



SRstackware®
Frequently Asked Questions
P/N: 6806800N91E
August 2022



Legal Disclaimer*

SMART Embedded Computing, Inc. (SMART EC), dba Penguin Solutions™, assumes no responsibility for errors or omissions in these materials. **These materials are provided "AS IS" without warranty of any kind, either expressed or implied, including but not limited to, the implied warranties of merchantability, fitness for a particular purpose, or non-infringement.** SMART EC further does not warrant the accuracy or completeness of the information, text, graphics, links, or other items contained within these materials. SMART EC shall not be liable for any special, indirect, incidental, or consequential damages, including without limitation, lost revenues or lost profits, which may result from the use of these materials. SMART EC may make changes to these materials, or to the products described therein, at any time without notice. SMART EC makes no commitment to update the information contained within these materials.

Electronic versions of this material may be read online, downloaded for personal use, or referenced in another document as a URL to a SMART EC website. The text itself may not be published commercially in print or electronic form, edited, translated, or otherwise altered without the permission of SMART EC.

It is possible that this publication may contain reference to or information about SMART EC products, programming, or services that are not available in your country. Such references or information must not be construed to mean that SMART EC intends to announce such SMART EC products, programming, or services in your country.

Limited and Restricted Rights Legend

If the documentation contained herein is supplied, directly or indirectly, to the U.S. Government, the following notice shall apply unless otherwise agreed to in writing by SMART EC.

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data clause at DFARS 252.227-7013 (Nov. 1995) and of the Rights in Noncommercial Computer Software and Documentation clause at DFARS 252.227-7014 (Jun. 1995).

SMART Embedded Computing, Inc., dba Penguin Solutions

2900 S. Diablo Way, Suite 190

Tempe, Arizona 85282

USA

*For full legal terms and conditions, visit <https://www.penguinsolutions.com/edge/legal/>

Table of Contents

About this Manual	7
1 General Information	11
1.1 Overview	11
1.1.1 What is the current protocol segregation between Basic and Enhanced SRstackware	11
1.1.2 How to quickly demonstrate a certain protocol of SRstackware at customer place	12
1.1.3 How to quickly configure SRstackware	13
1.1.4 What does the copy running-config startup-config command do	13
1.1.5 How to enable logging in SRstackware	13
1.1.6 How to dump the output of show command in CLI into a txt file	14
1.1.7 Does SRstackware provide API for configuration	14
2 Layer 2 Information	15
2.1 Overview	15
2.1.1 Interface Related	15
2.1.1.1 Behavior of LoopBack Interface in SRstackware®	15
2.1.1.2 How to detect whether the interface link is up/down	15
2.1.1.3 What is the meaning of the interface statistics shown in the show interface command	15
2.1.1.4 How to add an L2 static entry to a switchport interface	15
2.1.1.5 How storm control functions in SRstackware	16
2.1.1.6 How to see the L2 table MAC entries	16
2.1.2 VLAN	17
2.1.2.1 What are the configurable VLAN IDs in SRstackware	17
2.1.2.2 How to remove a tagged port from VLAN 1	17
2.1.2.3 How to configure an untagged vlan	17
2.1.2.4 How to configure both tagged and untagged vlan	17
2.1.2.5 How to configure only tagged vlan	17
2.1.2.6 How does ingress filter affect the port properties	18
2.1.2.7 Does SRstackware 2.1 software support double VLAN also called as VLAN Stacking	18
2.1.2.8 How to configure VLAN Stacking in SRstackware and are there any limitations	18
2.1.3 LACP	19
2.1.3.1 What are the scenarios to configure LACP	19

Table of Contents

2.1.3.2	Does SRstackware support Split-LACP	19
2.1.3.3	What are the load-balancing techniques supported by SRstackware	19
2.1.3.4	What is the procedure to be followed while configuring LACP through channel-group	20
2.1.3.5	What is the procedure to be followed while configuring static-channel-group	20
2.1.3.6	What are the other procedural assumptions while configuring channel-group or static-channel-group	20
2.1.3.7	How does unicast or non-unicast traffic affect load balancing	20
2.1.3.8	How to know the bandwidth of an Aggregator Port	21
2.1.3.9	What are the limitations imposed on Load balancing by Broadcom	21
2.1.4	STP, RSTP, and MSTP	21
2.1.4.1	How to disable STP per port	21
2.1.4.2	Can I enable MSTP over LACP interface	22
2.1.5	SNMP	22
2.1.5.1	Are there any known limitations for SNMP	22
2.1.5.2	What are the SNMP MIBs supported in SRstackware 2.1 release	22
2.1.5.3	How to access L2 SNMP MIB objects for multiple bridges	23
2.1.5.4	How to access and configure individual switch ports using L2 MIBs	24
3	Layer 3 Information	27
3.1	Overview	27
3.1.1	Interface Related	27
3.1.1.1	Explain the routerport concept. How to make a switchport to use an L3 interface	27
3.1.1.2	Explain vlnx.y interface and its working in linux and SRstackware®	27
3.1.1.3	How to configure switch to send higher data traffic to an application running on management processor	28
3.1.1.4	How does an SRstackware convert a switchport to a routerport	28
3.1.1.5	How to configure the DOT1X on VLAN Interface	28
3.1.2	Static Routing	29
3.1.2.1	How to add a static arp entry	29
3.1.2.2	Is there a command that can register the MAC table and the ARP table in Static	29
3.1.2.3	How to add a static route	29
3.1.2.4	How to add a static route with next hop address directly to FIB (without slow path)	30
3.1.2.5	How to add a static route with just interface name (without next hop)	

	directly to FIB (without slow path)	30
3.1.2.6	How to add static route with next hop address on a network configured on a vlanx.y interface (without slow path)	30
3.1.3	QoS	31
3.1.3.1	Why some of the class-maps throw error when configured at the same time	31
3.1.4	Access Control Lists and Match lists	31
3.1.4.1	What is the difference between ip-access-list and access-list	31
3.1.4.2	What are Match lists	31
3.1.4.3	What are rules and actions	32
3.1.4.4	What is the difference between outervlan and innervlan in match L2 param criteria	32
3.1.4.5	Are there any limitations to match lists	32
3.1.4.6	When multiple rules exist with the same set of qualifiers, which of these rules will be applied	32

Table of Contents

About this Manual

Overview of Contents

This document explains the SRstackware® behavior that occurs due to specific implementation of either software or hardware. This should be kept in mind while using SRstackware on any hardware or platform until explicitly mentioned.

This document is based on the features available in SRstackware 2.1 and higher version.

This manual consists of the following chapters:

[Chapter 1, General Information on page 11](#)

[Chapter 2, Layer 2 Information on page 15](#)

[Chapter 3, Layer 3 Information on page 27](#)



Abbreviations






This document uses the following abbreviations:

Abbreviation	Definition
CLI	Command Line Interface
MAC	Media Access Control
OSPF	Open Shortest Path First
QoS	Quality of Service
RIP	Routing Information Protocol
SNMP	Simple Network Management Protocol
SRstackware	Switching and Routing stackware
VLAN	Virtual LAN

Conventions

The following table describes the conventions used throughout this manual.

Notation	Description
0x00000000	Typical notation for hexadecimal numbers (digits are 0 through F), for example used for addresses and offsets
0b0000	Same for binary numbers (digits are 0 and 1)
bold	Used to emphasize a word
Screen	Used for on-screen output and code related elements or commands. Sample of Programming used in a table (9pt)
Courier + Bold	Used to characterize user input and to separate it from system output
<i>Reference</i>	Used for references and for table and figure descriptions
File > Exit	Notation for selecting a submenu
<text>	Notation for variables and keys
[text]	Notation for software buttons to click on the screen and parameter description
...	Repeated item for example node 1, node 2, ..., node 12
.	Omission of information from example/command that is not necessary at the time
..	Ranges, for example: 0..4 means one of the integers 0,1,2,3, and 4 (used in registers)
	Logical OR
	Indicates a hazardous situation which, if not avoided, could result in death or serious injury
	Indicates a hazardous situation which, if not avoided, may result in minor or moderate injury

Notation	Description
	Indicates a property damage message
	Indicates a hot surface that could result in moderate or serious injury
	Indicates an electrical situation that could result in moderate injury or death
<p>Use ESD protection</p> 	Indicates that when working in an ESD environment care should be taken to use proper ESD practices
	No danger encountered, pay attention to important information

Summary of Changes

This manual has been revised and replaces all prior editions.

Part Number	Publication Date	Description
6806800N91E	August 2022	Rebrand to Penguin Solutions.
6806800N91D	March 2020	Rebrand to SMART Embedded Computing. Updated Abbreviations table.
6806800N91C	July 2017	Added registered trademark to SRstackware.
6806800N91B	June 2014	Rebrand to Artesyn template.
6806800N91A	February 2012	EA Release.

General Information

1.1 Overview

This section provides FAQs on basic SRstackware® configuration.

1.1.1 What is the current protocol segregation between Basic and Enhanced SRstackware

Table 1-1 Supported Features

Features	Basic SRstackware	Enhanced SRstackware
Switch Management		
Command Line Interface (CLI)	Yes	Yes
Simple Network Management Protocol (SNMPv3)	Yes	Yes
Broadcast storm recovery	Yes	Yes
Flow control	Yes	Yes
Application Programming Interface (API) access	Yes	Yes
Layer 2 Switching		
Spanning Tree Protocol (STP)	Yes	Yes
Rapid Spanning Tree Protocol (RSTP)	Yes	Yes
Multiple Spanning Tree Protocol (MSTP)	Yes	Yes
Virtual LAN (VLAN) Tagging	Yes	Yes
Link Aggregation Control Protocol (LACP)	Yes	Yes
VLAN Classification by Protocol and Port	Yes	Yes
Class of Service (CoS)	Yes	Yes
Link Aggregation Control Protocol (LACP)	Yes	Yes
Generic Attribute Registration Protocol (GARP)	Yes	Yes
GARP Multicast Registration Protocol (GMRP)	Yes	Yes

General Information

Table 1-1 Supported Features

Features	Basic SRstackware	Enhanced SRstackware
GARP VLAN Registration Protocol (GVRP)	Yes	Yes
VLAN Stacking (Q-in-Q)	Yes	Yes
Static Filtering (ACL)	Yes	Yes
Enhanced Load Balancing (by TCP/UDP port)	Yes	Yes
Layer 3 Routing		
IPv4 Routing		
Internet Group Management Protocol (IGMP v1)	No	Yes
Internet Group Management Protocol (IGMP v2)	No	Yes
Internet Group Management Protocol (IGMP v3)	No	Yes
IGMP Snooping/Proxy	No	Yes
Routing Information Protocol (RIPv2)	No	Yes
Open Shortest Path First (OSPFv2)	No	Yes
Virtual Router Redundant Protocol (VRRP)	No	Yes
IPv6 Routing		
RIP Next generation (RIPng)	No	Yes

1.1.2 How to quickly demonstrate a certain protocol of SRstackware at customer place

Refer to Protocol Demo Guide supplied with the blade documentation. This guide provides the topology and the configuration steps to demonstrate a certain protocol of SRstackware.

1.1.3 How to quickly configure SRstackware

Refer to SRstackware section of the *Basic Blade Services Programmer's Reference Guide* for quick configuration tips.

Refer to Protocol Demo Guide supplied with the blade documentation to know the configuration for a particular protocol.

Refer to Command Reference documents for information on the CLI commands.

If you want to quickly configure your own VLANs and topology with protocols, refer to SRstackware default configuration, and also refer to Command Reference documents to get to know the CLI commands.

1.1.4 What does the `copy running-config startup-config` command do

The SRstackware configuration is persistent in all the blades currently supported. Read the blade documentation for any exceptions.

The `copy running-config startup-config` command copies the current running configuration to the startup configuration file.

Use this command if you want to write configurations to a file that will be used at startup time. This command is the same as the `write memory` command, and is available in the Privileged Exec mode.

Example:

```
CLI> enable
CLI# copy running-config startup-config
```

1.1.5 How to enable logging in SRstackware

Issue the following:

```
CLI#conf t
```

Enter configuration commands, one per line. End with CNTL+Z. CLI(config)#log ?

file	Logging to file
record-priority	Log the priority of the message within the message
syslog	Logging goes to syslog
trap	Limit logging to specified level

```
CLI(config)#log syslog
```

General Information

The syslog command will start logging in `/var/log/messages`.

```
CLI(config)#log file ?
```

```
FILENAME                               Logging filename
```

Logs the messages to the file of your choice. You can give the complete path.

```
CLI(config)#log file /var/log/test
```

1.1.6 How to dump the output of `show` command in CLI into a txt file

Use output redirection `>` symbol.

Example:

```
CLI#show interface > /var/frame.txt
```

1.1.7 Does SRstackware provide API for configuration

Yes, SRstackware provides API for configuration. Refer to API Developer Guide for more information. It also provides configuration through API from remote node. Check the features and protocols list in the *Basic Blade Services Programmer's Reference* of the specific blade release.

Layer 2 Information

2.1 Overview

This section provides FAQs on Layer 2 Protocols.

2.1.1 Interface Related

2.1.1.1 Behavior of LoopBack Interface in SRstackware®

Do not shutdown the loopback (lo) interface in the linux/SRstackware and save this configuration. The SRstackware daemons depend on this interface and cannot communicate with each other.

2.1.1.2 How to detect whether the interface link is up/down

The `show interface <intfname>` command shows the RUNNING flag if the link is up. If the RUNNING flag is not present in the output then the link is down.

2.1.1.3 What is the meaning of the interface statistics shown in the `show interface` command

input packets xxxx, bytes xxxx, dropped xxx, multicast packets xxxx

output packets xxxx, bytes xxxx, multicast packets xxx broadcast packets xx

input packets count = Input Unicast + Input NonUnicast Packets

output packets count = Output Unicast + Output NonUnicast Packets

multicast packets count = Input Multicast + Output Multicast Packets



You may see multicast count if the LHC is sending packets on the interface. All the numbers displayed are in decimal.

2.1.1.4 How to add an L2 static entry to a switchport interface

Assuming that the ge13 is in the same bridge and in the same vlanid, enable the interface.

```
(config)#interface ge13
(config-if)# no shutdown
```

Layer 2 Information

Configure MAC address on the switchport to allow forwarding of the packets destined to this MAC address.

```
(config)# bridge <num> address <macaddress> ge13 forward vlan  
<vlanid>
```

2.1.1.5 How storm control functions in SRstackware

SRstackware storm control works on the frame level, that is, it always forwards a fixed no of packets as per configured threshold irrespective of the frame size.

In Storm Control CLI, the threshold level specifies the percentage of the maximum packet rate supported on the interface with the packet size of 1512. Hence on an interface of 1Gbps:

maximum packet-rate = $(1000*1000*1000)/(1512*8) = 82670$ packets per second.

Example: Configuring storm-control on broadcast level to 0.10%:

```
(config-if)#storm-control broadcast level 0.10
```

As mentioned above if you configure a 1G port for 0.10% of broadcast level, then it will allow a forwarding of approximately 82 packets. Similarly if you configure a 1G port for 10.00% of broadcast level, then it allows a forwarding of only 8267 (approximate) packets.



Storm Control will not stop the packets transmitting to CPU.

2.1.1.6 How to see the L2 table MAC entries

Currently SRstackware does not have a CLI command which can display the Layer 2 forwarding table. However, SRS supports SNMP objects for displaying the L2 entries through the below dot1q tables.

To walk on L2 table:

```
Linux> snmpwalk -v3 -u admin -A adminpwd123 -Ob localhost  
dot1qTpFdbTable  
Description  
dot1qTpFdbPort.12.0.128.66.39.53.156 = INTEGER: 5046  
dot1qTpFdbPort.<vlanid>.<MACaddress in decimal> = <Port ifindex>  
dot1qTpFdbStatus.12.0.128.66.39.53.156 = INTEGER: learned(3)
```

Where, dot1qTpFdbStatus.<vlanid>.<mac address in decimal> = <status of the entry>

To know the ifindex to port name:

```
root@atca-f120:/root> snmpwalk -v3 -u admin -A adminpwd123 -Ob
localhost ifName
Description
IF-MIB::ifName.5020 = STRING: ge20
Where, ifName.<ifindex> = <port name>
```

If you want the MAC addresses in hex for easy searching, then you can use the below command `snmpwalk -v3 -u admin -A adminpwd123 -Ob localhost dot1dTpFdbTable`

2.1.2 VLAN

2.1.2.1 What are the configurable VLAN IDs in SRstackware

The VLAN IDs are in the range from 2 to 4022. If you want to enable intervlan routing, then the maximum number of intervlan routing enabled VLANs must be limited to 128.

The SNMP MIB sets does not allow operations on default vlan 1.

2.1.2.2 How to remove a tagged port from VLAN 1

By design, the VLAN 1 can be removed only when the port is assigned with at least one untagged vlan either through access or hybrid configuration.

2.1.2.3 How to configure an untagged vlan

Use the access mode configuration to configure untagged vlan. In this case, the port is removed from default vlan 1.

2.1.2.4 How to configure both tagged and untagged vlan

Use the hybrid mode configuration. The default vlan configured through hybrid is an untagged vlan. The other vlans configured are basically tagged vlans. If the default vlan is configured, then the VLAN 1 is removed from this port.

2.1.2.5 How to configure only tagged vlan

Use the trunk mode configuration. The default vlan 1 configured but it is a tagged vlan.

If Ingress filter is disabled and if you have set the native vlan then untagged packets are also received and it attaches default vlan for those packets.

Layer 2 Information

2.1.2.6 How does ingress filter affect the port properties

The ingress filtering is enabled by default in SRS. The access configuration receives only untagged packets. The hybrid configuration receives only tagged/untagged based on acceptable-frame-type field. The trunk configuration receives only tagged packets.

If the ingress filtering is disabled, then the port starts receiving both untagged and tagged packets irrespective of the mode but the vlans properties remain as per configuration.

2.1.2.7 Does SRstackware 2.1 software support double VLAN also called as VLAN Stacking

SRstackware 2.1 supports VLAN stacking configuration. But it does not support Provider Bridging.

2.1.2.8 How to configure VLAN Stacking in SRstackware and are there any limitations

To configure VLAN Stacking in SRstackware, use the following commands in interface mode: `#switchport vlan-stacking customer-edge-port` (for customer port)
When the node is customer-edge, the inner-vlan tag is added for the ingress packets and the inner-vlan tag is removed for the egress packets.

```
#switchport vlan-stacking provider-port (for provider port)
```

When the node is provider, the double tag is retained on both ingress and egress packets. The inner-tag is used for L2 forwarding. The provider-edge ports can also forward single tagged packets.

Following are some of the limitations imposed by BCM:

- VLAN Stacking can be enabled at chip level and cannot be enabled at port level. Because of this limitation, the following behavior is observed:
 - As soon as a port is configured as customer-edge port, all the ports on the chipset becomes customer-edge ports.
 - As soon as a port is configured as provider port, all the ports on the chip-set is set to customer-edge except the port configured to provider.
- On removing `vlan-stacking` configuration using `no vlan-stacking` command, it is removed on all the ports on the same chip-set.
- When `no switchport` command is executed on the port which has `vlan-stacking` configuration, it does not modify the configuration of `vlan-stacking` at BCM.

NOTE: The Provider bridging functionality is not supported.

2.1.3 LACP

2.1.3.1 What are the scenarios to configure LACP

The channel-group command is provided to configure LACP. The ports should have same properties and they should be physically present on the same chipset to configure LACP. An interface po# is created per aggregated interface for configuration.

The configuration can be done on individual ports and this configuration is propagated to aggregated interface after syncing.

2.1.3.2 Does SRstackware support Split-LACP

The static-channel-group is provided for the same. This does not follow link aggregation protocol and will not have LACP control packets. But this is basically trunking of ports. All the ports in the trunk should be present on the same chipset on one end, connected to ports (normal, not aggregated) on different ends. All the load balancing algorithms work on static-channel groups also.

In this case an interface sa# is created per static-channel for configuration.

2.1.3.3 What are the load-balancing techniques supported by SRstackware

Following load-balancing techniques are supported by SRstackware:

dst-mac	Destination Mac address based load balancing
src-dst-mac	Source and Destination Mac address based load balancing
src-mac	Source Mac address based load balancing
src-ip	Source IP based load balancing
src-dst-ip	Source and Destination IP based load balancing
dst-ip	Destination IP based load balancing

In addition to the above load balancing techniques, the SRstackware supports enhanced mode in which the switch chipset uses enhanced hash algorithm for better load balancing results. You can add more fields for the hashing algorithm in this mode. The load balancing can be visibly seen when there are multiple flows, typically more than three to four flows of random traffic.

load-balancing field-select normal/enhanced command has to be used to configure enhanced mode load-balancing options. This command should be executed on the ingress port and not on the aggregator or member port. Ingress port is the one from where the traffic is coming in to the chip and then the traffic is forwarded to the aggregator.

Check the Release Notes of the blade documentation if there are any limitations regarding LACP for a particular chipset.

Layer 2 Information

2.1.3.4 What is the procedure to be followed while configuring LACP through channel-group

When member ports are in autonegotiation mode: When the first port is up while adding to aggregator, the aggregator assumes that these are the master port properties. While adding second port onwards to the aggregator, the bandwidth and duplex have to be same as the first port added to the aggregator. This requires that the ports have to be linked up and auto negotiated properly to match the first port properties. To avoid the above situation, all the ports should be added to aggregator in shutdown state. After adding all the ports, bring up the aggregator port using no shutdown CLI.

When member ports are manually configured as duplex full: All the member ports should match the first port manual configuration.

2.1.3.5 What is the procedure to be followed while configuring static-channel-group

In manually configured duplex mode or in auto-negotiation mode, all:

- Member ports should be administratively up and in operationally running state
- Member ports properties should match the first port configuration

A maximum of eight ports can be configured in a static-channel-group(sa#).

2.1.3.6 What are the other procedural assumptions while configuring channel-group or static-channel-group

- The configuration of normal or enhanced load balanced algorithm can be performed only at chipset level.
- Save the aggregator configuration only when the aggregator is functionally up, in sync, and running. This is very important to have the configuration to be persistent.
- Bandwidth, duplex value, VLAN, and bridge configuration are only propagated to the aggregator port from the first member port. All other desired configuration should be done on the aggregator port (po#/sa#) itself.
- Maximum of eight aggregators (po# & sa#) can be configured at an instance.

2.1.3.7 How does unicast or non-unicast traffic affect load balancing

A unicast traffic is a learned traffic by the switch. There are hosts/routers connected to the blade which actually has that particular MAC/IP addresses.

A non-unicast traffic is an unlearned traffic by the switch. The switch does a flood to that particular vlan for this traffic. For example the broadcast and multicast traffic are considered non-unicast. But a unicast address for which there are no hosts/routers are also considered non-unicast.

The load balancing algorithm behaves differently for addresses which are learned and not learned. But it is easy to say from the topology whether the switch can learn a particular address or not, so that the load balancing behavior is predictable.

2.1.3.8 How to know the bandwidth of an Aggregator Port

Bandwidth of an Aggregator port (sa#/po#) can be known using `show interface sa#/po#` command. The bandwidth shown in this command is computed as follows:

(Number of member ports) * (Bandwidth of each member port)

The member ports here include all the active and passive members of an Aggregator Port.



The bandwidth displayed for the aggregated port is not the running bandwidth, that is, it always displays two times the bandwidth of single port, even if any of the port is down.

2.1.3.9 What are the limitations imposed on Load balancing by Broadcom

- When the traffic flow is fragmented, load balancing based on UDP/TCP port numbers is not supported.
- If IPv6 extension headers is enabled (in case of IPv6) or IP options are enabled (in case of IPv4) in the traffic flows, then the flows are not load balanced based on UDP/TCP port numbers.

2.1.4 STP, RSTP, and MSTP

2.1.4.1 How to disable STP per port

The STP is enabled on the bridge and to disable it per-port on that bridge, use the command `bridge # spanning-tree disable`.

Layer 2 Information

2.1.4.2 Can I enable MSTP over LACP interface

All the VLAN and other configuration should be done on the individual ports and then LACP should be configured. This configuration is then automatically configured on LACP interface. However, in case of MSTP, it is better to reconfigure the interface level configuration on the LAG port.

Example:

```
int po1
  bridge-group <bridge_id> instance <instance_id>
```

To remove/modify any of these configurations, remove the channel-group and then reconfigure it on member ports.

2.1.5 SNMP

2.1.5.1 Are there any known limitations for SNMP

- The default vlan ID 1 is not configurable through SNMP MIB.
- Refer the MIB constraints document that is provided along with the release for any other known limitations.

2.1.5.2 What are the SNMP MIBs supported in SRstackware 2.1 release

The following are the MIBs supported in SRstackware 2.1 release:

- BRIDGE-MIB (RFC 4188)
- P-BRIDGE-MIB (RFC 4363)
- Q-BRIDGE-MIB (RFC 4363)
- IF-MIB (RFC 2863)
- OSPF-MIB (RFC 1850)
- IEEE8023-LAG-MIB (RFC IEEE Std 802.3ad)
- RSTP-MIB (RFC 4318)
- RIP MIB (RFC)
- VRRP MIB (RFC)
- SRS-MIB (Penguin Solutions Proprietary MIB)

2.1.5.3 How to access L2 SNMP MIB objects for multiple bridges

For Bridge-based MIBs, the context parameter of the snmp commands (snmpget/snmpset/snmpwalk) is used to specify the bridge number in the format of Bridge<Number>.

In case of F120 and F125, in default configuration, Bridge1 represents Base switch and Bridge2 represents Fabric switch.

Example of an snmp command to access a Bridge1 (Base Switch) object:

```
snmpget -v3 -u admin -n "Bridge1" -l noAuthNoPriv -a MD5 -A  
adminpwd123 localhost dot1dBaseBridgeAddress.0
```

The following are the MIBs, among the currently supported MIBs, for which Bridge Number needs to be specified as context:

- BRIDGE-MIB (RFC 4188)
- P-BRIDGE-MIB (RFC 4363)
- Q-BRIDGE-MIB (RFC 4363)
- RSTP-MIB (RFC 4318)

Layer 2 Information

2.1.5.4 How to access and configure individual switch ports using L2 MIBs

Refer the mapping between port names and port ifindex. In case of ATCA-F120, refer below list.

Port Name	ifindex	Port Name	ifindex
ge1	5001	ge44	5065
ge2	5002	ge45	5066
ge3	5003	ge46	5067
ge4	5004	ge47	5068
ge5	5005	ge48	5069
ge6	5006	xe1	5025
ge7	5007	xe2	5026
ge8	5008	xe3	5027
ge9	5009	xe4	5028
ge10	5010	xe5	5029
ge11	5011	xe6	5030
ge12	5012	xe7	5031
ge13	5013	xe8	5032
ge14	5014	xe9	5033
ge15	5015	xe10	5034
ge16	5016	xe11	5035
ge17	5017	xe12	5036
ge18	5018	xe13	5037
ge19	5019	xe14	5038
ge20	5020	xe15	5039
ge21	5021	xe16	5040
ge22	5022	xe17	5041
ge23	5023	xe18	5042
ge24	5024	xe19	5043
ge25	5046	xe20	5044
ge26	5047	xe21	5045
ge27	5048		
ge28	5049		
ge29	5050		
ge30	5051		
ge31	5052		

ge32 5053
ge33 5054
ge34 5055
ge35 5056
ge36 5057
ge37 5058
ge38 5059
ge39 5060
ge40 5061
ge41 5062
ge42 5063
ge43 5064

Layer 2 Information

Layer 3 Information

3.1 Overview

This section provides FAQs on Layer 3 Protocols.

3.1.1 Interface Related

3.1.1.1 Explain the routerport concept. How to make a switchport to use an L3 interface

The routerport is an L3 interface in the OS associated to a switchport. The `no switchport` command on an interface makes it a routerport.

3.1.1.2 Explain `vlanx.y` interface and its working in linux and SRstackware®

The `vlanx.y` is an L3 interface which is used for vlan routing. When a vlan-aware bridge with vlan routing is enabled on a switchport interface, the SRstackware® adds an interface to the linux kernel as a network interface. The SRstackware then assigns a default MAC address to this interface and adds an entry in the L2 table as given below.

When this inter-vlan routing is enabled on a particular VLAN, then the ports with that vlan starts receiving untagged packets even though the port is a tagged only port.

Example:

```
mac=00:59:b3:8b:00:02 vlan=22 modid=0 port=28/cpu0 Static CPU L3
```

The above entry ensures that the packets coming to vlan 22 should be marked to CPU port where the network interface driver for `vlan1.22` receives these packets.

The packets can be received to `vlanx.y` interface only when a port has only one untagged vlan id `Y`, or the vlan tagged with `Y` along with many other tagged vlans.

This interface is primarily used for control traffic and cannot be used for data traffic. It shows low performance when it is used for data traffic.



If Inter VLAN routing is enabled on a particular VLAN, the ingress-filter should be disabled to allow untagged packets. Thus the Ingress Filter setting on an inter vlan routing cannot be independently changed.

Layer 3 Information

3.1.1.3 How to configure switch to send higher data traffic to an application running on management processor

For higher data traffic to management processor use the ethernet interface (ethx.y). Make sure the vlan x is configured on switch port and the BBP port on base or fabric chipsets on the blade. Refer to the blade-specific Programmer's Reference guide to find the port names which map to BBP port and switch ports of base and fabric.

Make sure that the ethx.y vlan interface is created in linux.

3.1.1.4 How does an SRstackware convert a switchport to a routerport

A switchport can be used only for L2 protocols. The L3 protocols cannot run on these interfaces. If you convert a switchport to a routerport by command `no switchport`, all the vlans and L2 configuration on that port are removed. None of L2 protocols can be enabled on that port. However, all the L3 protocols can run on this routerport. An interface `ge#` is added to the linux device drivers. It appears in `ifconfig` list.

The following is done in BCM to make it as a router port:

SRstackware adds a vlan starting from number 4023 and counting up the number of router ports configured.

Example:

```
vlan 4023 ports cpu,ge19 (0x0000000010080000), untagged ge19
(0x0000000000080000)
```

3.1.1.5 How to configure the DOT1X on VLAN Interface

VLAN router port network IP address access is decided based on authentication state of the actual physical interface, where it receives IP packet either destined or routed through the VLAN interface network IP address.

For example, if we receive a packet with destination IP address as 192.168.6.2 on vlan1.100 interface through physical interface ge22, the DOT1X should be configured as follows:

```
F140-3#show running-config interface vlan1.100
!
interface vlan1.100
 no switchport
 arp-ageing-timeout 3000
 ip address 192.168.6.2/24
 no shutdown
!
F140-3#
F140-3#show running-config interface ge22
```

```
!  
interface ge22  
  description Base RTM 1G Uplink 2 (ETH2)  
  bridge-group 1  
  switchport mode access  
  switchport access vlan 100  
  switchport mode access ingress-filter disable  
  no shutdown  
  dot1x port-control auto  
  dot1x port-control dir in  
!  
F140-3#
```

3.1.2 Static Routing

3.1.2.1 How to add a static arp entry

It is recommended to use the arp command available in the SRstackware CLI to configure an arp address as follows:

```
config# arp <ipaddr> <macaddr>
```

3.1.2.2 Is there a command that can register the MAC table and the ARP table in Static

Use the arp command to create a static ARP entry.

Example: (config)#arp 1.1.1.10 aaaa.bbbb.cccc

Use the bridge address command to register the MAC table.

Example: (config)#bridge 1 address aaaa.bbbb.cccc forward eth1

3.1.2.3 How to add a static route

It is mandatory to use the ip route command available in SRstackware CLI. The routes configured using the linux command are not reflected in the protocol operation. Use the below command.

```
config# ip route <destination network> <nextthop IP/interface name>
```

Layer 3 Information

3.1.2.4 How to add a static route with next hop address directly to FIB (without slow path)

The FIB entry is added only after the MAC address for the next hop entry is resolved. But if you know the next hop entry then you can add the static entry and avoid the slow path. Use the below commands.

```
config# ip route <destination network> <nexthop IP>
config# arp <nexthop ip> <macaddr>
```

3.1.2.5 How to add a static route with just interface name (without next hop) directly to FIB (without slow path)

Configure the ip address of the destination network so that it will be a connected route.

```
(config)# interface ge13
(config-if)# ip address 20.20.20.5
```

Otherwise the packet goes to the linux and you should enable ip forwarding and then the route will be added to L3 table of BCM. The interface name can be a router port or a vlan interface.

3.1.2.6 How to add static route with next hop address on a network configured on a vlan.x.y interface (without slow path)

1. Configure the IP address of the network that is part of next hop address on the vlan interface. Configure a virtual MAC address on this vlan interface.

Example:

```
(config)# interface vlan1.21
(config-if)# mac-address xxxx.xxxx.xxxx
config-if)# ip address 20.20.20.5
```

2. Configure the destination route with nexthop IP address

Example:

```
(config)# ip route 40.40.40.0/24 20.20.20.6
```

3. Configure the static arp entry for the next hop IP address

Example:

```
(config)# arp 20.20.20.6 000a.4556.6778
```

4. Configure this MAC address on the interface corresponding to vlan 21. Enable the interface.

Example:

```
(config)#interface ge13
(config-if)# no shutdown
(config)# bridge 1 address 000a.4556.6778 forward ge13 vlan 21
```

Assuming ge13 in the same bridge and same vlanid.

3.1.3 QoS

3.1.3.1 Why some of the class-maps throw error when configured at the same time

Class Map defines criteria to match against a specific traffic flow to further classify it.

There are six such match criteria. They are ip-dscp, access-group, ip-precedence, mpls exp-bit and layer4<UDP/TCP ports>, traffic-type.

All these types are parameters for match criteria. Of all these match types, only one match criteria can be defined for a single class-map. However, in addition to above types, Match VLAN command is used to specify VLAN ID while defining any match criteria together in a single class map. Only match vlan and vlan-range can be configured together with any one of the above match criteria together in a single class map. No other match criteria can be configured together in a single class map.

3.1.4 Access Control Lists and Match lists

3.1.4.1 What is the difference between ip-access-list and access-list

The `access-list` command configures an access list for filtering packets. This filtering list will be used by all protocol modules. The RIP and OSPF are the protocols that currently use `access-list`.

The `ip-access-list` command creates an IP access-control list (ACL), based on the source address, or creates an IP extended ACL, based on the source and destination address. This command is related to QoS, and until QoS is enabled using `mls qos enable`, an IP access list cannot be created.

3.1.4.2 What are Match lists

The Match lists configuration helps the user to match a particular field in the packet header to a value.

Layer 3 Information

3.1.4.3 What are rules and actions

A rule command combines a match list with an action. This helps you to configure the switch to take a specific action if the packet matches the match list criteria, configured using the match list commands.

3.1.4.4 What is the difference between outervlan and innervlan in match L2 param criteria

Use outervlan qualifier for normal/single tagged vlan packet flow. The innervlan qualifier should be used only in case of double tagged packets.

3.1.4.5 Are there any limitations to match lists

There are some limitations to match lists usage in the initial release. This section will be modified as and when the limitations change. This information will also be available in the release notes.

- When a match list is configured with a match criteria of double tag vlan then change priority action cannot be assigned.
- Multiple match lists cannot be associated to one action in the same command. Need to create separate match lists for an action.
- The match list or an action cannot be modified while the rule is installed. The rule should be deleted and added again.
- Match list ID is unique in the blade. The match list cannot be reused for other units. The Match lists configured and not associated to rule are not persistent.
- A maximum of five actions can be specified per match list entry.

3.1.4.6 When multiple rules exist with the same set of qualifiers, which of these rules will be applied

The rule with the lowest entry-id at BCM is executed when multiple rules exist with the same set of qualifiers.

PENGUIN[™]

SOLUTIONS 

Penguin Solutions is a trade name used by SMART Embedded Computing, Inc., a wholly owned subsidiary of SMART Global Holdings, Inc. Penguin Edge is a trademark owned by Penguin Computing, Inc., a wholly owned subsidiary of SMART Global Holdings, Inc. All other logos, trade names, and trademarks are the property of their respective owners. ©2022 SMART Embedded Computing, Inc.