540)
- MRP User Guide (33541)
- Static IP Routing User Guide (33542) (this manual)
- Release Notes (version specific)



Context-sensitive Help screens are built into the S3280 Web interface (click ) and the CLI (type ? or **Help**).

Check the TN web site at <http://www.transition.com/> for additional white papers, application notes, etc.

Check the S3280 landing page at <http://www.transition.com/TransitionNetworks/Landing/indura/indura.aspx> for Product Information, Application Notes, etc.

Check the S3280 product page at <http://www.transition.com/TransitionNetworks/Products2/Family.aspx?Name=S3280> for access to the latest S3280 datasheet, features, applications, specs, SKUs, etc.

When the procedures in this manual are successfully completed, refer to the *S3280 Web Interface User Guide* or the *S3280 CLI Reference* for configuration, monitoring, diagnostics, and maintenance information.

Static IP Routing Process

This section describes the functional aspects of static IP routing as implemented on the S3280 NID.

The S3280 NID (without the Static IP Routing function) is a Layer 2 forwarding device. Frames are forwarded only based on VLAN IDs and Destination MAC addresses, according to the VLAN configuration. Without L3 routing, VLANs configured on the device are independent broadcast domains. Frames from one VLAN cannot be forwarded to a different VLAN by the S3280 devices.

The Static IP Routing feature lets you define L3 routing endpoints over the VLANs configured in the S3280. Through such endpoints, the traffic is routed between the VLANs (e.g., a host from one VLAN can communicate through hosts located in different VLANs).

So from a functional perspective, the Static IP Routing feature resides above the L2 forwarding as shown in Figure 1 below.




Figure 1: IP Routing Feature Architectural Positioning

By default, frames are forwarded based on the L2 forwarding rules. Once a L3 routing endpoint is configured for a given VLAN, routing capabilities are enabled for that particular VLAN. That means that IP frames with the destination MAC address of the L3 routing endpoint are routed, by the software, between the IP subnets defined in the NID, while all other frames are forwarded according to the L2 rules. More details are provided in the following sections. See IETF [RFC 2863](https://www.rfc-editor.org/rfc/2863) for additional information.

System Functions

The Static IP Routing functions can be performed from the CLI or web interface; see section 2. [CLI \(Command Line Interface\)](#) or section 3. [Web Interface](#) for details.

The IP forwarding function is based on assigning subnets on a per-port or per-VLAN (C-TAG) basis. Each subnet will have an IP address assigned to the port or the VLAN for routing purposes (default gateway, etc). Each subnet will also have a default forwarding IP address. A common industry method for these configurations and functions is the Classless Inter-Domain Routing (CIDR). This allows a simplified configuration method for the network operator to instantiate forwarding rules on the S3280.

There are up to 8 subnets (size defined by a mask) that can be assigned per VLAN, with up to 16 VLANS supported. If an interface is to be supported (default VLAN), then 8 subnets per interface are supported. The forwarding tables allow a default IP address (0.0.0.0) that allows forwarding to an interface, such as one connected to a service provider, or a default server. This minimizes the network topology knowledge needed at the device or by the configuration methods.

Address Resolution Protocol (ARP) is used for association of the IP address with the Ethernet Address associated with each device on the network. ARP caches must be cleared on a periodic basis.

Gratuitous ARP sensing is supported such that if a device advertises itself, the S3280 adds it to the ARP tables in the event that there is a need for forwarding a packet.

Proxy ARPing across interfaces is supported, and can be enabled or disabled per the operator's policy.

The CLI interface is the most commonly used customer interface. SIR is implemented in the S3280 family of supported NIDs via the CLI, scripting, GUI, and SNMP.

Static IP Routing / CIDR Example

Business Ethernet customers often require specific subnets unique to their internal networks. Operators providing E-LAN services for business customers provide the required IP addresses for routing and interconnecting various customer sites for both private and public Internet access.

Classless Inter-Domain Routing (CIDR) allows a network to be sub-divided into variously sized subnets, providing the opportunity to size a network based on the actual subnets that the customer requires instead of allocating a full IPv4 8-bit block of IP addresses (Class A,B,C).

The new CIDR feature allows individual customer subnets to be statically assigned. Any port should be capable of being provisioned with an IPv4 CIDR notation that specifies a syntax of IP addresses and their associated routing prefix (e.g., 192.168.1.0/24).

In the example below, the first interface (ViF1) is the connection to the Operator's aggregation node with the /30 IP address space specified. Everything is routed to the edge device that is on/associated to the **.73** network. The second interface (ViF2) hosts the IP address for the customer with their gateway address **.81**.

CIDR Example:

```
interface vlan vif1
tag 10
ip 10.100.1.74/30
ports 1
management telnet 0.0.0.0/0
management snmp 0.0.0.0/0
!
interface vlan vif2
tag 20
ip 64.22.224..81/29
ports 2
!
ip route 0.0.0.0/0 10.100.1..73
```

SIR Details

S3280 Static IP routing is implemented with these components:

- a) *Routing management component*, which contains data routing table and routing endpoints database. It performs route lookup, slow path forwarding, and egress port lookups from the MAC table.
- b) *ARP component*, which processes incoming ARP packets, responds to ARP requests, generates ARP requests and maintains the ARP cache table;
- c) *ICMP component*, which processes incoming ICMP requests, responses to ICMP echo packets and generate ICMP packets with corresponding operational codes. It is also responsible for generation of ICMP packets for ping and traceroute utilities. It also manages ping and traceroute sessions initiated by the user and keeps sessions data.
- d) *Fast path component*, which performs validation, filtering, forwarding and editing of the IP packets. Major task of the fast path is to forward valid IP packets basing on the routing cache table.

There is a common northbound management API that is suitable for all management facilities (CLI, SNMP, and WEB) used in the system.

Southbound Static IP routing components interact with the S3280 Eth driver to perform fast path routing and with the S3280 L2 MAC table for getting the egress port when building routing cache entry.

High-Level Functional View

From a functional point of view (Figure 2) there are two separate tasks performing routing in the system:

1. A fast path high priority task, responsible for validation of received packets, and their dispatching according to routing cache.
2. A low priority slow task, responsible for configuration of endpoints and static routing table; handling ARP, ICMP and misrouted packets. It is also responsible for the audit of routing cache.

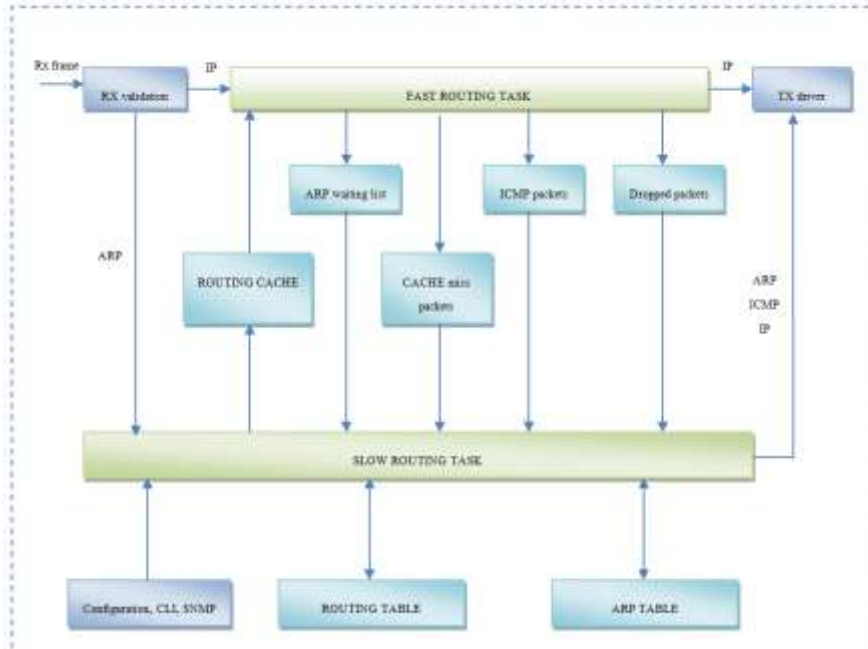


Figure 2. Static IP Routing Module Functional Diagram

In the figure above, the fast task is processing the RX frames as follows:

1. Error packets are dropped. MAC address is checked to match L3 routing pattern. If not match, the packet is returned to driver for processing by other applications (management, SOAM, RSTP etc). ARP packets are queued to the slow task.
2. IP packets are looked up in the routing cache and, if hit they are edited and sent to the egress port.
3. "Cache miss" IP packets are queued to the slow task for lookup in the slow path routing table and for necessary ARP/ICMP processing.
4. ICMP packets destined to L3 endpoint are queued to the slow task.
5. Dropped packets are queued to slow task.
6. Packets which are not subject for IP routing (inband management, SOAM or other control plane protocols with software implementation), are returned back to the Packet module of the S3280 framework without any changes.

The slow task is processing RX frames as follows:

- a) Handles cache miss packets, performs lookup in the routing table, adds lines to routing cache if lookup success.
- b) Handles ARP packets and builds the ARP table.
- c) Sends ICMP replies for dropped IP packets, replies to ICMP Echo packets and processes other incoming ICMP packets.
- d) Handles ARP waiting packets queue, sends pending packets when ARP requests are resolved and drops pending packets if timeout expired.
- e) Handles management configuration and diagnostic requests.

Component Descriptions

This section provides a high level description of the Static IP Routing components and the relationships between the components.

The Routing Manager Component

The Routing manager component works in the context of the Slow Forwarding thread. Major blocks of the Routing manager component are:

- *Routing endpoint configuration DB*, which is responsible for endpoint creation, deletion and modification by the asynchronous management API.
- *Routing Table* which is responsible for maintenance of the static routes including manipulations from management and creation/deletion of automatic per-endpoint connected routes.
- *Slow Path forwarder* which is responsible with making decision about packets that fast path failed to find a proper route. It performs a lookup in the static routing table. If a matching route is found, then it invokes the ARP component to resolve the destination MAC and search the switch L2 MAC table to obtain the egress port. When the forwarder eventually has all the required information, the packet is sent to destination. The packet is stored in the ARP waiting list until MAC address and port is resolved or certain period of time has expired. If no match is found the packet is dropped, counted and an ICMP packet sent back if required.

The Fast Path Component

Fast path is a part of Static IP Routing module, which runs in the context of high priority task, executing receive and transmit routines provided by the S3280 *Packet* module. The FDMA driver and *Packet* module are modified to achieve the appropriate underlying support for the fast path functionality. The details about those modifications are described in the corresponding section.

Fast path receives the captured frame from the FDMA driver. At first, the fast path checks if Static IP Routing module is the recipient of the frame, i.e. if destination MAC address of the frame is the same as one of the configured routing endpoints. All packets, which are not subject of the Static IP routing, (like inband management packets or packets for control plane protocols), are processed by the Packet module in the usual way. Every incoming packet, which is subject of the Static IP routing, goes through the validation. All packets that are neither IP nor ARP nor ICMP, are dropped. If frames are ICMP, ARP or IP dropped due to TTL=0 or failed lookup, they are forwarded to the Slow Routing task for further processing.

Packets which have successfully passed the validation and filtering go further to the fast forwarder. It performs a lookup into the routing cache table, to find the entry matching the destination IP address of the packet. A routing cache record contains all the information required for frame forwarding, which are egress port, next hop MAC address, egress VLAN and VLAN operations. If a matching record is found for the packet, then the packet is modified accordingly, and sent back to the switch towards the egress port. Every time a packet uses a routing cache entry, the timestamp of the entry is updated, allowing slow task to perform audit and aging of the cache. Timestamp is a 32 bit field which is copied from a global counter updated every 1ms by the fast path polling function.

IP packets, for which no corresponding entry is found in the routing cache, are sent for processing to the slow path task. The rate of slow path processing is limited 64 packets/sec. Cache misses exceeding this limit are dropped.

The ARP Component

ARP component works in the context of the Slow Forwarding task. It consists of the ARP cache table and the ARP protocol logic. All incoming ARP packets detected by the fast paths are forwarded to the ARP subsystem. ARP component responds to the ARP requests by looking for IP addresses of the enabled L3 endpoints. If Proxy ARP is enabled, ARP responses are also sent for IP addresses of the hosts outside of the ingress subnet but with a destination path known to the router.

Slow forwarder signals ARP component to send an ARP request if fast path failed to find a matching entry in the routing cache, but a route exists in the static routing table (meaning that the route is known, but not the MAC address of the destination/next hop). New ARP record is added to the ARP cache table with sending an ARP request.

ARP component updates existing ARP records if it receives an ARP announcement or a GARP message.

The ICMP Component

ICMP component works in the context of the Slow Forwarding task. All incoming ICMP packets captured by the data path come to this component for processing. Slow path forwarder also signals the ICMP component to send an ICMP packet if an ingress packet is dropped.

ICMP module sends ICMP responses to the Echo messages used for ping and trace purposes.

Diagnostic tools are part of the ICMP component. There are two diagnostic tools supported by the Static IP routing module. Those are ping and traceroute. ICMP component is responsible for managing ping and traceroute sessions. It holds information for each active diagnostic session and performs appropriate actions based on their states.

Hardware Layer Interface

General Information

Basically, the S3280 software interface to the CPU port consists of the FDMA driver, which deals directly with CPU FDMA part, and the *Packet* module, which configures the FDMA driver, manages it and provides an interface for applications to the driver. Each application, which needs to receive frames forwarded by the S3280 hardware switch to the CPU port, should register with the *Packet* module providing a callback for frame receive, and use the appropriate API for transmit.

Usually *Packet* module configures the FDMA driver to use interrupts that are generated by the FDMA CPU's part.

There are two types of events that are triggers for the FDMA interrupts. Those are:

- block done;
- transfer done.

The extraction channel *block done* event means that a frame (or a part of frame for jumbo frames) is ready to be processed by the software. *Transfer done* event for the extraction channel means that all elements of the current chain of DCBs that the FDMA was filling have been used, and another chain should be provided. It also may be called as an end-of-chain event.

For the injection channel, the *block done* event is not used. *Transfer done* event for the injection channel means the last chain (usually it is only one DCB) is sent and the FDMA channel is ready to send another chain that the software should provide.

When the FDMA driver receives a new frame, it invokes the appropriate callback that the *Packet* module registers with the driver. Then the *Packet* module adds the received DCB to its internal queue. All these actions are done in the DSR context. Furthermore, the RX thread in *Packet* module processes that queue by looking for the packets' recipients in its list of subscribers. When such recipient is found, *Packet* module invokes the corresponding callback that recipient has registered with the *Packet* module.

When an application needs to send a packet, it uses the corresponding API provided by the *Packet* module. The *Packet* module allocates memory for a packet, initializes a DCB and passes it to the FDMA driver, which in turn injects it to the FDMA channel.

Modifications to the CPU Port

The mechanism described above is quite suitable for processing slow protocols traffic, but has certain [disadvantages](#) that make it inapplicable for static IP routing with sufficiently high rates of incoming IP traffic. These [disadvantages](#) are:

- a lot of context switches are performed before a packet reaches its recipient on receive side;
- a lot of context switches are performed while a packet passes from an application to the FDMA;
- the S3280 chip has a bug in the FDMA interrupts system that results in the generation of unexpected extra interrupt per frame passing through the interface.

The frequent context switches consume a lot of CPU resources.

Each interrupt is a context switch as well. By having three interrupts per a frame (including one extra) a significant part of the CPU resources would be spent on context switches.

Therefore, with the Static IP Routing feature enabled, the S3280 software interface to the CPU port will work in a so-called *polling mode* for the extraction and injection FDMA channels, using only end-of-chain interrupt for the extraction channel for rapid reaction on this type of events.

However, the *Packet* module configures three other FDMA channels specifically for CCM generation purposes. These channels can also use *transfer done* interrupts per frame that would decrease efficiency of the polling scheme used for extraction and injection.

As implementation of static IP routing functionality requires making a lot of changes in these two S3280 modules, those changes should be demarcated by using the compilation flag `TN_OPT_FDMA_POLLING_MODE`.

The Packet Module

To implement the polling mode, the *Packet* module creates a periodic timer with 1 millisecond period using the *Timer* module. As the fast path needs to work in real time, with the speed as high as possible, we should secure a deterministic behaviour of the system. So actually that *timer* generates a kind of master synchronization pulse for entire software system.

Processing the extraction and injection of DCB chains is done in the appropriate FDMA driver API functions described in the related section. The polling timer handler (called every 1ms), invokes these functions one after another, in several iterations implementing a multi-pass processing algorithm in order to free DCBs faster. For each iteration, the polling timer handler invokes first the extraction channel polling API function, and then the injection channel polling API function. The number of iterations is set for high performance (currently 4 iterations per 1ms handler run).

The system scheduler is locked the entire time that the loop is working.

Interaction among Packet module, FDMA driver, L3 Routing module, and other applications is shown below.

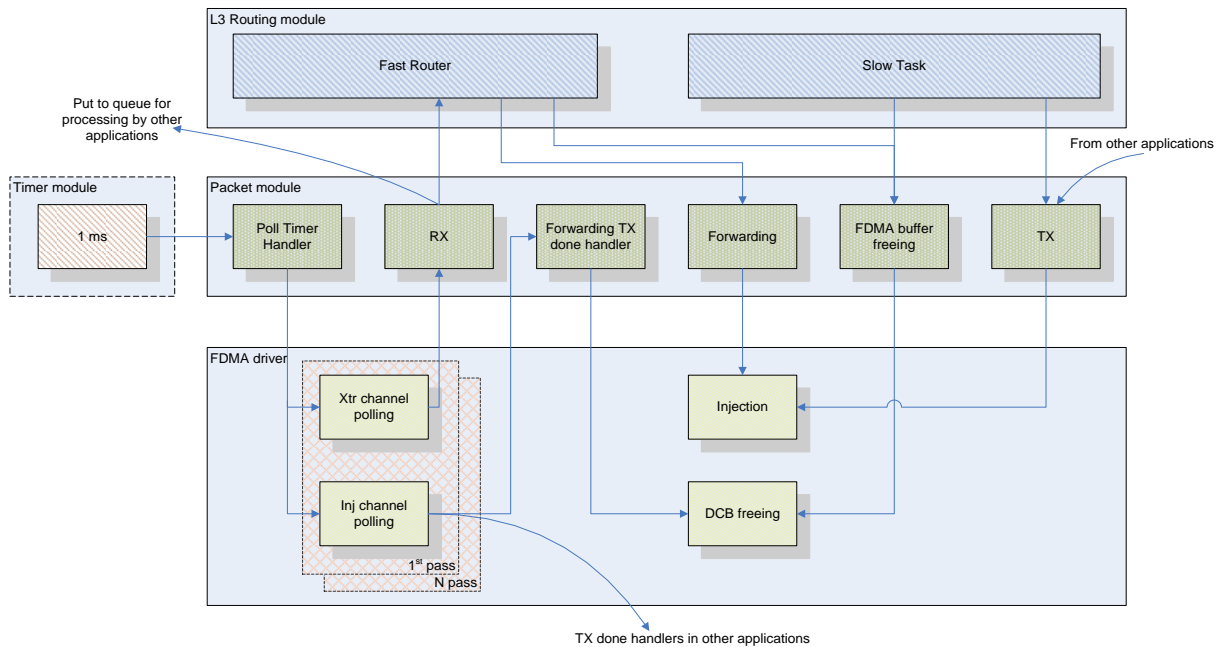


Figure 3. Software Data Path

Each of the polling API functions is working until either it completes its work, or its time limit has expired. Polling functions process DCBs one by one traversing a DCB chain. On each 10th DCB the API function calculates how many microseconds were spent processing previous 10 DCBs and decreases the given time limit with that value. If time limit runs out, processing is aborted until the next 1ms interval. Function *hal_time_get* is used for determination of expired time. It reads a HW timer and returns how many microseconds are passed from the system start.

The time limit for each of those functions and number of passes are adjustable parameters and should be set so that CPU time remains available for other tasks. To reach the required throughput performance we need to give as much time as possible to the fast path. The current design assigns is up to 80% of CPU time for the fast path and at least 20% for others. If polling functions consume less time than the configured limits, they return the CPU to the other applications.

When the extraction polling API function receives a frame, it passes the DCB pointed by that frame to the Packet module by calling the appropriate callback, as well as extraction interrupt handler does. Each received frame can be either forwarded by the Static IP Routing fast path or may go to the Packet module's receiving queue. At first, the Packet module should call the new callback provided by the fast path. If the fast path doesn't consume the frame, then the frame goes to the receiving queue and it is processed as per current procedures.

The fast path can consume a received frame in different ways; typically (cache hit case) the frame is edited and forwarded to a port (a port and a VLAN). In this case the fast path determines the egress port and egress VLAN and it invokes the corresponding new Packet module API function that performs injection operation for given DCB.

The following CPU port interface parameters are also changed to support static IP routing:

- Length of allocated chain of DCBs and data buffers for extraction is increased from 256 to 2048;
- Additional extraction queue is configured for frames subject to static IP routing. The normal priority is selected for **that queue that lower** than Management traffic priority.

2. CLI (Command Line Interface)

The SIR feature lets you define L3 routing endpoints over the VLANs configured in the NID. SIR can be configured via the CLI commands or via the web GUI. This section describes the CLI commands used for Static IP Routing configuration and maintenance.

The Static IP Routing CLI is for experienced network administrators who are responsible for configuring and maintaining the S3280. The CLI offers a comprehensive set of management features for use during initial setup (set IPs, etc.) and troubleshooting, as well as for day-to-day management (device management, firmware upgrades, managing security features, etc.).

The S3280 offers a rich set of commands through its CLI for performing configuration and status monitoring. The CLI is accessible through the RS-232 serial console, telnet and SSH. The CLI incorporates user authentication for security purposes. The CLI interface can be accessed via Secure Shell (SSH) interface. This provides a more secure interface as SSH uses public-key cryptography for authentication. When the SSH server is enabled, normal telnet access can be enabled or disabled to avoid potential security holes.

All of the other CLI commands and their functions are described in the S3280 CLI Reference Guide manual.

Static IP routing CLI commands are located in the “SIR” section of the main CLI tree. Listed below is the full set of CLI commands used for the Static IP Routing feature.

>sir ?

Available Commands:

1. **SIR Endpoint Configuration** [<ep-id>]
 2. **SIR Endpoint Create** <ep-id> <ip-addr> (vlan-aware|vlan-unaware|untagged) [<vid>] [(horizon (uplink|downlink))]
 3. **SIR Endpoint Delete** <ep-id>
 4. **SIR Endpoint Admin** [<ep-id>] [enable|disable]
 5. **SIR Endpoint Port Add** <ep-id> <port-list>
 6. **SIR Endpoint Port Remove** <ep-id> [<port-list>]
 7. **SIR Endpoint ProxyARP Admin** [<ep-id>] [enable|disable]
 8. **SIR Endpoint ProxyARP Group** [<ep-id>] [<group>]
 9. **SIR Endpoint SplitHorizon** [<ep-id>] [enable|disable]
 10. **SIR Endpoint Horizon** [<ep-id>] [uplink|downlink]
 11. **SIR Endpoint Statistics** [<ep-id>] [clear]
 12. **SIR Route Configuration**
 13. **SIR Route Add** <subnet> <gw>
 14. **SIR Route Remove** <subnet> <gw>
 15. **SIR Route DefaultGW Add** <gw>
 16. **SIR Route DefaultGW Remove** <gw>
 17. **SIR ARP Show**
 18. **SIR ARP Flush** [<subnet>]
 19. **SIR ARP Configuration** [(wait <wait-timeout>)] [(age <ageing-timeout>)] [(refresh <refresh-timeout>)] [(garp (enable|disable))]
 20. **SIR Ping** <ip-addr> [(length <ping-length>)] [(count <ping-count>)] [(interval <ping-interval>)] [(origin-ip-addr <origin-ip-addr>)]
 21. **SIR Traceroute** <ip-addr> [(hops <trace-hops>)] [(min-ttl <trace-min-ttl>)] [(queries <trace-queries>)] [(timeout <trace-timeout>)] [(origin-ip-addr <origin-ip-addr>)]
 22. **SIR Alarm Show** [<ep-id>]
- >

Changes to Existing CLI Commands

Help (?) command: the SIR command is added to the list of supported CLI commands.

System Version (sys ver) command has minor changes for SIR support.

System Configuration (sys config) command has minor changes for SIR support.

The Static IP Routing CLI commands are explained in the following sections.

Routing Endpoint Commands

Command: Show Routing Endpoint Configuration

Syntax: `sir endpoint config [<ep-id>]`

Description: Display the current Routing Endpoint configuration. If the Endpoint ID (ep-id) parameter is not specified, then the table view displays. The parameters are:

<ep-id>: Endpoint ID (L3EP1-16).

Example 1:

```
>sir endpoint config
L3 Static Routing Endpoint Configuration:
=====
ID          IP Address      Admin   Vlan Tag   VID   Ports      ProxyARP  Group  Split Horizon  Horizon
-----
L3EP1      10.100.5.20/24 disabled  untagged   0     none       enabled   1     disabled       downlink
>
```

Example 2:

```
>sir endpoint config L3EP1
ID          : L3EP1
IP Address  : 10.100.5.20/24
Admin       : disabled
Vlan Usage  : untagged
VID         : 0
Ports       : none
ProxyARP Mode : enabled
ProxyARP Group: 1
Split Horizon : disabled
Horizon     : downlink
MAC Address  : 00-c0-f2-00-00-0c
>
```

Messages: *No created L3 endpoints (if none exist), or
L3 Static Routing Endpoint L3EP1 has not been created yet*

Command: Create New Routing Endpoint

Syntax: **sir endpoint create** <ep-id> <ip-addr> (vlan-aware|vlan-unaware|untagged) [<vid>] [(horizon|uplink|downlink)]

Description: Create a new L3 Static IP Routing Endpoint with default parameters, where:

- <ep-id> : Endpoint ID (L3EP1-16) to be created.
- <ip-addr> : IPv4 address in CIDR notation (a.b.c.d/m).
- vlan-aware|vlan-unaware|untagged:** VLAN mode (aware, unaware, or untagged).
- <vid> : VLAN ID (0-4094), if 0 then unaware or untagged.
- horizon** : Horizon direction keyword
- uplink|downlink** : Horizon direction.

Example:

```
>sir endpoint create L3EP1 10.100.5.20/24
>sir endpoint config
```

L3 Static Routing Endpoint Configuration:

```
=====
ID   IP Address   Admin  Vlan Tag  VID Ports  ProxyARP  Group Split Horizon Horizon
-----
L3EP1 10.100.5.20/24 disabled untagged 0 none    enabled   1 disabled downlink
>
```

Messages: *Invalid <ep-id> parameter: 2 (if the VID does not exist).*

*W tn_sip/l3 05:31:02 70/tn_sip_routing_endpoint_add#494: Warning: l3 match mgmt IP
Creating of L3 Endpoint is failed, L3 endpoint IP equal management IP*

Note: A newly created endpoint by default is disabled, ARP proxy enabled, not uplink, split horizon disabled, belongs to Proxy ARP group 1. "Split-Horizon" here means "forced forwarding of received IP frames on customer side interfaces to a designated network side interface" so that routing of frames does not occur on the S3280, but at a routing point higher in the hierarchy.

Validation: a VLAN with a VID must exist; the IP address must be unique, and no overlapping subnets can exist.

Command: Delete Existing Routing Endpoint

Syntax: **sir endpoint delete** <ep-id>

Description: Delete an existing L3 Static IP Routing Endpoint, where:

- <ep-id>: Endpoint ID (L3EP1-16) to be deleted.

Example:

```
>sir endpoint delete L3EP1
>
```

Messages: *Invalid <ep-id> parameter: 1 if the endpoint does not exist.*

Validation: The routing endpoint with configured static routes cannot be deleted.

Command: Enable or Disable Routing Endpoint

Syntax: `sir endpoint admin [<ep-id>] [enable|disable]`

Description: Set or show L3 Static IP Routing Endpoint administrative state, where:
 <ep-id> : Endpoint ID (L3EP1-16) to be displayed or enabled/disabled.
 enable|disable: enable or disable the specified Endpoint. The default is 'Show administrative state'.

```
Example 1: >sir endpoint admin
           L3 Static Routing Endpoint Admin State:
           =====
           No created L3 endpoints
           >
```

```
Example 2: >sir endpoint admin
           L3 Static Routing Endpoint Admin State:
           =====
           ID          IP Address          Admin
           -----
           L3EP1      0.0.0.0          disabled
           >
```

Messages:
 L3 Endpoint has not been created yet
 Changing admin status is failed, Cannot admin up L3 endpoint with no ports assigned to it

Command: Add Port to Routing Endpoint

Syntax: `sir endpoint port add <ep-id> <port-list>`

Description: Add a port or list of ports to L3 Static IP Routing Endpoint, where:
 <ep-id> : Endpoint ID (L3EP1-16) of the port.
 <port-list>: Port or ports list separated by comma (a,b,...).

```
Example: >sir endpoint port add L3EP1 1,5,8
         >sir endpoint admin l3ep1 enable
         Changing admin status is failed, Cannot admin up L3 endpoint with no
         ports assigned to it
         >sir endpoint port add L3EP1 1,5,8
         >sir endpoint admin l3ep1 enable
         >sir endpoint admin
           L3 Static Routing Endpoint Admin State:
           =====
           ID          IP Address          Admin
           -----
           L3EP1      10.100.5.20/24    enabled
           >
```

Messages: L3 Endpoint has not been created yet
Validation: The Port must be a member of the VLAN of the L3 Endpoint.

Command: Remove Port from Routing Endpoint

Syntax: `sir endpoint port remove <ep-id> [port-list]`

Description: Remove a port or list of ports from L3 Static IP Routing Endpoint, where:
 <ep-id> : Endpoint ID (L3EP1-16).

port-list : Port or ports list separated by comma (a,b,...: the default is 'remove all ports').

Example:

```
>sir endpoint port remove L3EP1 5,8
>
```

Command: Enable/Disable Proxy ARP

Syntax: `sir endpoint proxyarp admin [<ep-id>] [enable|disable]`

Description: Enable/Disable or show Proxy ARP for L3 Static IP Routing Endpoint, where:
 <ep-id> : Endpoint ID (L3EP1-16).

enable|disable: enable or disable the show proxy ARP admin state for the routing endpoint.

Example 1: `>sir end proxyarp admin`

```
L3 Static Routing Endpoint Proxy ARP Configuration:
=====
```

```
No created L3 endpoints
>
```

Example 2: `>sir endpoint proxyARP admin`

```
>
ID      ProxyARP  Group
=====  =====  =====
L3EP1   enabled    1
L3EP5   disabled   1
L3EP8   enabled    1
L3EP9   enabled    2
```

Messages: *Setting proxy ARP admin status is failed, L3 endpoint is enabled*

Command: Set Proxy ARP Group

Syntax: `sir endpoint proxyarp group [<ep-id>] [<group>]`

Description: Configure Proxy ARP Group for L3 Static IP Routing Endpoint, where:
 <ep-id>: Endpoint ID (L3EP1-16).

<group>: ARP group (1-16; the default is group 1).

Example: `>sir endpoint proxyARP group L3EP1 16`

```
>
>sir endpoint proxyarp group
```

```
L3 Static Routing Endpoint Proxy ARP Configuration:
=====
```

```
No created L3 endpoints
```

```
>sir endpoint proxyARP group L3EP1 16
```

```
L3 Endpoint has not been created yet
```

```
>
```

Messages: *Changing proxy ARP group is failed, L3 endpoint is enabled*

Command: **Enable/Disable Split Horizon**
Syntax: **sir endpoint splithorizon** [<ep-id>] [enable|disable]
Description: Enable/Disable Split Horizon for L3 Static IP Routing Endpoint, where:
 <ep-id> : Endpoint ID (L3EP1-16).
enable|disable: enables or disables the selected split horizon for routing endpoint.
Enabled = IP routes learned from a neighbor are not advertised back in updates to that neighbor.
Disabled: all routes in the route table are advertised through this interface with no change to the routing metric.

Example 1: SIR>**endpoint splithorizon L3EP1 enable**
 SIR>

Example 2: SIR>**endpoint splithorizon**

ID	Split Horizon	Horizon
L3EP1	disabled	downlink
L3EP5	disabled	downlink
L3EP8	disabled	downlink
L3EP9	enabled	uplink

Example 3: >**sir endpoint splithorizon L3EP1 enable**
 >**sir endpoint splithorizon**

L3 Static Routing Endpoint Split Horizon Configuration:
 =====

ID	IP Address	Split Horizon	Horizon
L3EP1	0.0.0.0	enabled	downlink

>

Messages:
 No created L3 endpoints
 Setting Split Horizon mode is failed, L3 endpoint is enabled

Command: **Set Horizon**
Syntax: **sir endpoint horizon** [<ep-id>] [uplink|downlink]
Description: Configure or show Horizon for L3 Static IP Routing Endpoint, where:
 <ep-id> : Endpoint ID (L3EP1-16).
uplink|downlink: Horizon direction (uplink or downlink direction).

Example: >**sir endpoint horizon L3EP1 uplink**
 >**sir endpoint horizon**

L3 Static Routing Endpoint Split Horizon Configuration:
 =====

ID	IP Address	Split Horizon	Horizon
L3EP1	0.0.0.0	enabled	uplink

>

Messages:
 L3 Endpoint has not been created yet
 Setting Horizon direction is failed, L3 endpoint is enabled

Command: **Routing Endpoint Statistics**

Syntax: **sir endpoint statistics** [<ep-id>] [clear]

Description: Show L3 Static IP Routing Endpoint Statistics. Show routing endpoint statistics or clear and show if 'clear' command is specified, where:

 <ep-id>: Endpoint ID (L3EP1-16).

clear : Clear the log file.

Example:

```
>sir endpoint statistics l3ep1
```

```
L3 Static Routing Endpoint L3EP1 Statistics
```

Rx Packets:	0	Tx Packets:	1
Rx Octets:	0	Tx Octets:	42
Rx Unicast:	0	Tx Unicast:	0
Rx Multicast:	0	Tx Multicast:	0
Rx Broadcast:	0	Tx Broadcast:	1
Rx Forw Datagrams:	0	Tx Forw Datagrams:	0
		Tx Requests:	0
Rx Discards:	0	Tx Discards:	0
Rx No Routes:	0		
Rx Addr Errors:	0		
Rx Unknown Protos:	0		
Rx Header Errors:	0		
Rx ICMP Packets:	0	Tx ICMP Packets:	0
Rx ICMP Discards:	0		
>			

Routing Table Commands

Command: Show Routing Table

Syntax: `sir route config`

Description: Show the current L3 Static IP Routing Table. Displays the current SIR route configuration in terms of Destination, Gateway, Egress Endpoint, VID, and Horizon information.

Example: `>sir route configuration`

```
Routing Table:
=====
Destination          Gateway              Egress Endpoint    VID    Horizon
-----
10.100.5.0/24        *                   L3EP1              0     downlink
>
```

Note: The first entry in this example is automatically created when routing endpoint 1 (L3EP1) is created in the system.

Messages: *Routing Table is empty (if none configured)*

Command: Add New Routing Entry

Syntax: `sir route add <subnet> <gw>`

Description: Add a new static route entry to the routing table, where:

<subnet>: IPv4 subnet address in CIDR notation (a.b.c.d/m).

<gw> : IPv4 gateway address (a.b.c.d).

Example: `>sir route configuration`

```
Routing Table:
=====
Destination          Gateway              Egress Endpoint    VID    Horizon
-----
10.100.5.0/24        *                   L3EP1              0     downlink
>sir route add 10.100.6.0/24 10.100.5.46
>sir route configuration

Routing Table:
=====
Destination          Gateway              Egress Endpoint    VID    Horizon
-----
10.100.5.0/24        *                   L3EP1              0     downlink
10.100.6.0/24        10.100.5.46        L3EP1              0     downlink
>
```

Validation: Gateway IP address must be a member of one of the subnet directly connected to S3280.

Messages: *Adding of a static route is failed, Gateway not on connected net*

Command: **Remove Existing Routing Entry**

Syntax: **sir route remove** <subnet> <gw>

Description: Remove existing static route entry from the routing table, where:
 <subnet>: IPv4 subnet address in CIDR notation (a.b.c.d/m).
 <gw> : IPv4 gateway address (a.b.c.d).

Example: >**sir route configuration**

Routing Table:

=====

Destination	Gateway	Egress Endpoint	VID	Horizon
10.100.5.0/24	*	L3EP1	0	downlink
10.100.6.0/24	10.100.5.46	L3EP1	0	downlink

>**sir route remove 10.100.6.0/24 10.100.5.46**

>**sir route configuration**

Routing Table:

=====

Destination	Gateway	Egress Endpoint	VID	Horizon
10.100.5.0/24	*	L3EP1	0	downlink

>

Messages: *Removing of a route is failed, Object is not found*

Command: **Add Default Gateway**

Syntax: **sir route defaultgw add** <gw>

Description: Add default gateway to the routing table, where:
 <gw> : IPv4 gateway address (a.b.c.d) to be added.

Example: >**sir route defaultgw add 10.100.6.1**

>

Messages: *Adding of default gateway is failed, Gateway not on connected net*

Note: Up to two default gateways can be defined per system as any combination of common uplink or downlink horizon gateways.

Validation: The Gateway IP address must be a member of one of the subnets directly connected to the S3280 NID.

Command: **Remove Default Gateway**

Syntax: **sir route defaultgw remove** <gw> [uplink|downlink]

Description: Remove the default gateway from the routing table, where:
 <gw> : IPv4 gateway address (a.b.c.d).

Example: >**sir route defaultgw remove 10.100.6.1**

>

Messages: *Removing of default gateway is failed, Object is not found if default gateway does not exist.*

ARP Table Commands

Command: Show ARP Table

Syntax: sir arp show

Description: Display the current ARP table in terms of its Destination, Endpoint, VID, and MAC Address information.

Example 1: > **sir arp show**

```
L3 Static Routing ARP table:
=====

ARP Table is empty
>
```

Command: Flush ARP Table

Syntax: sir arp flush [subnet]

Description: Remove entries from an ARP table for given subnet. If the subnet is empty, then remove all ARP entries from the table. The parameters are:

subnet : The IPv4 subnet address in CIDR notation (a.b.c.d/m) to flush.

Example: > **sir arp flush 10.100.6.0/24**
> **sir arp flush 192.168.1.110**
> **sir arp flush 192.168.1.110/32**
>

Command: Set / Show SIR ARP Configuration Parameters

Syntax: **sir arp config** [(wait <wait-timeout>)] [(age <ageing-timeout>)] [(refresh <refresh-timeout>)] [(garp (enable disable))]

Description: Show or set and show ARP configuration parameters. The default is “Show GARP operations state”. The parameters are:

wait : ARP wait timeout keyword.

<**wait-timeout**> : Time between ARP request send and declaration of failed (1-10; the default is 2 seconds).

age : ARP aging timeout keyword.

<**ageing-timeout**> : Time to keep unused ARP record in the table (60-86400 seconds; the default is 120 seconds)

refresh : ARP refresh timeout keyword.

<**refresh-timeout**>: Time between consequential refreshing ARP requests sent for the record in the table (the valid range is 60-86400 seconds; the default is 600 seconds).

garp : GARP operations state keyword. Status of Gratuitous ARP. If it is disabled, all received GARP messages are ignored and no GARP messages are generated by the system. By default, Gratuitous ARP is enabled.

enable : Enable GARP operations

disable: Disable GARP operations

Example:

```
>sir arp config
Wait Timeout      : 2
Ageing Timeout    : 120
Refresh Timeout   : 60
GARP              : enabled
>sir arp config wait 4 age 60 refresh 300 garp enable
Wait Timeout      : 4
Ageing Timeout    : 60
Refresh Timeout   : 300
GARP              : enabled
>
```

Network Diagnostic Commands

The SIR module provides the *ping* and *traceroute* commands for diagnosing network problems.

Traceroute is a network diagnostic tool for displaying the route (path) and measuring transit delays of packets across an IP (Internet Protocol) network.

Ping is a computer network tool used to test whether a specified host is reachable across an IP network.

Command: **Ping**
Syntax: **sir ping** <ip-addr> [(length <ping-length>)] [(count <ping-count>)] [(interval <ping-interval>)] [(origin-ip-addr <origin-ip-addr>)]
Description: Ping IP address (ICMP echo), where:
 <ip-addr> : IPv4 destination address (a.b.c.d).
 length : Ping length keyword.
 <ping-length> : Ping ICMP data length (2-1452; the default is 56), excluding MAC, IP, and ICMP headers.
 count : Ping count keyword.
 <ping-count> : Transmit ECHO_REQUEST packet count (1-60; the default is 5 packets).
 interval : Ping interval keyword
 <ping-interval>: Ping interval (0-30; the default is 0).
 origin : Ping/Trace origin IP address keyword
 <ip-addr>: Ping/Trace origin IPv4 address in CIDR notation (a.b.c.d/m).

Example 1: **SIR>ping 10.100.5.20**
 PING server 10.100.5.20, 56 bytes of data.
 64 bytes from 10.100.5.20: icmp_seq=0, time=10ms
 64 bytes from 10.100.5.20: icmp_seq=1, time=0ms
 64 bytes from 10.100.5.20: icmp_seq=2, time=0ms
 64 bytes from 10.100.5.20: icmp_seq=3, time=0ms
 64 bytes from 10.100.5.20: icmp_seq=4, time=0ms
 Sent 5 packets, received 5 OK, 0 bad
 SIR>

Example 2: **SIR>ping 10.100.5.20 length 1000 interval 5**
 PING server 10.100.5.20, 1000 bytes of data.
 64 bytes from 10.100.5.20: icmp_seq=0, time=10ms
 64 bytes from 10.100.5.20: icmp_seq=1, time=0ms
 64 bytes from 10.100.5.20: icmp_seq=2, time=0ms
 64 bytes from 10.100.5.20: icmp_seq=3, time=0ms
 64 bytes from 10.100.5.20: icmp_seq=4, time=0ms
 Sent 5 packets, received 5 OK, 0 bad
 SIR>

Example: **>sir ping 10.100.5.20 length 1000 interval 5**
 PING server 10.100.5.20, 1000 bytes of data.
 Request timed out
 Request timed out
 Request timed out
 Request timed out
 Request timed out
 Sent 5 packets, received 0 OK, 0 bad
 >

Command: Traceroute

Syntax: **sir traceroute** <ip-addr> [(hops <trace-hops>)] [(min-ttl <trace-min-ttl>)] [(queries <trace-queries>)] [(timeout <trace-timeout>)] [(origin-ip-addr <origin-ip-addr>)]

Description: Trace IP address, where:

<ip-addr> : IPv4 destination address (a.b.c.d).

hops : Trace hops keyword.

<trace-hops> : MAX number of traced hops (1-50; the default is 10).

min-ttl : Trace MIN TTL keyword.

<trace-min-ttl>: transmit TTL value of first packet (1-255; the default is 1).

queries : Trace queries keyword

<trace-queries>: Transmit packet tries (1-10 tries; the default is 1 attempt).

timeout : Trace timeout keyword.

<trace-timeout>: Reply waiting time (1-10 seconds; the default is 2 seconds).

origin : Ping/Trace origin IP address keyword

<ip-addr>: Ping/Trace origin IPv4 address in CIDR notation (a.b.c.d/m)

Example 1:

```
SIR>traceroute 10.100.5.20
trace to 10.100.5.20, 10 hops max, start from TTL 1
1: 10.10.10.1          0.092ms
2: 10.10.100.121      0.815ms
3: 10.10.100.254      1.645ms
4: 10.100.5.254       2.357ms
5: 10.100.5.20        3.257ms
Resume: reached, hops 5
>
```

Example 2:

```
SIR>traceroute 10.100.5.20 min-ttl 4 timeout 10
trace to 10.100.5.20, 10 hops max, start from TTL 4
1: 10.100.5.254       2.357ms
2: 10.100.5.20        3.257ms
Resume: reached, hops 2
```

Example 3:

```
>sir traceroute 10.100.5.20
TRACE to 10.100.5.20, 10 hops max, start from TTL 1.
1:
2:
3:
4:
5:
6:
7:
8:
9:
10:
Resume: unreachable, hops 9
>
```

Example 4:

```
>sir traceroute 192.168.1.110 hops 6 min-ttl 2 queries 2 timeout 1 origin
192.168.1.10
Start trace is failed, Invalid parameter
>
```

Command: Show SIR Alarm

Syntax: **sir alarm show** [<ep-id>]

Description: Show the current SIR Endpoint alarms configured, where:
<ep-id>: Endpoint ID (L3EP<1-16>). If no entry, all existing L3EPs are displayed.

Example 1: >**sir alarm show**

```
L3 Static Routing Endpoint Alarms:
```

```
=====
```

ID	IP Address	Alarms
-----	-----	-----
L3EP1	10.100.5.20/24	No alarms

```
>
```

Example 2: >**sir alarm show**

```
L3 Static Routing Endpoint Alarms:
```

```
=====
```

```
No created L3 endpoints
```

```
>
```

SIR Configuration Example (via CLI)

The example below uses a single test setup for all test cases. There are six NIDs in the setup, running Static IP routing. “SB” devices are either IXIA or SmartBits devices, which allow transmission of multiple flows in parallel, with editable IP/Ethernet headers. PCs can be either Windows- or Linux-based, and they can be virtual machines to save resources. PC1 and PC2 have FTP servers enabled, and at least PC4 must be able to capture received packets.

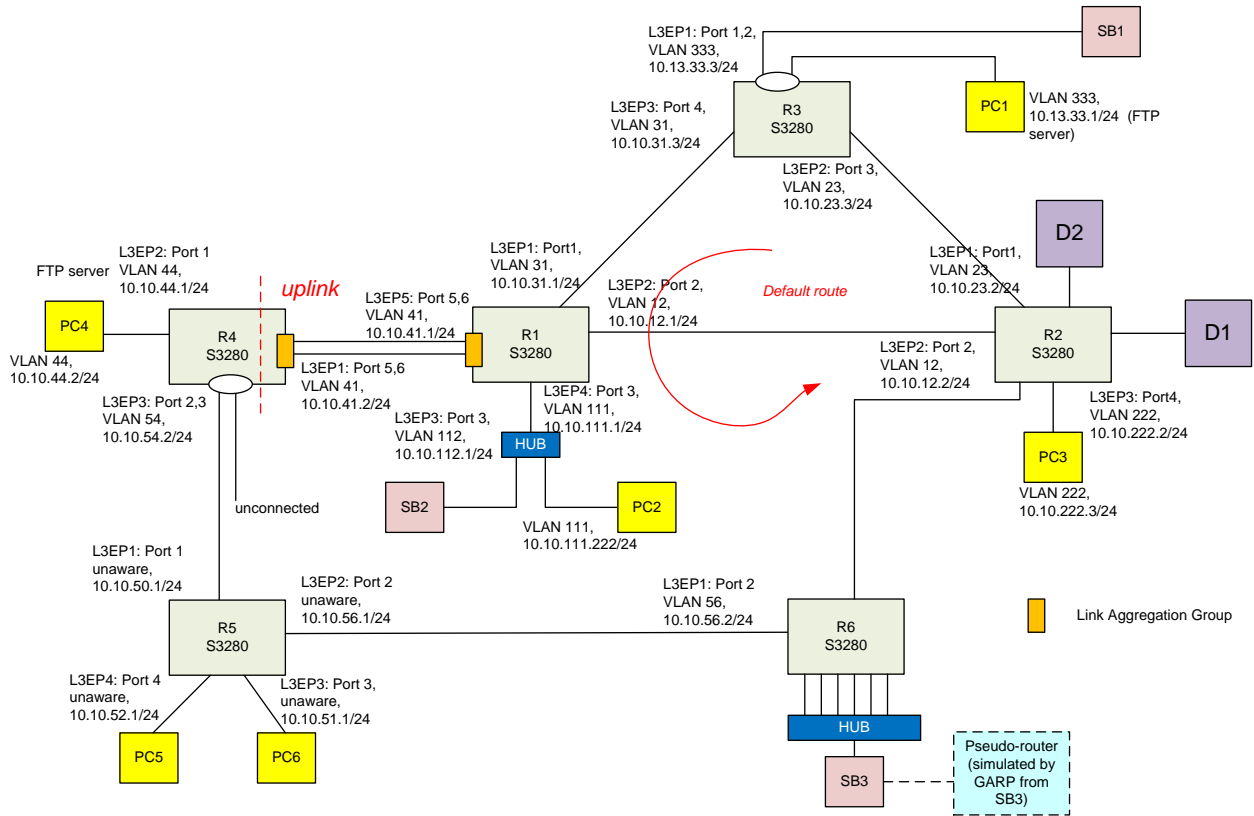


Figure 4. Test Setup

The devices in the setup are configured such that:

- The Default route configuration follows a counter-clockwise direction over the big loop (see the red arrow in the figure above).
- R5 works in VLAN unaware mode.
- R4 works in split horizon mode. The Uplink L3EP is connected via a LAG to R1.
- The link between R6 and R2 is asymmetric (there is one VLAN for R6->R2 traffic and a different VLAN for R2->R6 traffic).
- The route between PC1 and PC5/PC6 is asymmetric. The PC1->PC4 direction goes through R3->R1->R4->R5, while the opposite direction goes through R5->R6->R2->R3.
- R6 has 16 L3EPs, and all ports are connected. Some of the ports connected to the Hub have two L3EPs, on different VLANs. The subnets facing the Hub and SB3 device are large subnets (22 bits subnet mask). Also, on R6 there are two static routes for large subnets (22 bits subnet masks) pointing toward a pseudo-router on the SB3 side (dashed box in the diagram).
- R5 and PC5/PC6 are configured for Proxy ARP.

- R3 and R4 have L3EPs covering 2 ports (in the same VLAN).
- All other known subnets are configured in the routers based on the shortest path (least amount of nodes to traverse); D1 and D2 are devices which allow the measurement of latency and throughput at L3 level.

Also, there are two VLANs configured across all devices; one purely for L2 traffic (and for regression tests) and the other one for PTP synchronization.

For PTP purposes, the R3 is the master, and all NIDs are synchronized in the same PTP domain.

Background Operations

During the tests, the following operations are executed continuously by scripts, in order to simulate a stress in the configuration process.

- On R3, one L3EPs on VLAN 500 (a VLAN not used otherwise for traffic tests) will be continuously created and removed on the port facing R2, roughly every 2 minutes. This should trigger gratuitous ARPs, which should be visible on the R2 node (the R2 node will have a L3EP on the same VLAN 500 not used for other purposes).
- On R1, a script creates and deletes 50 static routes towards R2. These routes are not used for any particular routing purpose.
- A script on R5 issues periodically (every 10 minutes) ARP flush commands.
- A continuous SNMP walk will be executed on R6 and R4. (R6 is the router with the largest routing/ARP tables, and R4 has the split horizon configuration).
- CLI automation scripts are provided later in this chapter.

A detailed CLI configuration is provided below.

Detailed Configuration (CLI)

R1 S3280

*/*R1->R3*/*

```
sir endpoint create l3ep1 10.10.31.1/24 vlan-aware 31
sir endpoint port add l3ep1 1
sir endpoint admin l3ep1 enable
```

*/*R1->R2*/*

```
sir endpoint create l3ep2 10.10.12.1/24 vlan-aware 12
sir endpoint port add l3ep1 2
sir endpoint admin l3ep2 enable
```

*/*R1->HUB->SB2*/*

```
sir endpoint create l3ep3 10.10.112.1/24 vlan-aware 112
sir endpoint port add l3ep3 3
sir endpoint admin l3ep3 enable
```

*/*R1->HUB->PC2*/*

```
sir endpoint create l3ep4 10.10.111.1/24 vlan-aware 111
sir endpoint port add l3ep4 3
sir endpoint admin l3ep4 enable
```

*/*R1->R4*/*

```
sir endpoint create l3ep5 10.10.40.1/24 vlan-aware 40
sir endpoint port add l3ep5 5
sir endpoint port add l3ep5 6
sir endpoint admin l3ep5 enable
```

*/*Default*/*

```
sir route defaultgw add 10.10.41.2
```

/->R2*/*

```
sir route add 10.10.222.0/24 10.10.12.2
sir route add 10.10.201.0/24 10.10.12.2
sir route add 10.10.202.0/24 10.10.12.2
sir route add 10.10.26.0/24 10.10.12.2
```

/->R3*/*

```
sir route add 10.13.33.0/24 10.10.31.3
```

/->R4*/*

```
sir route add 10.10.44.0/24 10.10.41.2
sir route add 10.10.54.0/24 10.10.41.2
```

/->R5*/*

```
sir route add 10.10.51.0/24 10.10.41.2
sir route add 10.10.52.0/24 10.10.41.2
sir route add 10.10.56.0/24 10.10.41.2
```

/->R6*/*

```
sir route add 10.10.62.0/24 10.10.12.2
sir route add 10.100.4.0/22 10.10.12.2
sir route add 10.101.4.0/22 10.10.12.2
```

R2 S3280

*/*R2->R3*/*

```
sir endpoint create l3ep1 10.10.23.2/24 vlan-aware 23
sir endpoint port add l3ep1 1
sir endpoint admin l3ep1 enable
```

```
/*R2->R1*/
```

```
sir endpoint create l3ep2 10.10.12.2/24 vlan-aware 12  
sir endpoint port add l3ep2 2  
sir endpoint admin l3ep2 enable
```

```
/*R2->PC3*/
```

```
sir endpoint create l3ep3 10.10.222.2/24 vlan-aware 222  
sir endpoint port add l3ep3 3  
sir endpoint admin l3ep3 enable
```

```
/*R2->R6*/
```

```
sir endpoint create l3ep4 10.10.62.2/24 vlan-aware 62  
sir endpoint port add l3ep4 4  
sir endpoint admin l3ep4 enable
```

```
/*R2->R6*/
```

```
sir endpoint create l3ep5 10.10.26.2/24 vlan-aware 26  
sir endpoint port add l3ep5 5  
sir endpoint admin l3ep5 enable
```

```
/*R2->D1*/
```

```
sir endpoint create l3ep6 10.10.201.1/24 vlan-aware 201  
sir endpoint port add l3ep6 6  
sir endpoint admin l3ep6 enable
```

```
/*R2->D2*/
```

```
sir endpoint create l3ep7 10.10.202.1/24 vlan-aware 202  
sir endpoint port add l3ep7 7  
sir endpoint admin l3ep7 enable
```

```
/*Default*/
```

```
sir route defaultgw add 10.10.23.3
```

```
/*->R1*/
```

```
sir route add 10.10.112.0/24 10.10.12.1
```

```
sir route add 10.10.111.0/24 10.10.12.1
```

```
/*->R3*/
```

```
sir route add 10.13.33.0/24 10.10.23.3
```

```
/*->R4*/
```

```
sir route add 10.10.41.0/24 10.10.12.1
```

```
sir route add 10.10.44.0/24 10.10.12.1
```

```
/*->R5*/
```

```
sir route add 10.10.54.0/24 10.10.12.1
```

```
sir route add 10.10.51.0/24 10.10.12.1
```

```
sir route add 10.10.52.0/24 10.10.12.1
```

```
/*->R6*/
```

```
sir route add 10.10.56.0/24 10.10.26.2
```

```
sir route add 10.10.100.4.0/22 10.10.26.2
```

```
sir route add 10.10.101.4.0/22 10.10.26.2
```


R3 S3280*/*R3->SB1/PC1*/*

```
sir endpoint create l3ep1 10.10.33.3/24 vlan-aware 33
sir endpoint port add l3ep1 1
sir endpoint port add l3ep1 2
sir endpoint admin l3ep1 enable
```

*/*R3->R2*/*

```
sir endpoint create l3ep2 10.10.23.3/24 vlan-aware 23
sir endpoint port add l3ep2 3
sir endpoint admin l3ep2 enable
```

*/*R3->R1*/*

```
sir endpoint create l3ep3 10.10.31.3/24 vlan-aware 31
sir endpoint port add l3ep3 4
sir endpoint admin l3ep3 enable
```

*/*Default*/*

```
sir route defaultgw add 10.10.31.1
```

/->R1*/*

```
sir route add 10.10.112.0/24 10.10.31.1
sir route add 10.10.111.0/24 10.10.31.1
sir route add 10.10.41.0/24 10.10.31.1
```

/->R2*/*

```
sir route add 10.10.222.0/24 10.10.23.2
sir route add 10.10.201.0/24 10.10.23.2
sir route add 10.10.202.0/24 10.10.23.2
sir route add 10.10.26.0/24 10.10.23.2
```

/->R4*/*

```
sir route add 10.10.41.0/24 10.10.31.1
sir route add 10.10.44.0/24 10.10.31.1
```

/->R5*/*

```
sir route add 10.10.54.0/24 10.10.31.1
sir route add 10.10.51.0/24 10.10.31.1
sir route add 10.10.52.0/24 10.10.31.1
sir route add 10.10.56.0/24 10.10.31.1
```

/->R6*/*

```
sir route add 10.10.62.0/24 10.10.31.1
sir route add 10.10.101.0/22 10.10.31.1
sir route add 10.10.104.0/22 10.10.31.1
```

R4 S3280*/*R4->R1*/*

```
sir endpoint create l3ep1 10.10.41.2/24 vlan-aware 41
sir endpoint port add l3ep1 5
sir endpoint port add l3ep1 6
sir endpoint horizon l3ep1 uplink
sir endpoint splithorizon l3ep1 enable
sir endpoint admin l3ep1 enable
```

*/*R4->PC4*/*

```
sir endpoint create l3ep2 10.10.44.1/24 vlan-aware 44
sir endpoint port add l3ep2 1
sir endpoint horizon l3ep2 downlink
sir endpoint splithorizon l3ep2 enable
sir endpoint admin l3ep2 enable
```

/*R4->R5*/

```
sir endpoint create l3ep3 10.10.54.2/24 vlan-aware 500
sir endpoint port add l3ep3 2
sir endpoint port add l3ep3 3
sir endpoint horizon l3ep3 downlink
sir endpoint splthorizon l3ep3 enable
sir endpoint admin l3ep3 enable
```

/*Default*/

```
sir route defaultgw add 10.10.54.1
```

/*->R1*/

```
sir route add 10.10.112.0/24 10.10.41.1
sir route add 10.10.111.0/24 10.10.41.1
sir route add 10.10.12.0/24 10.10.41.1
```

/*->R2*/

```
sir route add 10.10.222.0/24 10.10.41.1
sir route add 10.10.201.0/24 10.10.41.1
sir route add 10.10.202.0/24 10.10.41.1
sir route add 10.10.26.0/24 10.10.41.1
```

/*->R3*/

```
sir route add 10.13.33.0/24 10.10.41.1
sir route add 10.13.31.0/24 10.10.41.1
sir route add 10.13.23.0/24 10.10.41.1
```

/*->R5*/

```
sir route add 10.10.51.0/24 10.10.41.1
sir route add 10.10.52.0/24 10.10.41.1
sir route add 10.10.56.0/24 10.10.54.1
```

/*->R6*/

```
sir route add 10.10.62.0/24 10.10.41.1
sir route add 10.100.4.0/22 10.10.41.1
sir route add 10.101.4.0/22 10.10.41.1
```

R5 S3280

/*R5->R4*/

```
sir endpoint create l3ep1 10.10.54.1/24 vlan-unaware
sir endpoint port add l3ep1 1
sir endpoint proxyarp group l3ep1 5
sir endpoint proxyarp admin l3ep1 enable
sir endpoint admin l3ep1 enable
```

/*R5->R6*/

```
sir endpoint create l3ep2 10.10.56.1/24 vlan-unaware
sir endpoint port add l3ep2 2
sir endpoint proxyarp group l3ep2 5
sir endpoint proxyarp admin l3ep2 enable
sir endpoint admin l3ep2 enable
```

/*R5->PC5*/

```
sir endpoint create l3ep3 10.10.51.1/24 vlan-unaware
sir endpoint port add l3ep3 3
sir endpoint proxyarp group l3ep3 5
sir endpoint proxyarp admin l3ep3 enable
sir endpoint admin l3ep3 enable
```

/*R5->PC6*/

```
sir endpoint create l3ep4 10.10.52.1/24 vlan-unaware
sir endpoint port add l3ep4 4
sir endpoint proxyarp group l3ep4 5
sir endpoint proxyarp admin l3ep4 enable
sir endpoint admin l3ep4 enable
```

/*Default*/

```
sir route defaultgw add 10.10.56.2
```

/*->R1*/

```
sir route add 10.10.112.0/24 10.10.54.2
sir route add 10.10.111.0/24 10.10.54.2
sir route add 10.10.12.0/24 10.10.54.2
```

/*->R2*/

```
sir route add 10.10.222.0/24 10.10.56.2
sir route add 10.10.201.0/24 10.10.56.2
sir route add 10.10.202.0/24 10.10.56.2
sir route add 10.10.26.0/24 10.10.56.2
```

/*->R3*/

```
sir route add 10.13.33.0/24 10.10.56.2
sir route add 10.13.31.0/24 10.10.56.2
sir route add 10.13.23.0/24 10.10.56.2
```

/*->R4*/

```
sir route add 10.10.41.0/24 10.10.56.2
sir route add 10.10.44.0/24 10.10.56.2
```

/*->R6*/

```
sir route add 10.10.62.0/24 10.10.56.2
sir route add 10.100.4.0/22 10.10.56.2
sir route add 10.101.4.0/22 10.10.56.2
```

R6 S3280**/*R6->R5*/**

```
sir endpoint create l3ep1 10.10.56.2/24 vlan-aware 500
sir endpoint port add l3ep1 2
sir endpoint admin l3ep1 enable
```

/*R6->R2*/

```
sir endpoint create l3ep2 10.10.62.1/24 vlan-aware 62
sir endpoint port add l3ep2 1
sir endpoint admin l3ep2 enable
```

```
sir endpoint create l3ep3 10.10.26.1/24 vlan-aware 26
sir endpoint port add l3ep3 1
sir endpoint admin l3ep3 enable
```

/*R6->HUB->SB3/Pseudo-router*/

```
sir endpoint create l3ep4 60.60.60.1/22 vlan-aware 601
sir endpoint port add l3ep4 3
sir endpoint admin l3ep4 enable

sir endpoint create l3ep5 60.60.64.1/22 vlan-aware 602
sir endpoint port add l3ep5 4
sir endpoint admin l3ep5 enable

sir endpoint create l3ep6 60.60.68.1/22 vlan-aware 603
sir endpoint port add l3ep6 5
sir endpoint admin l3ep6 enable

sir endpoint create l3ep7 60.60.72.1/22 vlan-aware 604
sir endpoint port add l3ep7 6
sir endpoint admin l3ep7 enable

sir endpoint create l3ep8 60.60.76.1/22 vlan-aware 605
sir endpoint port add l3ep8 7
sir endpoint admin l3ep8 enable

sir endpoint create l3ep9 60.60.80.1/22 vlan-aware 606
sir endpoint port add l3ep9 8
sir endpoint admin l3ep9 enable

sir endpoint create l3ep10 60.60.84.1/22 vlan-aware 610
sir endpoint port add l3ep10 3
sir endpoint admin l3ep10 enable

sir endpoint create l3ep11 60.60.88.1/22 vlan-aware 620
sir endpoint port add l3ep11 4
sir endpoint admin l3ep11 enable

sir endpoint create l3ep12 60.60.92.1/22 vlan-aware 630
sir endpoint port add l3ep12 5
sir endpoint admin l3ep12 enable

sir endpoint create l3ep13 60.60.96.1/22 vlan-aware 640
sir endpoint port add l3ep13 6
sir endpoint admin l3ep13 enable

sir endpoint create l3ep14 60.60.100.1/22 vlan-aware 650
sir endpoint port add l3ep14 7
sir endpoint admin l3ep14 enable

sir endpoint create l3ep15 60.60.104.1/22 vlan-aware 660
sir endpoint port add l3ep15 8
sir endpoint admin l3ep15 enable

sir endpoint create l3ep16 60.60.108.1/22 vlan-aware 699
sir endpoint port add l3ep16 3
sir endpoint admin l3ep16 enable
```

```
/*Default*/
sir route defaultgw add 10.10.62.2
/*->R1*/
sir route add 10.10.112.0/24 10.10.62.2
sir route add 10.10.111.0/24 10.10.62.2
sir route add 10.10.12.0/24 10.10.62.2
/*->R2*/
sir route add 10.10.222.0/24 10.10.62.2
sir route add 10.10.201.0/24 10.10.62.2
sir route add 10.10.202.0/24 10.10.62.2
/*->R3*/
sir route add 10.13.33.0/24 10.10.62.2
sir route add 10.13.31.0/24 10.10.62.2
sir route add 10.13.23.0/24 10.10.62.2
/*->R4*/
sir route add 10.10.41.0/24 10.10.62.2
sir route add 10.10.44.0/24 10.10.62.2
/*->R5*/
sir route add 10.10.51.0/24 10.10.56.1
sir route add 10.10.52.0/24 10.10.56.1
sir route add 10.10.54.0/24 10.10.56.1
/*->PseudoRouter*/
sir route add 60.70.60.0/22 60.60.60.2
sir route add 60.70.64.0/22 60.60.60.2
sir route add 10.100.4.0/22 60.60.68.10
sir route add 10.101.4.0/22 60.60.68.10
```

CLI Automation Scripts

Automatic background operations implemented using the Tera Term freeware tool. Script files are written in Tera Term macro language.

Start Background Script

1. Connect to router via Tera Term telnet.
2. On Tera Term click control/macro/
3. Double click script_file.ttl

Script Files used for Background Operations

TeraTerm script file r3.ttl

```

;-----
; on R3 one L3EPs on VLAN 500 (a VLAN not used otherwise for traffic tests) shall
; be continuously created and removed on the port facing R2, roughly every 2 minutes.
; This should trigger gratuitous ARPs, which should be visible on the R2 node
;-----
for i 1 100000 ;forever
sendln 'sir endpoint create l3ep4 10.20.31.1/24 vlan-aware 500'
wait '>'
sendln 'sir endpoint port add l3ep4 3'
wait '>'
sendln 'sir endpoint admin l3ep4 enable'
wait '>'
pause 120
sendln 'sir endpoint admin l3ep4 disable'
wait '>'
sendln 'sir endpoint delete l3ep4 '
wait '>'
pause 2
next

```

TeraTerm script file r2.ttl

```

;-----
;on R2 node shall have a L3EP on the same VLAN 500 not used for other purposes
;-----
sendln 'sir endpoint create l3ep6 10.20.31.2/24 vlan-aware 500'
wait '>'
sendln 'sir endpoint port add l3ep6 1'
wait '>'
sendln 'sir endpoint admin l3ep6 enable'
wait '>'

```

TeraTerm script file r1.ttl

```

;-----
; on R1 script shall create and delete 50 static routes towards R2. The routes
; will not be used for any particular routing purpose.
;-----
for k 1 100000 ;forever
;create 50 routes
for i 1 50
sprintf 'sir route add 10.10.%d.22/24 10.10.12.%d' i i
sendln inputstr
wait '>'
next
pause 2
;delete 50 routes
for i 1 50
sprintf 'sir route delete 10.10.%d.22/24 10.10.12.%d' i i
sendln inputstr
wait '>'
next
pause 2
next

```


TeraTerm script file r5.ttl

```
-----  
; on R5 every 10 min issue arp flush  
-----  
for i 1 100000 ;forever  
  sir arp flush  
  pause 600  
next
```

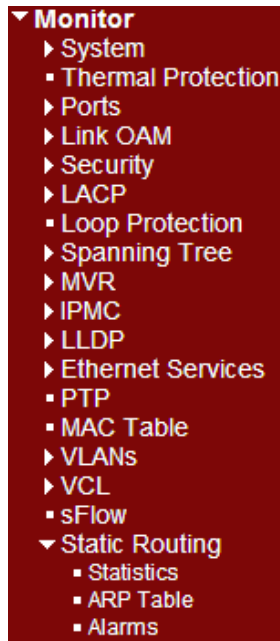
3. Web Interface

The SIR feature lets you define L3 routing endpoints over the VLANs configured in the NID. SIR can be configured via the CLI commands or via the web GUI. This section describes the web GUI screens used for Static IP Routing functionality.


The S3280 menus are updated for Static IP Routing to include the **Static Routing** item in the **Configuration**, **Monitor**, and **Diagnostics** menu paths. The Configuration, Monitor, and Diagnostics menu paths are shown below with the Static Routing dropdowns:



Configuration > Static Routing



Monitor > Static Routing



Diagnostics > Static Routing

The Static IP Routing web GUI menus and parameters are described in the following sections.

Configuration > Static Routing > Endpoints Configuration

The S3280 **Configuration > Static Routing > Endpoints** menu path lets you configure L3 static routing endpoints. (The S3280 **Monitor > Static Routing > Statistics** menu path lets you monitor L3 static Routing statistics for a specified Endpoint.)

L3 Static Routing Endpoints Configuration

This page allows for adding and deleting Static IP Routing Endpoints and changing its properties such as admin admin, etc. Up to sixteen L3 Endpoints are supported.

After you click the “Add L3 Endpoint” button, a new row is added to the default table where you can enter and change all configurable L3 Endpoint parameters.



The L3 Static Routing Endpoints Configuration table’s configurable parameters are explained below.

Delete: check the checkbox to delete an existing row from the table at the next Save. To delete an L3 Endpoint entry, select the checkbox. The entry will be deleted during the next Save. If the checkbox is disabled it means that Static Routing is configured through this L3 Endpoint and cannot be deleted.


L3 Endpoint: the configured name for the L3 endpoint (e.g., *L3EP1* above). Enter a unique L3 Endpoint identifier. The valid range is 1 to 16.

Admin: check the checkbox to enable Static IP Routing Endpoint administrative status. Routing operations are not performed if Admin status is disabled. To modify admin status, check appropriate checkbox. The modification will be applied during the next Save. By default, admin status is disabled, and for every new Endpoints entry the checkbox is unchecked.

IP Address: the IP address configured for the L3 endpoint (e.g., *10.100.5.20/24* above). The IP address / subnet size of L3 Endpoint, using the CIDR notation. The IP Mask length valid range is 0 – 32 characters.

MAC Address: Initially blank. Displays the current MAC address after a save is performed (e.g., *00-c0-f2-00-00-0c*).

VLAN Tag: the VLAN tag configured for the L3 endpoint (e.g., VLAN *unaware* above). The combo box lets you select VLAN Tag. The options are *untagged*, *C-Tag*, and *unaware*.



VLAN ID: the VLAN ID configured for the L3 endpoint (e.g., VID *0* above). Indicates the VLAN ID for this L3 Endpoint. For VLAN Tag set to *untagged* or *unaware*, the VLAN ID should be *0*. The valid range is 1-4094. **Note:** You can not set the L3 endpoint VLAN equal to the existing Management VLAN ID.

Port Members: checkboxes identifying which ports are included in each L3 endpoint configuration. A row of check boxes for each port is displayed for each L3 Endpoint. To include a port in L3 Endpoint, check the box. To remove or exclude the port from the L3 Endpoint, make sure the box is unchecked. By default, no ports are members, and for every new Endpoint entry all boxes are unchecked.

Proxy ARP Enabled: checkbox identifying if the configured L3 endpoint is enabled (checked) for Proxy ARP administrative status. Proxy ARP allows the NID to respond to an ARP request for a particular host with its own MAC address, if it knows how to forward frames for the real destination host. To modify admin status, check the appropriate checkbox. The modification will be applied during the next Save. By default, Proxy ARP admin status is enabled, and for every new Endpoints entry the checkbox is checked.

Proxy ARP Group: displays the Proxy ARP (address resolution protocol) configured group number (e.g., Groups 1 and 2 above). The valid range is 1 to 16. By default, the Proxy ARP Group number is 1 for all L3 Endpoints.

Split Horizon - Enabled: checkbox identifying if Split Horizon administrative status is configured for the L3 endpoint (checked). By default, Split Horizon is disabled for all interfaces, which means that the Endpoints are part of the same horizon. To modify Split Horizon, check the appropriate checkbox. The modification will be applied during the next Save.

Split Horizon - Horizon: Horizon for L3 Endpoint. To modify Horizon, select new value from combo box. The modification will be applied during the next Save. This is the direction (uplink or downlink) using the split horizon method of preventing a routing loop in a network by ensuring that information about the routing for a particular packet is never sent back in the direction from which it was received. This is a facility of a Domain Name System (DNS) implementation that provides different sets of DNS information that are typically selected by the source address of the DNS request. Split horizon provides a mechanism for security and privacy management by logical or physical separation of DNS information for network-internal access (within an administrative domain, e.g., a company) and access from an insecure, public network (e.g., the Internet).

Adding a New L3 Endpoint

Click the **Add L3 Endpoint** button to add a new L3 Endpoint. An empty row is added to the table, where all parameters can be configured as needed.

The **Delete** button can be used to undo the addition of new L3 Endpoint.

Buttons

Add L3 Endpoint: Click to add a new L3 endpoint to the table and configure it. An empty row is added to the table, where all parameters can be configured as needed.

Refresh: Click to refresh the page immediately. Note that unsaved changes will be lost.

Save: Click to save the current settings / changes.

Reset: Sets the settings back to the defaults. Click to undo any changes made locally and revert to previously saved values.

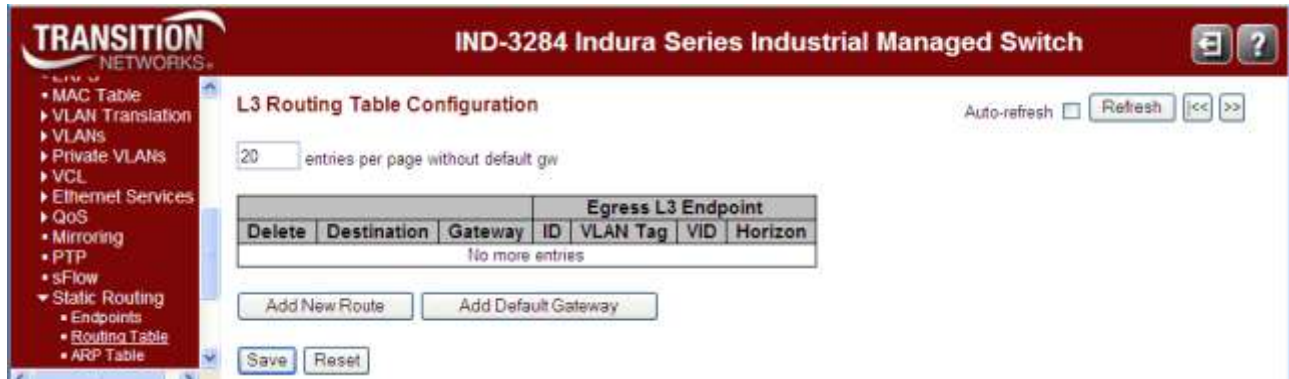
Delete: Can be used to undo the addition of new L3 Endpoint.

Configuration > Static Routing > Routing Table

The S3280 **Configuration > Static Routing > Routing Table** menu path lets you add and configure new routes and gateways in terms of the route destination and gateway and L3 endpoint ID, VID and horizon.

L3 Routing Table Configuration

This page allows for configuring Static IP Routing Table. The default L3 Routing Table is shown below.



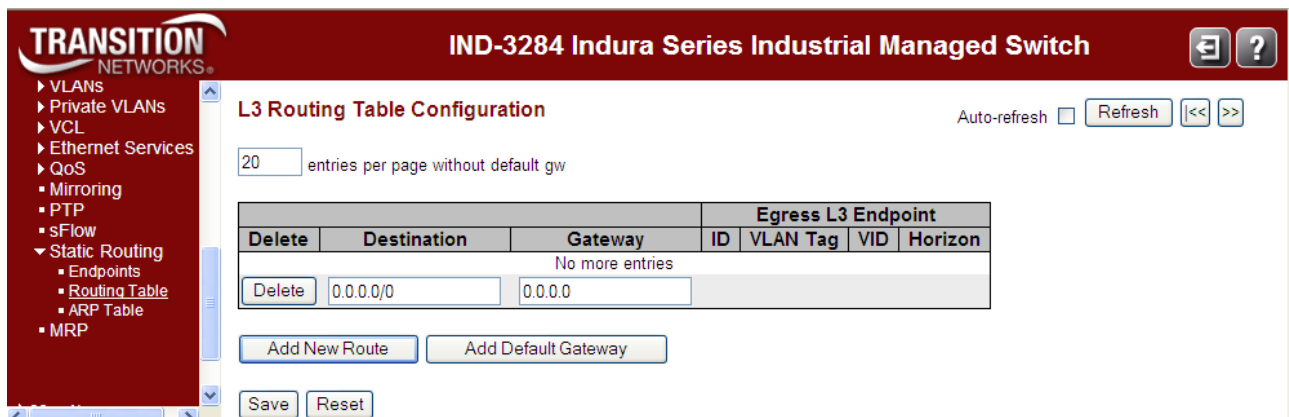
Each page shows up to 99 entries from the Static IP Routing table (the default is 20 entries) selected through the "entries per page" input field. When first visited, the web page will show the first 20 entries from the beginning of the Static IP Routing Table.

Clicking the **Refresh** button will update the displayed table. The >> button will use the last entry of the currently displayed Static IP Routing entry as a basis for the next lookup. When the end is reached the text "No more entries" is shown in the displayed table. Use the << button to start over.

Click the "Add New Route" or the "Add Default Gateway" button to add a new row (instance) to the current table; here you can enter or change all L3 Routing table configurable parameters.

Add New Route

Click the **Add New Route** button to add a new Static Route. An empty row is added to the table, where all parameters can be configured as needed.



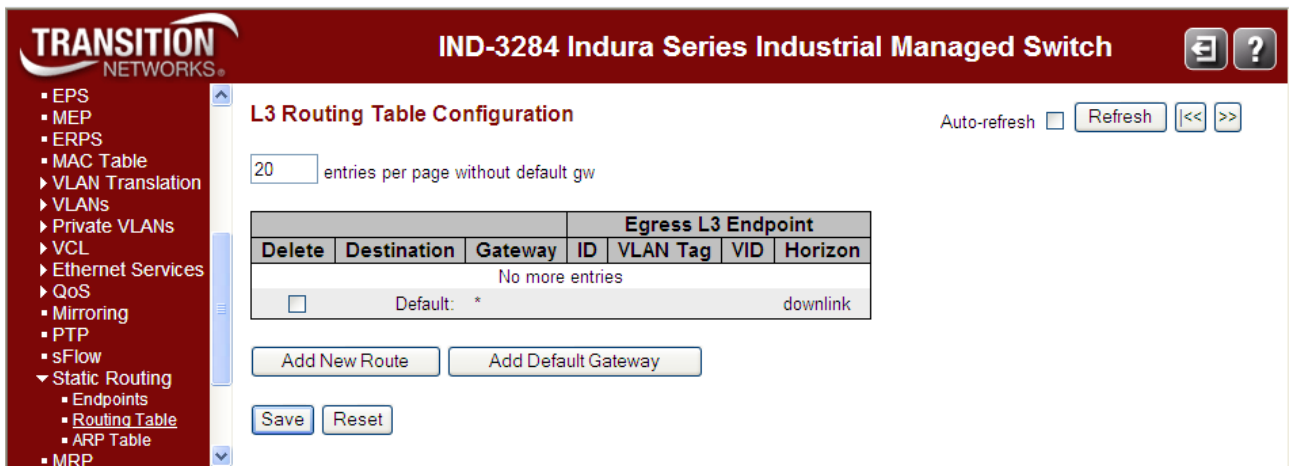
Up to 128 Static Routes are supported. You can use the **Delete** button to undo the addition of new Static Route.

Add Default Gateway

Select the **Configuration > Static Routing > Routing Table** menu path to display the default L3 Routing Table Configuration page.



Click the **Add Default Gateway** button to add a new Default gateway. An empty row is added to the table, where all parameters can be configured as needed.



Up to 2 Default Gateways are supported. You can use the **Delete** button to undo the addition of a new Default Gateway.

The L3 Routing Table parameters are explained below.

Delete: Check the checkbox to delete the instance at the next Save. The instance can not be deleted if the checkbox is grayed out. To delete Route entry, select the checkbox. The entry will be deleted during the next Save. If the checkbox is disabled it means that this route is the 'connected route' and cannot be deleted.

Destination: The destination IP address/subnet size, using the CIDR notation. Specify either an IP address or host name for the destination network or host (e.g., *10.100.5.0/24* above).

Gateway: The IP address. The gateway must belong to one of the subnets directly connected to the NID. Specify either an IP address or host name for the gateway or router to use when forwarding (e.g., * or *10.110.5.100* above).

Egress L3 Endpoint ID: a unique identifier for this L3 endpoint (e.g., *L3EP1* above).

Egress L3 Endpoint VID: a unique VLAN ID for this L3 endpoint (e.g., *100* above)

Egress L3 Endpoint Horizon: the direction for this L3 endpoint (either *downlink* or *uplink*).

Buttons

entries per page without default gw: Sets the maximum number of instances that can be displayed on a single page (e.g., *20* above).

Add New Route: Click to display the entry table to enter a new row (instance) to the Routing table and define its parameter settings.

Add Default Gateway: Click to add a new row (instance) to the current table; here you can change or enter Gateway configurable parameters.

Save: Click to save the current settings.

Reset: Click to set the settings back to the defaults.

Auto-refresh: Click to automatically refresh the page every 3 seconds.

Refresh: Click to automatically refresh the page immediately. Clicking the **Refresh** button will update the displayed table.

|<<: Updates the table starting from the first entry in the MVR Channel Configuration for a specific MVR VLAN. Use the |<< button to start over.

>>: Updates the table, starting with the entry after the last entry currently displayed. Clicking the >> button will use the last entry of the currently displayed Static IP Routing entry as a basis for the next lookup. When the end is reached the text "No more entries" is shown in the displayed table.

Messages

Message: *Routing Table Configuration failed - invalid IP parameter value*

Meaning: You tried to save with an invalid IP address configured.

Recovery: Click the browser's back button and retry the operation with a different IP address.

Message:

Routing Table Configuration failed - Adding of Route is failed, Route already exists

Routing Table Configuration failed - Adding of Route is failed, Static route match existing on the same horizon

Routing Table Configuration failed - Adding of Route is failed, Gateway not on connected net

Meaning: You tried to save with two endpoints having the same configuration.

Recovery: Click the browser's back button and retry the operation with a different (unique) static route.

Message: L3 Static Routing Endpoints Configuration failed - Creating of L3 Endpoint is failed, Subnet overlaps with other L3 endpoint

Meaning: You tried to save with two endpoints having the same configuration.

Recovery: Click the browser's back button and retry the operation with a different (unique) static route.

Message: Message from webpage - L3EP ID: 1 is already used

Example: You tried to save with two endpoints having the same Egress L3 Endpoint ID.



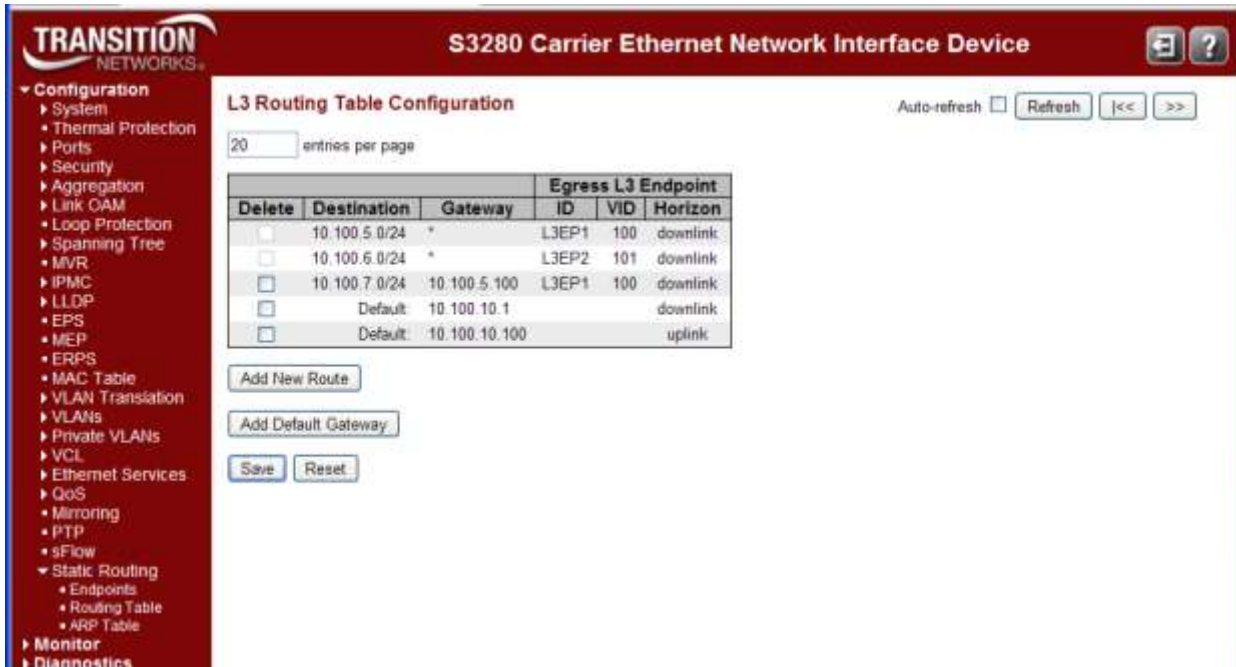
Meaning:

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Retry the operation with a different (unique) Egress L3 Endpoint ID. Refer to the preceding section.

Routing Table Config Example

The screen below shows several destination and gateway configurations.



Note that you can not add a new Static route that matches an existing Static route on the same horizon.

Configuration > Static Routing > ARP Table

The S3280 **Configuration > Static Routing > ARP Table** menu path lets you define the ARP Aging timeout, Wait reply timeout, and Refresh timeout values. This page also lets you flush (reset to defaults) the table values for a selected SIR instance or All SIR instances.

The ARP (Address Resolution Protocol) protocol is used to convert an IP address into a physical address, such as an Ethernet address. ARP allows a host to communicate with other hosts when only the Internet address of its neighbors is known. Before using IP, the host sends a broadcast ARP request containing the Internet address of the desired destination system.

L3 Static Routing ARP Config

This page allows for configuring Static IP Routing ARP parameters. The default L3 Static Routing ARP Config table is shown below. Note that if no L3 end points have been created, the message “No created L3 Endpoints” displays.

The L3 Static Routing ARP Config table is described below.

Aging Timeout (sec): The amount of time in seconds to maintain an instance in the table. The valid range is 1-1500 seconds (e.g., 60 shown above). This is the amount of time to keep unused ARP record in the table. The valid range is 60 to 86400 seconds. By default, the aging timeout value is 120 seconds.

Wait Reply Timeout (sec): The amount of time in seconds to wait for a reply before timing out. This is the amount of time between ARP request send and declaration of failed. The valid range is 1 to 10 seconds (e.g., 60 seconds shown above). By default, the wait reply timeout value is 2 seconds.

Refresh Timeout (sec): The amount of time in seconds to wait for a refresh. This is the amount of time between consequential refreshing ARP requests sent for the record in the table. The valid range is 60 to 86400 seconds (e.g., 300 shown above). By default, refresh timeout value is 60 seconds.


GARP enabled: The status of Gratuitous ARP. If disabled all received GARP messages are ignored and no GARP messages are generated by the system. By default, Gratuitous ARP is enabled (checkbox checked). You can uncheck the checkbox to disable Gratuitous ARP.

Flushing of ARP Table

Select the L3 Static Routing Endpoint from the combobox () to remove ARP entries for an endpoint, or click the 'All' checkbox to clear the entire ARP Table.

Buttons

Save: Click to save the currently displayed settings.



Selection dropdown: (e.g., **L3EP1** for Layer 3 End Point 1).

All: Check the **All** checkbox and click the **Flush** button to flush (reset) all instances to default values.

Flush: Click to flush a selected instance or all instances. Select the L3 Static Routing Endpoint from the combobox to remove ARP entries for a specific endpoint, or click the 'All' checkbox to clear all ARP table endpoints.

Monitor > Static Routing > Statistics

The S3280 **Monitor > Static Routing > Statistics** menu path lets you view L3 static Routing statistics for a specified endpoint. (The **Configuration > Static Routing > Statistics** menu path lets you configure L3 static Routing statistics for a specified endpoint.)

L3 Static Routing Statistics

This page displays the current static routing statistics collected for a specified Endpoint (e.g., L3EP1).

TRANSITION NETWORKS S3280 Carrier Ethernet Network Interface Device

L3 Static Routing Statistics Auto-refresh Refresh Clear

L3EP1

Receive Total		Transmit Total	
Rx Packets	0	Tx Packets	0
Rx Octets	0	Tx Octets	0
Rx Multicast	0	Tx Multicast	0
Rx Broadcast	0	Tx Broadcast	0
Rx Unicast	0	Tx Unicast	0
Rx Forw Datagrams	0	Tx Forw Datagrams	0
		Tx Requests	0
Receive Errors		Transmit Errors	
Rx Discards	0	Tx Discards	0
Rx No Routes	0		
Rx Unknown Protos	0		
Rx Addr Errors	0		
Rx Header Errors	0		
Receive ICMP		Transmit ICMP	
Rx Packets	0	Tx Packets	0
Rx Discards	0		

The L3 Static Routing Endpoints statistics fields (Counters) are explained below.

Receive Total

- Rx Packets*: the total number of input IP datagrams received, including those received in error.
- Rx Octets*: the total number of octets received in input IP datagrams, including those received in error.
- Rx Multicast Packets*: the number of IP multicast datagrams received.
- Rx Broadcast Packets*: the number of IP broadcast datagrams received.
- Rx Forw(ard) Datagrams*: the number of input datagrams for which this entity was not their final IP destination and for which this entity attempted to find a route to forward them to that final destination.

Receive Errors

- Rx Discards*: the number of input IP datagrams for which no problems were encountered to prevent their continued processing, but were discarded (e.g., for lack of buffer space).
- Rx No Routes*: The number of input IP datagrams discarded because no route could be found to transmit them to their destination.
- Rx Address Errors*: the number of input IP datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (e.g., ::0).
- Rx Unknown Protos*: the number of locally-addressed IP datagrams received successfully but discarded because of an unknown or unsupported protocol.
- Rx Header Errors*: the number of input IP datagrams discarded due to errors in their IP headers, including version number mismatch, other format errors, hop count exceeded, errors discovered in processing their IP options, etc.

Transmit Total

Tx Packets: the total number of IP datagrams that this entity supplied to the lower layers for transmission. This includes datagrams generated locally and those forwarded by this entity.

Tx Octets: the total number of octets in IP datagrams delivered to the lower layers for transmission.

Tx Multicast Packets: the number of IP multicast datagrams transmitted.

Tx Broadcast Packets: the number of IP broadcast datagrams transmitted.

Tx Forw(ard) Datagrams: the number of datagrams for which this entity was not their final IP destination and for which it was successful in finding a path to their final destination.

Tx Requests: The total number of IP datagrams that local IP user-protocols (including ICMP) supplied to IP in requests for transmission.

Transmit Errors

Tx Discards: the number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but were discarded (e.g., for lack of buffer space).

Receive ICMP

Rx (Icmp) Packets: the total number of ICMP datagrams received.

Rx (Icmp) Discards: the total number of discarded ICMP datagrams.

Transmit ICMP

Tx (Icmp) Packets: the total number of ICMP datagrams transmitted.

Buttons

Selection dropdown: (e.g., **L3EP1** for Layer 3 End Point 1) (),

Auto-refresh: Click to automatically refresh the page every 3 seconds.

Refresh: Click to automatically refresh the page immediately.

Clear: Clears (zeros out) the counters on the page.

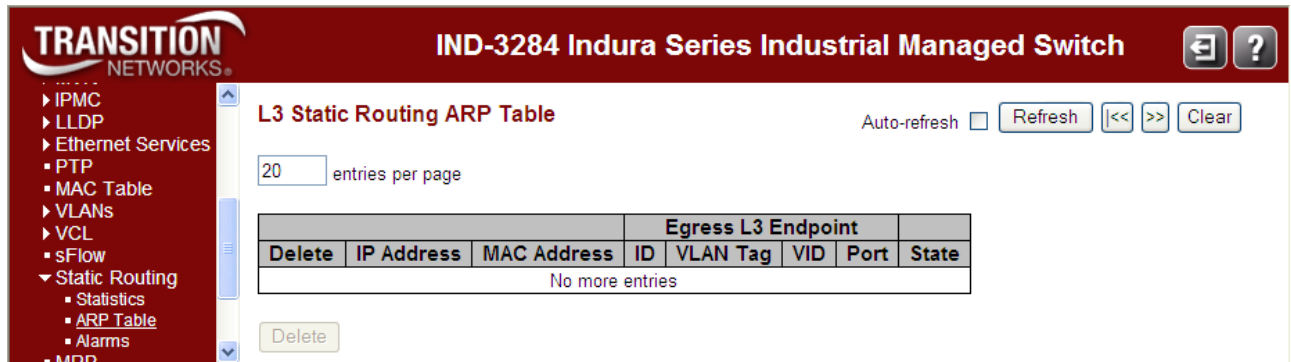
Monitor > Static Routing > ARP Table

The **Configuration > Static Routing > ARP Table** menu path displays the L3 Static Routing ARP Table information (read only), and lets you delete a selected ARP table instance.

L3 Static Routing ARP Table

The default L3 Routing ARP Table is shown below. Each page shows up to 99 entries from the Static IP Routing ARP table, default being 20, selected through the "**entries per page**" input field. When first visited, the web page will show the first 20 entries from the beginning of the Static IP Routing ARP Table.

Clicking the **Refresh** button will update the displayed table. The **>>** button will use the last entry of the currently displayed Static IP Routing entry as a basis for the next lookup. When the end is reached the text "*No more entries*" is shown in the displayed table. Use the **<<** button to start over.



The L3 Static Routing ARP Table parameters are explained below.

Delete: Check the checkbox to delete the ARP entry (nstance) at the next Save.

IP Address: The IP address assigned to the Ethernet management interface; it must be within the same subnet and subnet mask as the IP address on the router for this Ethernet connection (e.g., 10.100.5.11 above).

MAC Address: the Media Access Control address (e.g., 00.01.02.03 .04.aa) of the Ethernet management interface. The MAC address corresponding to the IP address.

Egress L3 Endpoint ID: The configured Egress L3 Endpoint VLAN (e.g., L3EP1).

Egress L3 Endpoint VLAN Tag: The Egress L3 Endpoint VLAN Tag type configured (i.e., *unaware*, *untagged*, or *C-Tag*).

Egress L3 Endpoint VID: The configured VLAN ID (e.g., VID 100) for this Egress L3 Endpoint.

Egress L3 Endpoint Port: The configured port for this for this Egress L3 Endpoint.

State: The ARP table entries resolving status.

References: The count of routing cache entries depending from the record.

Buttons

entries per page: Sets the maximum number of instances that can be displayed on a single page (e.g., 20 above).

Auto-refresh: Click to automatically refresh the page every three seconds.

Refresh: Click to automatically refresh the page immediately.

|<<: Updates the table starting from the first entry in the MVR Channel Configuration for a specific MVR VLAN.

>>: Updates the table, starting with the entry after the last entry currently displayed.

Clear: Clears (zeros out) the counters on the page.

Delete: The button can be used to undo the addition of a new L3 Endpoint.

Monitor > Static Routing > Alarms

The **Monitor > Static Routing > Alarms** menu path lets you view any existing Endpoint alarm indications.

The L3 Static Routing Endpoints Alarms table is shown below.

The screenshot shows the configuration page for the S3280 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with the following items: MAC Table, VLANs, VCL, sFlow, Static Routing (expanded), Statistics, ARP Table, and Alarms. The main content area is titled "L3 Static Routing Endpoints Alarms" and includes an "Auto-refresh" checkbox (unchecked) and a "Refresh" button. Below this is a table with two columns: "L3 Endpoint" and "Alarms".

L3 Endpoint	Alarms
L3EP1	No alarms

The screenshot shows the configuration page for the IND-3284 Indura Series Industrial Managed Switch. The left sidebar contains a navigation menu with the following items: VCL, sFlow, Static Routing (expanded), Statistics, ARP Table, Alarms, and MRP. The main content area is titled "L3 Static Routing Endpoints Alarms" and includes an "Auto-refresh" checkbox (unchecked) and a "Refresh" button. Below this is a table with two columns: "L3 Endpoint" and "Alarms".

L3 Endpoint	Alarms
L3EP1	No alarms

The L3 Static Routing Endpoints Alarms table parameters are explained below.

L3 Endpoint: (e.g., *L3EP1* shown above).

Alarms: (e.g., *No alarms* if no alarm conditions exist).

Buttons

Auto-refresh: Click to automatically refresh the page every three seconds.

Refresh: Click to refresh the page immediately.

Diagnostics > Static Routing Utilities

The **Diagnostics > Static Routing** menu path provides the Ping and Trace network utilities for L3 Static Routing.

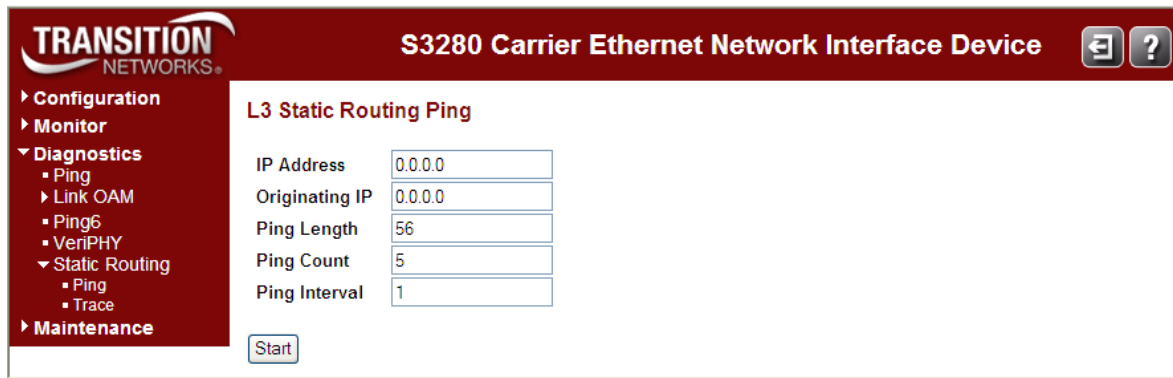
Diagnostics > Static Routing > Ping

This page lets you issue ICMP PING packets to troubleshoot IP connectivity issues.

After you press the **Start** button, ICMP packets are transmitted, and the sequence number and round trip time are displayed upon reception of a reply. The page refreshes automatically until responses to all packets are received, or until a timeout occurs. For example:

```
PING server 10.10.101.100, 56 bytes of data.
64 bytes from 10.10.101.100: icmp_seq=0, time=0ms
64 bytes from 10.10.101.100: icmp_seq=1, time=0ms
64 bytes from 10.10.101.100: icmp_seq=2, time=0ms
64 bytes from 10.10.101.100: icmp_seq=3, time=0ms
64 bytes from 10.10.101.100: icmp_seq=4, time=0ms
Sent 5 packets, received 5 OK, 0 bad
```

The L3 Static Routing Ping Table is shown below.



The L3 Static Routing Ping table parameters are explained below.

IP Address: The destination IP Address. Enter a valid IPv4 address (e.g., *10.10.101.1*). The default is *0.0.0.0*.

Originating IP: If more than one Egress L3 Endpoint is configured for a given destination IP address, you must specify this parameter. The default is *0.0.0.0*.

Ping Length: The payload size of the ICMP packet. Valid values range from 2 bytes to 1452 bytes. Enter the packet size in bytes (e.g., *56*).

Ping Count: The count of the ICMP packet. Valid values range from 1 ping to 32 pings. Enter the number of pings to be sent. The default is *5* pings.

Ping Interval: The interval of the ICMP packet. Valid values range from 0 second to 30 seconds. Enter the interval to be inserted between pings. The default is *1* seconds.

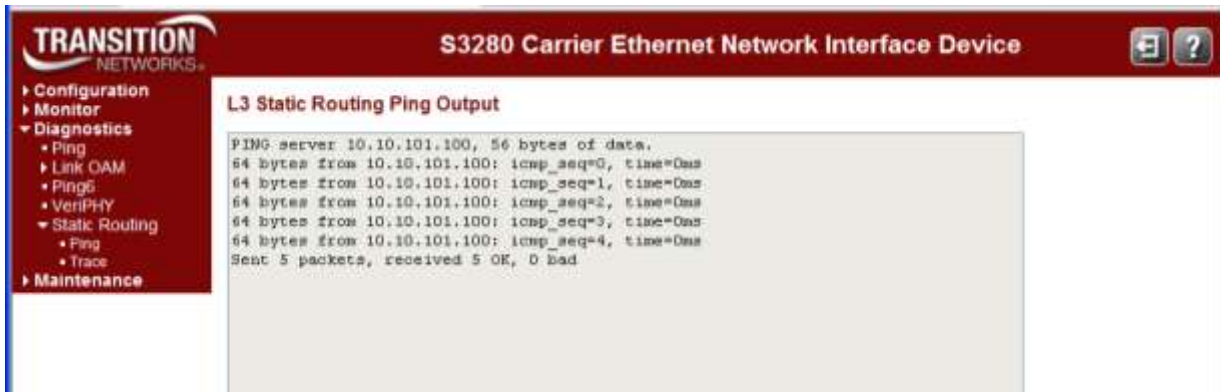
Buttons

Start: Click to start the ping process (i.e., start transmitting ICMP packets). When done, the ping output displays.

New Ping: Click to re-start diagnostics with PING. The **New Ping** button only displays after an initial Ping was performed.

Ping Output

A sample L3 Static Routing ping output is shown below.



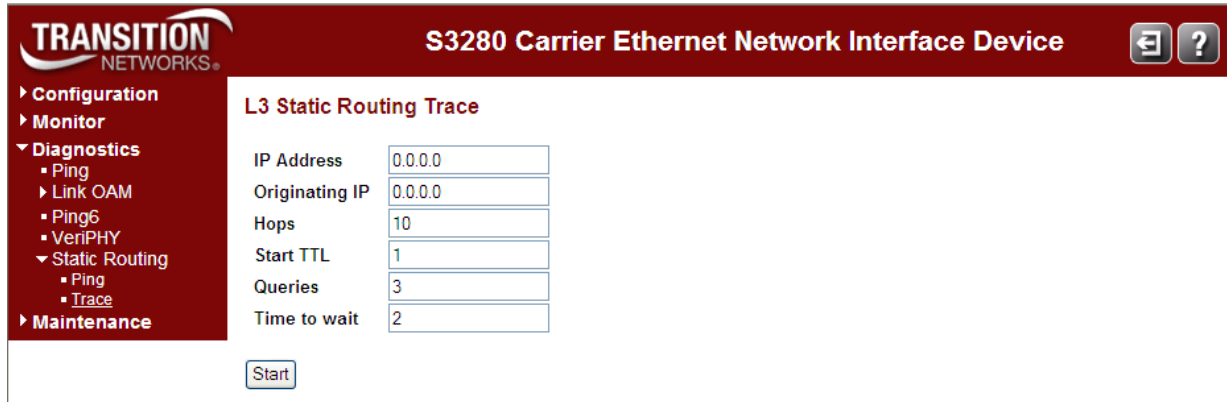
Diagnostics > Static Routing > Trace

Send Trace Parameters

The **Diagnostics > Static Routing > Trace** menu path lets you configure and send L3 static routing trace parameters. This page lets you issue ICMP traceroute packets to troubleshoot IP connectivity issues.

After you press the **Start** button, ICMP packets are transmitted, and the sequence number and round trip time are displayed upon reception of a reply. The page refreshes automatically until responses to all packets are received, or until a timeout occurs. For example:

```
TRACE to 10.10.101.100, 5 hops max, start from TTL 1.
1:      0.0ms      0.0ms      0.0ms
2:      0.0ms      0.0ms      0.0ms
3:      0.0ms      0.0ms      0.0ms
4:      0.0ms      0.0ms      0.0ms
5:      0.0ms      0.0ms      0.0ms
Resume: reached, hops 5
```



IP Address: Enter the destination IP Address. The IPv4 address in CIDR notation (a.b.c.d/m) (e.g., 10.10.101.1).

Originating IP: If more than one Egress L3 Endpoint is configured for given destination IP address, you must specify this parameter. The default is 0.0.0.0.

Hops: Enter the maximum number of traced hops. Enter the MAX number of traced hops. The default is 10 hops traced. Valid values range from 1-32.

Start TTL: Enter the Transmit TTL value of the first packet. The default is 1 second. Enter the “Start Time To Live” in seconds (e.g., 1 above). Valid values are 1-255.

Queries: Enter the Transmit packet tries. The default is 3 queries (e.g., 1 above). Transmit packet tries. Valid values are 1-10 tries (query attempts).

Time to wait: Enter the Reply waiting time. The default is 2 seconds (e.g., 1 above). Valid values are 1-10 seconds.

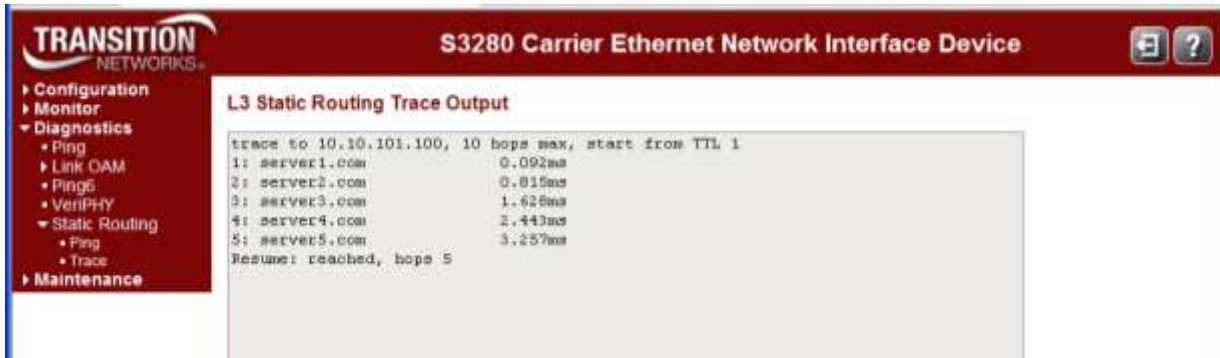
Buttons

Start: Click to start the Trace process (i.e., to start transmitting ICMP packets). When done, the trace output displays.

New Trace: Click to re-start diagnostics with traceroute. The **New Trace** button only displays after an initial Trace was performed.

Trace Output

A sample L3 Static Routing Trace output is shown below.



The screenshot shows the web interface of a Transition Networks S3280 Carrier Ethernet Network Interface Device. The interface has a dark red header with the logo and title. A left-hand navigation menu is visible, with 'Diagnostics' expanded to show 'Static Routing' and 'Trace' selected. The main content area displays the following trace output:

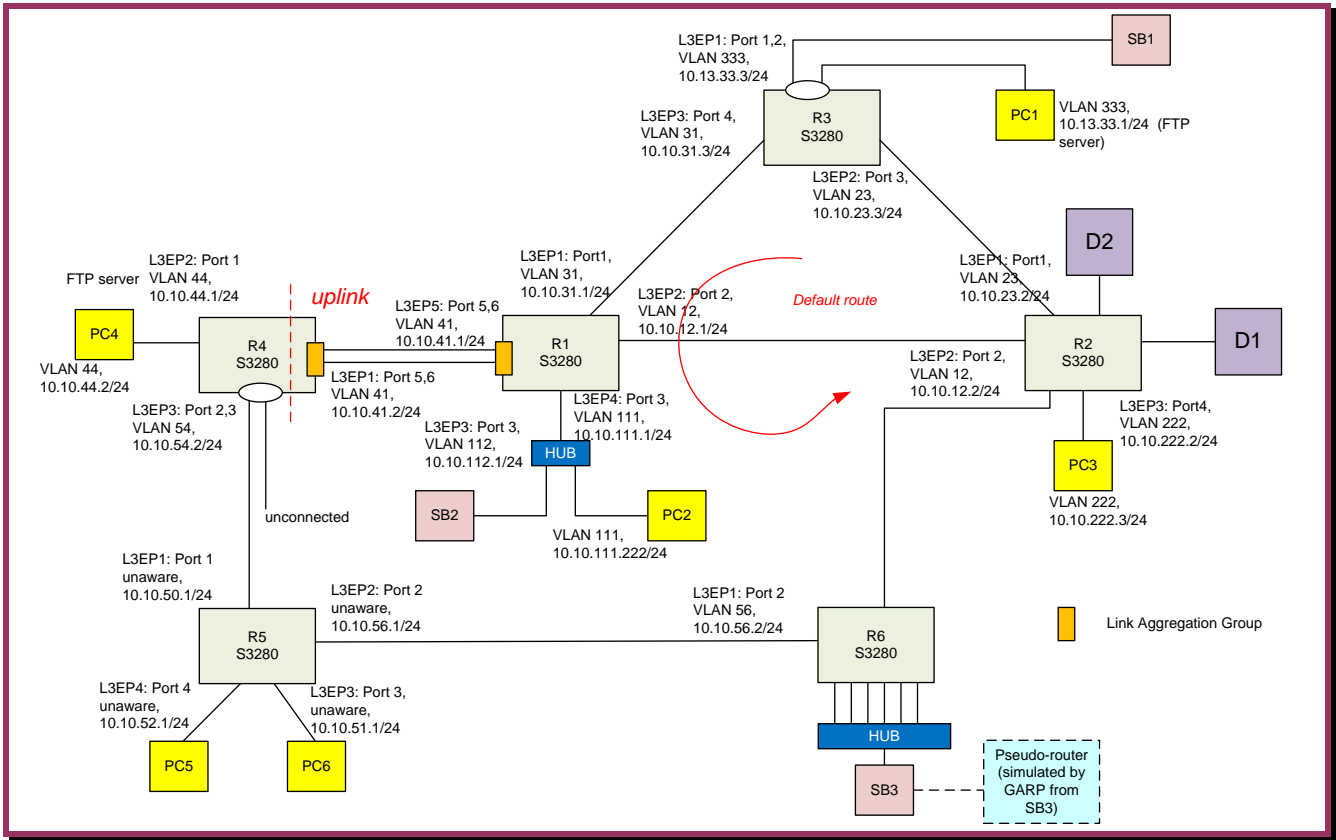
```
L3 Static Routing Trace Output

trace to 10.10.101.100, 10 hops max, start from TTL 1
1: server1.com          0.092ms
2: server2.com          0.815ms
3: server3.com          1.628ms
4: server4.com          2.443ms
5: server5.com          3.257ms
Resume: reached, hops 5
```

Note that the static route trace to IP address *10.10.101.100* was reached in 5 hops in the example above.

SIR Configuration Example (via Web GUI)

The configuration example below shows an SIR configuration.



1. Navigate to the **Configuration > Static Routing** menu path.
 - a. Select the **Endpoints** menu item. Configure the L3 Static Routing Endpoints.
 - b. Select the **Routing Table** menu item. Configure the L3 Routing Table (“New Route” and “Default Gateway” parameters).
 - c. Select the **ARP Table** menu item. Configure the L3 Static Routing ARP parameters.
2. Navigate to the **Monitor > Static Routing** menu path.
 - a. Select the **Statistics** menu item. View the L3 Static Routing Statistics.
 - b. Select the **ARP Table** menu item. View the L3 Static Routing ARP Table.
 - c. Select the **Alarms** menu item. View the L3 Static Routing Endpoints Alarms.
3. Navigate to the **Diagnostics > Static Routing** menu path.
 - a. Select the **Ping** menu item. Configure and run the L3 Static Routing Ping.
 - b. Select the **Trace** menu item. Configure and run the L3 Static Routing Trace.

4. Messages and Troubleshooting

Static IP Routing Messages

This section provides general and specific Static IP Routing problem solving suggestions, general error recovery steps, and specific web interface messages, meanings, examples, and possible recovery steps.

For functional level troubleshooting for S3280 web interface operation, see the “Web Interface Messages” section of the S3280 Web Interface User Guide, 33506. Topics there include EPS Troubleshooting, ERPS Troubleshooting, IPv6 Troubleshooting, Troubleshooting High CPU Load Conditions, etc.

CLI Messages and Web GUI Messages

Message:

W tn_sip 01:51:31 61/tn_sip_web_routing_table_handler#747:

Warning: invalid IP parameter value

Meaning: You entered an invalid configuration parameter.

Recovery:

1. Change the SIR IP address.
2. Continue operation.
3. See the related manual or section for more information.

Message:

W tn_sip 02:18:26 61/tn_sip_web_routing_table_handler#793:

Warning: Adding of Route is failed, Static route match existing on the same horizon

Meaning: You tried to configure an L3 route, but it has not yet been created.

Recovery:

1. Make sure the route is created and exists in the table.
2. Continue operation.
3. See the related manual or section for more information.

Message:

W tn_sip/l3 03:51:09 70/tn_sip_routing_endpoint_add#534:

Warning: l3 match mgmtvlan

Meaning: You entered an invalid configuration parameter.

1. Change the L3 Endpoint VLAN ID.
2. Continue operation.
3. See the related manual or section for more information.

Message:

W tn_sip 03:43:46 61/tn_sip_web_endpoints_handler#529:

Warning: Creating of L3 Endpoint is failed, L3 endpoint VLAN equal to management VLAN

>W tn_sip/l3 00:35:56 70/tn_sip_routing_endpoint_add#537: Warning: l3 match mgmt vlan

Meaning: You entered an invalid configuration parameter.

Recovery:

1. Change the L3 Endpoint VLAN ID.
2. Continue operation.
3. See the related manual or section for more information.

Message:

W tn_sip/l3 03:49:17 70/tn_sip_routing_endpoint_add#577:

Warning: l3 sub overlap i=1

Meaning: You entered an invalid configuration parameter.

Recovery:

1. Change the L3 Static Routing Endpoints subnet configuration.
2. Continue operation.
3. See the related manual or section for more information.

Message:

W tn_sip 00:36:20 61/tn_sip_web_endpoints_handler#589: Warning: Enabling of L3 Endpoint is failed,

Cannot admin up L3 endpoint with no ports assigned to it

Meaning: You tried to enable an L3 endpoint configuration, but you have not first created the required port assignment.

Recovery:

1. Create the port configuration (e.g., via the “*SIR Endpoint Port Add*” command or “*Port Members*” parameter in the L3 Static Routing Endpoints Configuration table via the **Configuration > Static Routing > Endpoints** menu path).
2. Continue operation.
3. See the related manual or section for more information.

Message:

L3 Static Routing Endpoints Configuration failed

Modifying of L3 Endpoint is failed, Vlan type not match L2 port type

Meaning: You entered an invalid VLAN configuration parameter.

Recovery:

1. Verify the L3 Static Routing Endpoints VLAN Tag (*untagged*, *C-Tag*, or *unaware*) configuration.
2. Continue operation.
3. See the related manual or section for more information.


Message: L3EP ID: xis already used

Meaning: You entered an invalid L3EP ID parameter (e.g., at **Configuration > Static Routing > Endpoints**).

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Verify the “*L3 Static Routing Endpoints Configuration*” table parameter or “*SIR Endpoint*” command parameter.
3. Continue operation.
4. See the related manual or section for more information.

Message: IP Address must be in CIDR notation (a.b.c.d/m)



Meaning: You entered an invalid L3EP ID parameter (e.g., at **Configuration > Static Routing > Endpoints**).

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Verify the “L3 Static Routing Endpoints Configuration” table parameter or “SIR Endpoint” command parameter.
3. Continue operation.
4. See the related manual or section for more information.

Message:

```
E tn_sip 00:00:31 70/tn_sip_icmp_dropped_packet_process#521: Error: No route for reply
E api/fdma_irq 31.957,842 0/xtr_connect_and_restart#492: Error: Assert failed: ch_state->cur_tail &&
ch_state->cur_tail->next == NULL
```

Meaning: S3280 hangs when the PC is unplugged from the Management port and plugged into other port. With certain configurations, if a PC that is plugged into port 9 gets unplugged from port 9 and then plugged into another port, the system load gets stuck at 100%, IP management is lost, and errors are reported on the CLI. The Load average: (100ms, 1s, 10s): 100%, 99%, and 76%.

Recovery:

1. Upgrade to the latest S3280 software version.
2. Contact TN Tech Support.

For More Error Message Information

For specific messages resulting from S3280 web interface operation, see the “Web Interface Messages” section of the S3280 Web Interface User Guide, 33506.

For specific messages resulting from S3280 CLI operation, see the “CLI Messages” section of the S3280 CLI Reference Guide, 33507.

Static IP Routing (SIR) Troubleshooting

If you have problems trying to configure SIR on the S3280-TST, first check for:

- Incorrect IP address or subnet mask (interfaces on the same network must have IP addresses within the same subnet to be able to communicate directly). The IP address must be within the same subnet and subnet mask as the IP address on the router for this Ethernet connection.
- Duplicate IP addresses (IP addresses must be unique; the network portion can be the same for all IP addresses, but the host portion must be unique).
- Incorrect or missing default gateway address: add or correct as necessary.
- Incorrect or missing static routes; configure the destination IP address in the static route to point to the IP address assigned to the workstation PC. The static route is then automatically provided to the other nodes.

Problem: Ingress traffic is dropped on port at maximum rate.

Meaning: PTP incompatibility is causing ingress drops at maximum rate.

Recovery:

1. Assign different priorities to PTP and UDP traffic and enable tag QoS on the devices (e.g., PTP PCP=0 and UDP PCP=5).
2. Enable tag QoS with the CLI command ***Qos port class tag [<port_list>] enable.***

Example:

1. In SB1, define GARP message on VLAN 33 with self IP address 10.10.33.2, mode is Single Burst and count is one.
2. In R3, reconfigure ARP aging timeout to set it to max value 86400sec with the command ***sir arp configuration age 86400.***
3. Send GARP from SB1.
4. In SB2, define GARP message on VLAN 112 with self IP address 10.10.112.222, mode is Single Burst and count is one.
5. In R1, reconfigure ARP aging timeout to set it to max value 86400sec with the command: ***sir arp configuration age 86400.***
6. Send GARP from SB2;
7. In SB1 define UDP traffic on VLAN 33 with destination IP address 10.10.112.222.
8. Send UDP traffic from SB1 with rate ~900Mbps.
9. In SB2 define UDP traffic on VLAN 112 with destination IP address 10.10.33.2.
10. Send UDP traffic from SB2 with rate ~900Mbps.
11. Check statistics in R3 on L3EP1 and L3EP3.

Result: After approx. 1 minute all received frames on R3 port4 are dropped, L3EP3 receive counter stops.

5. SNMP MIBs and Traps

Supported MIBs

This section describes the SNMP MIBs used to implement the Static IP Routing features on the S3280 device, including standard MIBs supported by the feature as well as proprietary addendums to standard MIBs.

Standard MIBs

Standard MIB entities supported are:

- 1) IF-MIB (Per RFC 2863)
 - a. ifTable;
 - b. ifXTable;
- 2) IP-MIB (Per RFC 4293)
 - a. ipv4InterfaceTable;
 - b. ipIfStatsTable;
 - c. ipAddressTable;
 - d. ipDefaultRouterTable;
 - e. icmpStatsTable;
 - f. icmpMsgStatsTable;
- 3) FORWARD-MIB (Per RFC 4292)
 - a. inetCidrRouteTable;

Interface indexes (ifIndex) for routing endpoint are allocated according to current Transition SNMP agent implementation. Start index for routing endpoint *ifIndexes* is defined as 80000.

Proprietary MIBs

There are three major areas that need description in the proprietary MIBs:

- 1) An addendum to the IP interface, which allows configuring project-specific parameters of the routing endpoint. Those parameters are:
 - a. Split-horizon routing enable or disable;
 - b. Horizon of the routing endpoint;
 - c. Proxy-ARP enable or disable for the endpoint;
- 2) A proprietary MIB required for ARP table and ARP configuration. ipNetToPhysicalTable from the IP-MIB is taken as a reference for the ARP table. The main difference is that there are parameters related to VLAN configuration added to the table: VLAN type and VLAN id. Configuration of ARP table contains configurable timeout parameters for ARP module tuning:
 - a. ARP reply timeout;
 - b. ARP ageing timeout;
 - c. ARP refresh timeout;
- 3) The inetCidrRouteTable from the FORWARDING-MIB also needs to be extended with route horizon information.

A list of the above MIB addendums is provided below.

SIP-MIB header

```
tnSipMIB MODULE-IDENTITY
  DESCRIPTION
    "A MIB module provides L3 Static Routing management interface."
  ::= { enterprises transition(868) products(2) tnProducts(5) 120 }
```

IP-MIB extension

```
tnIpv4IfExtTable OBJECT-TYPE
  SYNTAX      SEQUENCE OF TnIpv4IfExtEntry
  MAX-ACCESS  not-accessible
  STATUS      current
  DESCRIPTION "This table contains additional objects for configuration of
              the routing endpoints. It is an addendum to the ipv4 table."
  ::= { tnSipMIB 1 }
```

```
tnIpv4IfExtEntry OBJECT-TYPE
  SYNTAX      TnIpv4IfExtEntry
  MAX-ACCESS  not-accessible
  STATUS      current
  DESCRIPTION "An entry containing additional management information
              applicable to a particular ipv4 interface."
  AUGMENTS   { ipv4InterfaceEntry }
  ::= { tnIpv4IfExtTable 1 }
```

```
TnIpv4IfExtEntry ::=
  SEQUENCE {
    tnIpv4IfExtHorizon          INTEGER,
    tnIpv4IfExtSplitHorizonStatus  INTEGER,
    tnIpv4IfExtArpProxyGroup     Unsigned32,
    tnIpv4IfExtArpProxyStatus    INTEGER,
  }
```

```
tnIpv4IfExtHorizon OBJECT-TYPE
  SYNTAX      INTEGER { downlink(1),
                       uplink(2) }
  MAX-ACCESS  read-write
  STATUS      current
  DESCRIPTION "Specifies horizon of the routing endpoint is Split Horizon
              routing enabled."
  DEFVAL { downlink }
  ::= { tnIpv4IfExtEntry 1 }
```

```
tnIpv4IfExtSplitHorizonStatus OBJECT-TYPE
  SYNTAX      INTEGER { enable(1),
                       disable(2) }
  MAX-ACCESS  read-write
  STATUS      current
  DESCRIPTION "Specifies enable or disable status for the Split Horizon
              routing. "
  DEFVAL { disable }
  ::= { tnIpv4IfExtEntry 2 }
```

```
tnIpv4IfExtArpProxyGroup OBJECT-TYPE
  SYNTAX      Unsigned32 (1..16)
  MAX-ACCESS  read-write
  STATUS      current
  DESCRIPTION "Specifies the number of the ARP proxy group of the routing
              entry."
  DEFVAL { 1 }
  ::= { tnIpv4IfExtEntry 3 }
```

```
tnIpv4IfExtArpProxyStatus OBJECT-TYPE
  SYNTAX      INTEGER { enable(1),
                       disable(2) }
  MAX-ACCESS  read-write
  STATUS      current
  DESCRIPTION "Specifies enable or disable status for the Split Horizon
              routing. "
  DEFVAL { disable }
```

```
::= { tnIpv4IfExtEntry 4 }
```

FORWADR-MIB extension

```
tnInetCidrRouteExtTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF TnInetCidrRouteExtEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION "This table contains additional objects for configuration of
                the routing table. It is an extension to the CIDR Internet
                route table."
    ::= { tnSipMIB 2 }
```

```
tnInetCidrRouteExtEntry OBJECT-TYPE
    SYNTAX      TnInetCidrRouteExtEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION "An entry containing additional management information
                applicable to a particular inet route."
    AUGMENTS   { inetCidrRouteEntry }
    ::= { tnInetCidrRouteExtTable 1 }
```

```
TnInetCidrRouteExtEntry ::=
    SEQUENCE {
        tnInetCidrRouteExtHorizon    INTEGER,
    }
```

```
tnInetCidrRouteExtHorizon OBJECT-TYPE
    SYNTAX      INTEGER { downlink(1),
                          uplink(2) }
    MAX-ACCESS  read-write
    STATUS      current
    DESCRIPTION "Specifies horizon of the route for Split Horizon routing."
    ::= { tnInetCidrRouteExtEntry 1 }
```

ARP-MIB extension

ARP table

```
tnRoutingArpTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF TnRoutingArpEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION "This table contains ARP table information with respect to
                Static IP routing."
    ::= { tnSipMIB 3 }
```

```
tnRoutingArpEntry OBJECT-TYPE
    SYNTAX      TnRoutingArpEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION "An entry containing information of the ARP table record."
    INDEX {
        tnRoutingArpDestNetAddr,
        tnRoutingArpEgrVlanId,
        tnRoutingArpEgrVlanType }
    ::= { tnRoutingArpTable 1 }
```

```
TnRoutingArpEntry ::=
    SEQUENCE {
        tnRoutingArpDestNetAddr    IpAddress,
        tnRoutingArpEgrVlanId      Unsigned32,
        tnRoutingArpEgrVlanType    Unsigned32,
        tnRoutingArpDestMacAddr    PhysAddress,
        tnRoutingArpEgressPort      Unsigned32,
        tnRoutingArpLastUpdated    TimeStamp,
        tnRoutingArpState          INTEGER
    }
```

```
tnRoutingArpDestNetAddr OBJECT-TYPE
    SYNTAX      InetAddress
    MAX-ACCESS  read-only
```

```

STATUS      current
DESCRIPTION "The IP Address corresponding to the media-dependent
            physical address."
::= { tnRoutingArpEntry 1 }

tnRoutingArpEgrVlanId OBJECT-TYPE
SYNTAX      Unsigned32 (0..4095)
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION "Specifies the egress vlan id of the address."
::= { tnRoutingArpEntry 2 }

tnRoutingArpEgrVlanType OBJECT-TYPE
SYNTAX      Unsigned32 (0..4294967295)
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION "Specifies the egress vlan ethernet type of the address."
::= { tnRoutingArpEntry 2 }

tnRoutingArpDestMacAddr OBJECT-TYPE
SYNTAX      PhysAddress (SIZE(0..65535))
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION "The media-dependent (MAC) physical address."
::= { tnRoutingArpEntry 4 }

tnRoutingArpEgressPort OBJECT-TYPE
SYNTAX      Unsigned32 (0..16)
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION "Egress port to reach the destination host."
::= { tnRoutingArpEntry 5 }

tnRoutingArpLastUpdated OBJECT-TYPE
SYNTAX      TimeStamp
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION "The value of sysUpTime at the time this entry was last
            updated. If this entry was updated prior to the last
            re-initialization of the local network management subsystem,
            then this object contains a zero value."
::= { tnRoutingArpEntry 6 }

tnRoutingArpState OBJECT-TYPE
SYNTAX      INTEGER { ok(1),
                    resolving(2),
                    refreshing(3),
                    failed(4) }
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION "Host reachability state."
DEFVAL { disable }
::= { tnRoutingArpEntry 7 }

```

ARP configuration parameters

```

tnRoutingArpConfigReplyWaitTime OBJECT-TYPE
SYNTAX      Unsigned32 (1..10)
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION "Specifies timeout in seconds between ARP request send and
            declaration of ARP failed."
::= { tnSipMIB 4 }

tnRoutingArpConfigAgeingTime OBJECT-TYPE
SYNTAX      Unsigned32 (60..86400)
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION "Specifies timeout in seconds to keep unused ARP record in
            the ARP table."
DEFVAL { 120 }
::= { tnSipMIB 5 }

```

```
tnRoutingArpConfigRefreshTime OBJECT-TYPE
    SYNTAX      Unsigned32 (60..86400)
    MAX-ACCESS  read-write
    STATUS      current
    DESCRIPTION "Specifies timeout in seconds between consequential
                refreshing ARP request sent for the record in the ARP
                table."
    DEFVAL { 600 }
    ::= { tnSipMIB 6 }
```

For Additional MIB Information

For the list of all of the S3280 SNMP Traps see related section in the S3280 User Guide.

For more information on the SNMP Agent, Network Management Station (NMS), MIBS, MIB modules and MIB Variables, the Object ID (OID), the MIB Tree / branch /node, MIB Table Indices, values, notations and transaction types, etc., see the SNMP Primer at <http://www.transition.com/pshelp/snmp.html#indices>.

Appendix A: Service, Warranty & Compliance Information

Service

See the "Service" section in the online device User Guide manual for regulatory agency compliance information.

Warranty

See the "Warranty" section in the online device User Guide manual for regulatory agency compliance information.

Compliance Information

See the "Compliance Information" section in the online device User Guide manual for regulatory agency compliance information.

Cautions and Warnings

Definitions

Cautions indicate that there is the possibility of poor equipment performance or potential damage to the equipment. Warnings indicate that there is the possibility of injury to a person.

Cautions and Warnings appear here and may appear throughout this manual where appropriate. Failure to read and understand the information identified by this symbol could result in poor equipment performance, damage to the equipment, or injury to persons.

See "Electrical Safety Warnings" in the online device User Guide manual for Electrical Safety Warnings translated into multiple languages.

Glossary of SIR Terms

This section describes many of the terms and mnemonics used in this manual. Note that the use of or description of a term does not in any way imply support of that feature or of any related function(s). See the related User Guide for descriptions of terms other than these related to SIR.

ARP

Address Resolution Protocol (ARP) is a telecommunications protocol used for resolution of network layer addresses into link layer addresses, a critical function in multiple-access networks. ARP was defined by IETF RFC 826 in 1982. ARP is designated by Internet Standard STD 37, and is updated by RFCs [5227](#) and [5494](#). ARP is also the name of the program for manipulating these addresses in most operating systems. ARP is a required TCP/IP standard defined in RFC 826, "Address Resolution Protocol (ARP)."

ARP is used to translate protocol addresses to hardware interface addresses (RARP provides the complementary translation). ARP is implemented in many network / internetwork technologies, (e.g., IPv4 using IEEE 802 standards, FDDI, X.25, Frame Relay and ATM, IPv4 over IEEE 802.3 and IEEE 802.11). ARP uses a simple message format that contains one address resolution request or response. In Internet Protocol Version 6 (IPv6) networks, ARP functionality is provided by the Neighbor Discovery Protocol (NDP). See also "Gratuitous ARP".

CIDR

Classless Inter-Domain Routing (CIDR) is a method for allocating IP addresses and routing Internet Protocol packets. The IETF introduced CIDR in 1993 to replace the previous addressing architecture of classful Internet network design. Their goal was to slow the growth of routing tables on routers across the Internet, and to help slow the rapid exhaustion of IPv4 addresses. CIDR was initially defined in IETF [RFC 1518](#). IP addresses are described as consisting of two groups of bits in the address: the more significant part is the network address, which identifies a whole network or subnet, and the less significant portion is the host identifier, which specifies a particular interface of a host on that network. This division is used as the basis for traffic routing between IP networks and for address allocation policies. Classful network design for IPv4 sized the network address as one or more 8-bit groups, resulting in the blocks of Class A, B, or C addresses. CIDR allocates address space to Internet service providers and end users on any address bit boundary, instead of on 8-bit segments. In IPv6, however, the interface identifier has a fixed size of 64 bits by convention, and smaller subnets are never allocated to end users. CIDR notation specifies IP addresses and their associated routing prefix. It appends to the address a slash character and the decimal number of leading bits of the routing prefix (e.g., *192.0.2.0/24* for IPv4 or *2001:db8::/32* for IPv6).

Cost: assigns an integer cost metric (ranging from 1 to 9,999) to be used in calculating the fastest, most reliable, and/or least expensive route. If a cost metric is not specified, 1 is used.

Destination

Specifies either an IP address or host name for the destination network or host (e.g., *10.100.5.0/24*).

Destination IP address

The IP address of the PC. The next hop IP address is the Ethernet interface directly opposite. The IP address assigned to the Ethernet management interface; it must be within the same subnet and subnet mask as the IP address on the router for this Ethernet connection.

FDMA

(Frame Direct Memory Access) frame-based form of DMA, which is the hardware mechanism that allows peripheral components to transfer their I/O data directly to and from main memory without the need to involve the system processor. Use of this mechanism can greatly increase throughput to and from a device, because a great deal of computational overhead is eliminated.

GARP

Gratuitous ARP - a gratuitous ARP reply is a reply to which no request has been made. It involves a request/reply that is not normally needed according to the ARP specification (RFC 826) but could be used in some cases. A gratuitous ARP request is an Address Resolution Protocol request packet where the source and destination IP are both set to the IP address of the machine issuing the packet, and the destination MAC is the broadcast address `ff:ff:ff:ff:ff:ff`. No reply packet would occur. See IETF [RFC 3927](#).

Gateway address

Specifies either an IP address or host name for the gateway or router to use when forwarding.

ICMP

Internet Control Message Protocol is part of the Internet Protocol Suite, as defined in IETF [RFC 792](#). ICMP messages are typically used for diagnostic or control purposes or generated in response to errors in IP operations (per RFC 1122). ICMP errors are directed to the source IP address of the originating packet. Each device (e.g., an intermediate router) forwarding an IP datagram first decrements the time to live (TTL) field in the IP header by one. If the resulting TTL is 0, the packet is discarded and an ICMP Time To Live exceeded in transit message is sent to the datagram's source address. ICMP differs from transport protocols such as TCP and UDP in that it is not typically used to exchange data between systems, nor is it regularly employed by end-user network applications (other than diagnostic tools such as ping and traceroute). RFC 792 was updated by RFCs [777](#) and [760](#).

Interface

Specifies the interface to be used for the route that uses the interface number. If an interface is not specified, the interface to be used for the route is determined from the gateway IP address.

ISR

(Interrupt Service Routine) a software routine that hardware invokes in response to an interrupt. An ISR examines an interrupt and determines how to handle it. ISRs handle the interrupt, and then return a logical interrupt value. If no further handling is required because the device is disabled or data is buffered, the ISR notifies the kernel. An ISR must perform very fast to avoid slowing down the operation of the device and the operation of all lower priority ISRs. External events trigger an interrupt; the normal control flow is interrupted and an ISR is called. Such events can be triggered by hardware or software.

Linktrace Protocol (Trace)

Protocol used for path discovery between a pair of maintenance points. Linktrace messages are triggered by an administrator using the traceroute command to verify the path between a pair of maintenance end points (MEPs) under the same maintenance association. Linktrace messages can also be used to verify the path between an MEP and a maintenance intermediate point (MIP) under the same maintenance domain. The operation of IEEE 802.1ag linktrace request and response messages is similar to the operation of Layer 3 traceroute commands.

MAC

(Media Access Control) address of the Ethernet management interface.

Proxy ARP

Proxy ARP (Proxy Address Resolution Protocol) - a technique where a device on a given network answers the ARP queries for a network address that is not on that network. The ARP Proxy is aware of the location of the traffic's destination, and offers its own MAC address in reply. Serving as an ARP Proxy for another host effectively directs LAN traffic to the Proxy. Typically, the "captured" traffic is then routed by the Proxy to the intended destination via another interface or via a tunnel. The process which results in the node responding with its own MAC address to an ARP request for a different IP address for proxying purposes is sometimes called 'publishing'. See IETF [RFC 925](#) and [RFC 1027](#) for more information.

SIR

(Static IP Routing) Static routing is a concept describing a way to configure path selection of routers in computer networks. This type of routing is characterized by the lack of communication between routers regarding the current topology of the network. This is achieved by manually adding routes to a routing table. Contrast “dynamic routing”.

Split-horizon DNS

(AKA split-horizon DNS, split-view DNS, or split DNS) the facility of a Domain Name System (DNS) implementation that provides different sets of DNS information, usually selected by the source address of the DNS request. This facility provides a mechanism for security and privacy management by logical or physical separation of DNS information for network-internal access (within an administrative domain such as a company) and access from an insecure public network (e.g., the Internet).

Static route

A “hard coded” path in a router that specifies how the router will get to a certain subnet via a particular path.

Static IP address

A fixed IP address that is configured on the settings of a computer or router. Typically involves as IP address, Subnet Mask, Default Gateway, and Domain Name Server (DNS).

Static IP route cost

Assigns an integer cost metric (ranging from 1 through 9,999) to be used in calculating the fastest, most reliable, and/or least expensive routes. If cost metric is not specified, 1 is used.

Static IP route destination

Specifies either an IP address or host name for the network or host.

Static IP route gateway

Specifies either an IP address or host name for the gateway or router to use when forwarding.

Static IP route interface

Specifies the interface to be used for the route that uses the interface number. If an interface is not specified, the interface to be used for the route is determined from the gateway IP address.

subnet mask

All computers that belong to a subnet are addressed with a common, identical, most-significant bit-group in their IP address. In IPv4, the routing prefix is specified in the form of the subnet mask, which is expressed in quad-dotted decimal representation like an address. For example, 255.255.255.0 is the network mask for the 192.168.1.0/24 prefix. Specifies a subnet mask to be associated with a route entry; if a subnet mask is not specified, 255.255.255.255 is used.

Subnet mask length

Should be 8 for Class A, 16 for Class B, and 24 for Class C networks.

Tera Term

A free open-source software implemented terminal emulator (communications) program. Tera Term emulates various types of computer terminals (e.g., DEC VT100, DEC VT382). Tera Term supports telnet, SSH 1, SSH 2, and serial port connections. Tera Term (or TeraTerm) also has a built-in macro scripting language and other useful plugins. Tera Term has features similar to PuTTY and SecureCRT.



Transition Networks

10900 Red Circle Drive

Minnetonka, MN 55343 USA

Tel: 952-941-7600 or 1-800-526-9267

Fax: 952-941-2322

Copyright© 2012, 2013 Transition Networks

All rights reserved.

Printed in the U.S.A.

Static IP Routing (SIR) User Guide, 33542 Rev. A