1911

FAQ Tracking List :

Edgecore Technical Support FAQ Tracking List (20201105).xlsx

Getting Started

Backup and Restore default configuration (/etc/sonic /config_db.json)

Example:

Reference model:

- Switch model name: AS7326-56X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20201123_130028_ec202006_74

Procedure :

Method 1: Backup and Restore default configuration manually

Backup default configuration "config_db.json" after SONiC installation

```
admin@sonic:~$ sudo cp /etc/sonic/config_db.json /etc/sonic/config_db.json.bk
```

Restore default configuration

Step 1. Replace current config_db.json by backup config

```
admin@sonic:~$ sudo cp /etc/sonic/config_db.json.bk /etc/sonic/config_db.json
```

Step 2. config reload or power cycle the switch

```
admin@sonic:~$sudo config reload -y
```

Method 2: Use sonic-cfggen command to generate default configuration which is identical to the file generated after ONIE installation:

```
sudo sonic-cfggen -H -p /usr/share/sonic/device/$Platform/platform.json --preset t1 -k $HwSKU > ~/default.log
```

Caution: You can get the values of $Platform and $HwSKU from "show platform" or "show version"

```
admin@sonic:~$ show platform summary
Platform: x86_64-accton_as7326_56x-r0
HwSKU: Accton-AS7326-56X
ASIC: broadcom
```

Step 1: Generate the default config_db.json to the home directory.

```
admin@sonic:~$ sudo sonic-cfggen -H -p /usr/share/sonic/device/x86_64-accton_as7326_56x-r0/platform.json --
preset t1 -k Accton-AS7326-56X > ~/default.log
```

Step 2: Move the "default.log" to the correct position.

```
admin@sonic:~$ sudo cp default.log /etc/sonic/config_db.json
```

Step 3. config reload or power cycle the switch

```
admin@sonic:~$sudo config reload -y
```

Method 3: Upload and Execute the script restore.sh on the SONiC

The script generates default configuration *config_db.json* automatically. Reboot the switch to take effect.

Step 1: Set the management IP

```
admin@sonic:~$ sudo ip address add 188.188.98.21/16 dev eth0
```

Step 2: Upload the Script to the SONiC

```
admin@sonic:~$ sudo scp root@188.188.36.36:/root/restore.sh ~/
```

Step 3: Change permission of script.

```
admin@sonic:~$ chmod +x restore.sh
```

Step 4: Running the script.

```
admin@sonic:~$ ./restore.sh
```

Step 5. config reload or power cycle the switch

```
admin@sonic:~$sudo config reload -y
```

Example:

```
admin@sonic:~$ ./restore.sh
Get HwSKU and Platfrom from the database
/etc/sonic/config_db.json is restored to default

admin@sonic:~$ sudo config reload -y
```

Edgecore SONiC Installation & Upgrade

1. Edgecore SONiC Installation via TFTP/HTTP

Topology: Connect TFTP and/or HTTP server to switch management port

Caution: USB installation is **only supported** in firmware version **SONiC.Edgecore-SONiC_20200722_070543_ec201911_141 and later**. For older firmware version installation, please use TFTP/HTTP procedure below. In older firmware version, an known issue related to GRUB menu update will cause root file system crash (refer to PR #4443).

Procedure :

Step 1. Enter the ONIE install mode

Note: Switch shall automatically enter the ONIE install mode if there's no NOS installed yet.

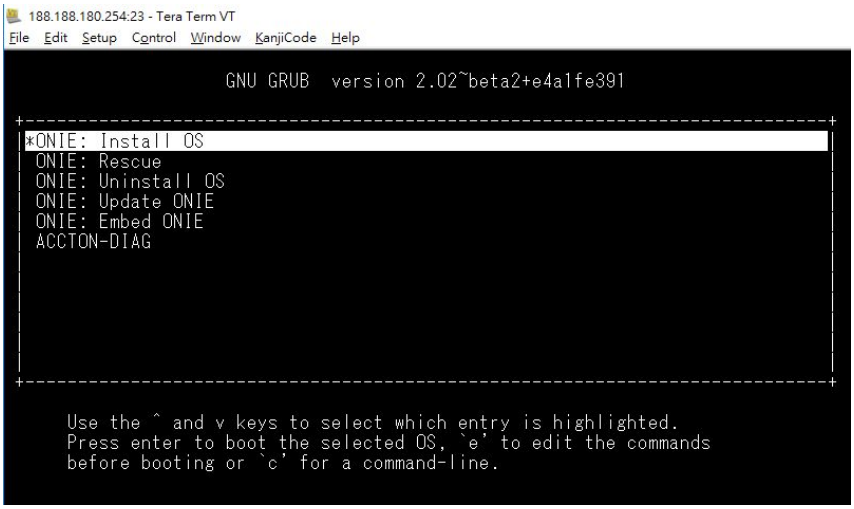

Step Ex. Stop the ONIE Service Discovery

```
ONIE:/ # onie-discovery-stop
discover: installer mode detected.
Stopping: discover... done.
```

Since switch will automatically start the ONIE Service Discovery, this command can make user types easily.

Note: This is not a necessary command. It won't affect installation no matter user execute it or not.

Step 2. Setup the ip address binding to switch management port

```
ONIE:/ # ifconfig eth0 192.168.1.2 netmask 255.255.255.0
```

Step 3. Install the image from remote URL via HTTP or TFTP

```
ONIE:/ # onie-nos-install tftp://192.168.1.3/Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178.bin
```

If the installation is successful, the device will reboot automatically and boot-up with SONiC.

After finish the installation, user might need to check the following two articles.

1. Username and Password

2. ecSONiC License Installation (License is no longer required from *SONiC.Edgecore-SONiC 201911.2 2020 Aug* version and later)

2. Edgecore SONiC Installation via USB

Procedure :

Step 1. Copy Edgecore SONiC installer file to USB flash drive with a name "**onie-installer-x86_64**".

Step 2. Reboot switch and plug in USB drive to switch

Step 3. Enter ONIE install mode

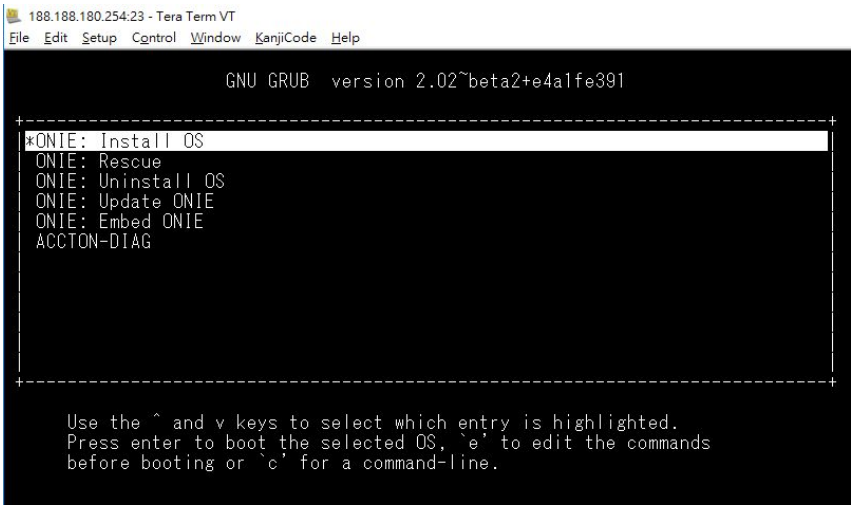

Step 4. Wait for ONIE to discover USB and start automatically uploading new firmware via USB drive.

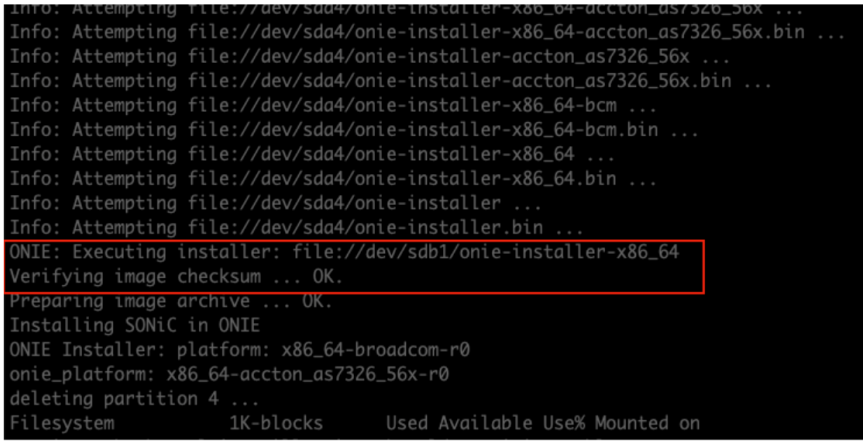

If the installation is successful, the device will reboot automatically and boot-up with SONiC.

After finish the installation, user might need to check the following two articles.

1. Username and Password

2. ecSONiC License Installation (License is no longer required from *SONiC.Edgecore-SONiC 201911.2 2020 Aug* version and later)

3. Edgecore SONiC Upgrade

Procedure :

Step 0. Check the image

```
admin@sonic:~$ sudo sonic_installer list
Current: SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
Next: SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
Available:
SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
```

Step 1. Set the manangement IP (Refer to Management and front port IPv4/IPv6 Address)

```
admin@sonic:~$ sudo ip address add 192.168.1.2/24 dev eth0
```

Step 2. Upgrade image from remote HTTP server

```
admin@sonic:~$ sudo sonic_installer install http://192.168.1.3Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178.bin -y
```

Caution: It could only use HTTP to upgrade.

Step 3. Check the images.

```
admin@sonic:~$ sudo sonic_installer list
Current: SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
Next: SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Available:
SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
```

Note: The "Next" boot-up image shall be set with new image automatically.

Step 4. Reboot device

```
admin@sonic:~$ sudo reboot
```

Step 5. Check the images status

```
admin@sonic:~$ sudo sonic_installer list
Current: SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Next: SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
Available:
SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
```

4. Change default boot-up image

Procedure :

Step 0. Check the images status

```
admin@sonic:~$ sudo sonic_installer list
Current: SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Next: SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
Available:
SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
```

Step 1. Change the default boot-up image

```
admin@sonic:~$ sudo sonic_installer set_default SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Command: grub-set-default --boot-directory=/host 1
```

Step 2. Check the image status

```
admin@sonic:~$ sudo sonic_installer list
Current: SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Next: SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Available:
SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
```

Then the "Next" boot-up image changed.

Extra: Choose the boot-up image in GRUB menu

Caution: You can choose the boot-up image in GRUB menu directly, but this way won't change the default boot-up image.

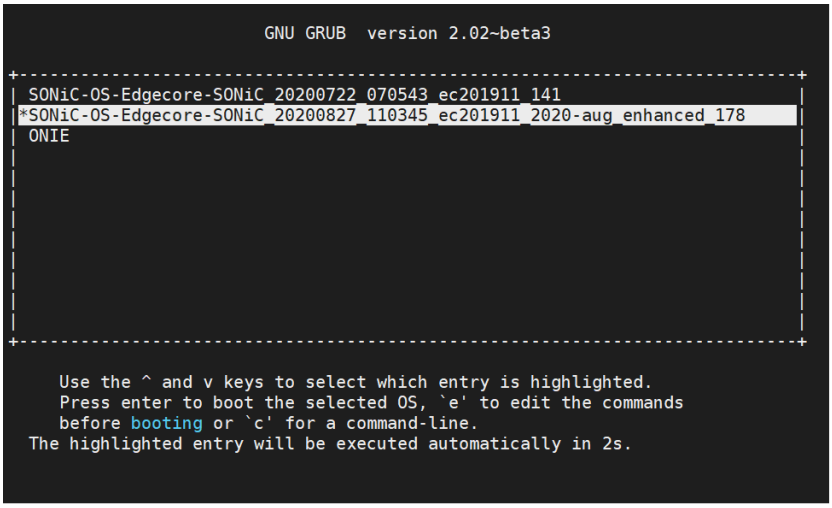

5. Remove installed SONiC image from SONiC installer list

Procedure :

Step 0. Check the images status

```
admin@sonic:~$ sudo sonic_installer list
Current: SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Next: SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Available:
SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141
```

Step 1. Remove the image

```
admin@sonic:~$ sudo sonic_installer remove SONiC-OS-Edgecore-SONiC_20200722_070543_ec201911_141 -y
```

Step 2. Check the images status

```
admin@sonic:~$ sudo sonic_installer list
Current: SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Next: SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
Available:
SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
```

JSON format & /etc/sonic/config_db.json

Example:

1. JSON Data Types and Introduction
2. Validate /etc/sonic/config_db.json in SONiC (TBD)

Reference model :

- Switch model name: All
- Edgecore SONiC version: All

Caution :

1. Content in /etc/sonic/config_db.json can be considered as starting config (refer to this article). It follows standard JSON format
2. The wrong JSON format in /etc/sonic/config_db.json may not apply the setting and/or cause the container not to work.

JSON Data Types:

- a string
- a number
- an object
- an array
- a boolean
- null

JSON Strings

Strings in JSON must be written with double quotes:

```
"platform"
```

JSON Numbers

Numbers must be an integer or a floating point.

```
1234567890 or 0.15
```

JSON Objects

1. JSON objects are surrounded by curly braces {}.
2. JSON objects are written in key/value pairs.
3. Keys must be strings, and values must be a valid JSON data type (string, number, object, array, boolean or null).
4. Keys and values are separated by a colon.
5. Each key/value pair is separated by a comma.

Caution: Object as values in JSON must follow the same results as JSON objects.

```
{"platform": "x86_64-accton_as7816_64x-r0"}
```

Example:

```
Object = {Key: Value}
```

Key must be strings. so, key must be written with double quotes

Value can be an object. Most of cases in SONiC (/etc/sonic/config_db.json), values are objects.

```
Object = {Key1: Value1, Key2: Value2}

       = {Key1: {key3: value3, key4: value4}, Key2: {key5: value5}}
```

i.e Value1 and Value2 are objects that consist of key3/value3, key4/value4 and key5/value5

Value1 = {key3: value3, key4: value4}

Value2= {key5: value5}

JSON Arrays

Values in JSON can be arrays. Array values must be of type string, number, object, array, boolean or *null*.

```
"ports": ["Ethernet0", "Ethernet1"]
```

JSON Booleans

Values in JSON can be true/false.

```
{"disable": true}
```

JSON null

Values in JSON can be null

```
{"platform": null}
```

Note: JSON Whitespace (Space, Horizontal tab, Line feed or New line or Carriage return)

Whitespace outside of string literals is ignored. i.g Insignificant whitespace may be included anywhere except within JSON Numbers and JSON String. Numbers can't have whitespace inside and strings would be interpreted as whitespace in the string or cause an error.

i.e. These two examples shown below are the same.

```
"BGP_NEIGHBOR":{"10.0.0.1":{"asn":"65200","holdtime":"180","keepalive":"60","local_addr":"10.0.0.0","name":"ARISTA01T2","nhopself":
0,"rrclient":0}}
```

```
"BGP_NEIGHBOR": {
"10.0.0.1": {
"asn": "65200",
"holdtime": "180",
"keepalive": "60",
"local_addr": "10.0.0.0",
"name": "ARISTA01T2",
"nhopself": 0,
"rrclient": 0
}
}
```

Validate /etc/sonic/config_db.json in SONiC(TBD)

Procedure :

Management and front port IPv4/IPv6 Address

Example 1: Configure IP address on management port "eth0"

Reference model:

- Switch model name: AS7326-56X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200722_070543_ec201911_141

Default setting:

1. DHCP address on management port"eth0"

Procedure :

Steps 1. Configure IP address on management port "eth0"

```
admin@sonic:~$ sudo config interface ip add eth0 188.188.97.16/16 188.188.1.1
admin@sonic:~$ sudo config interface ip add eth0 2001::8/64 2001::1
```

Steps 2. Check IP address

```
admin@sonic:~$ show ip interfaces
Interface   Master   IPv4 address/mask   Admin/Oper   BGP Neighbor   Neighbor
----------- -------- ------------------- ------------ -------------- ---------
Loopback0            10.1.0.1/32         up/up        N/A            N/A
docker0              240.127.1.1/24      up/down      N/A            N/A
eth0                 188.188.97.16/16    up/up        N/A            N/A
lo                   127.0.0.1/8         up/up        N/A            N/A

admin@sonic:~$  show ipv6 interfaces
Interface   Master   IPv6 address/mask                      Admin/Oper   BGP Neighbor   Neighbor IP
----------- -------- -------------------------------------- ------------ -------------- -------------
Bridge               fe80::4020:6aff:fee4:b0a3%Bridge/64    up/down      N/A            N/A
Ethernet0            fe80::ba6a:97ff:fee2:479c%Ethernet0/64 up/up        N/A            N/A
Ethernet1            fe80::ba6a:97ff:fee2:479c%Ethernet1/64 up/up        N/A            N/A
Ethernet48           fe80::ba6a:97ff:fee2:479c%Ethernet48/64 up/up       N/A            N/A
Ethernet52           fe80::ba6a:97ff:fee2:479c%Ethernet52/64 up/up       N/A            N/A
Loopback0            fe80::243e:e6ff:fee7:5ca1%Loopback0/64 up/up        N/A            N/A
eth0                 2001::8/64                             up/up        N/A            N/A
                     fe80::ba6a:97ff:fee2:479c%eth0/64
lo                   ::1/128                                up/up        N/A            N/A

admin@sonic:~$ show management_interface address
Management IP address = 188.188.97.16/16
Management Network Default Gateway = 188.188.1.1
Management IP address = 2001::8/64
Management Network Default Gateway = 2001::1
```

Note: The "show management_interface address" command and "BGP Neighbor information" do not support in ecSONiC-201904 version.

Steps 3. Save the setting to config_db.json

```
admin@sonic:~$ sudo config save -y
```

Steps 4. Check IP address setting on config_db.json

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{
    ...omitted
    "MGMT_INTERFACE": {
                "eth0|188.188.97.16/16": {
                        "gwaddr": "188.188.1.1"
                },
                "eth0|2001::8/64": {
                        "gwaddr": "2001::1"
                }
        },
    ...omitted
}
```

Example 2: Configure IP address on front port (aka physical interface).

Procedure :

Caution: Please make sure the port name is correct. Otherwise it does not work.

Since there's no warming messages for wrong command.

Steps 1. Configure IP address on front port. i.e Ethernet0

```
admin@sonic:~$ sudo config interface ip add Ethernet0 192.168.1.1/24
admin@sonic:~$ sudo config interface ip add Ethernet0 2002::1/64
```

Steps 2. Check IP address

```
admin@sonic:~$ show ip interfaces
Interface    IPv4 address/mask    Admin/Oper   BGP Neighbor   Neighbor IP
-----------  -------------------  -----------  -------------  -------------
Ethernet0    192.168.1.1/24       up/down      N/A            N/A
docker0      240.127.1.1/24       up/down      N/A            N/A
eth0         188.188.97.32/16     up/up        N/A            N/A
lo           127.0.0.1/8          up/up        N/A            N/A
             10.1.0.1/32

admin@sonic:~$  show ipv6 interfaces
Interface    IPv6 address/mask                      Admin/Oper   BGP Neighbor   Neighbor IP
-----------  -------------------------------------  -----------  -------------  -------------
Bridge       fe80::f4f6:a4ff:fe97:3c88%Bridge/64    up/up        N/A            N/A
Ethernet0    2002::1/64                             up/down      N/A            N/A
dummy        fe80::90cd:eff:fe7b:359f%dummy/64      up/up        N/A            N/A
eth0         fe80::aa2b:b5ff:fe9d:d7db%eth0/64      up/up        N/A            N/A
lo           ::1/128                                up/up        N/A            N/A
```

Steps 3. Save the setting to config_db.json

```
admin@sonic:~$ sudo config save -y
```

Steps 4. Check IP address setting on config_db.json

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{
    ...omitted
    "INTERFACE": {
        "Ethernet0|192.168.1.1/24": {},
        "Ethernet0|2002::1/64": {},
    ...omitted
    },
    ...omitted
}
```

Reset Password if forget username and password

Example: Reset Password

Reference model:

- Switch model name: AS7726-32X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178

Procedure :

Step 1: Reboot system and press “e” to edit SONiC grub option (e.g SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178)

```
                         GNU GRUB  version 2.02~beta3

+----------------------------------------------------------------------------+
|*SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178     |
| ONIE                                                                       |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
+----------------------------------------------------------------------------+

     Use the ^ and v keys to select which entry is highlighted.
     Press enter to boot the selected OS, `e' to edit the commands
     before booting or `c' for a command-line.
```

```
                         GNU GRUB  version 2.02~beta3

+----------------------------------------------------------------------------------------------+
|setparams 'SONiC-OS-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178'            |
|                                                                                              |
|        search --no-floppy --label --set=root SONiC-OS                                        |
|        echo    'Loading SONiC-OS OS kernel ...'                                              |
|        insmod gzio                                                                           |
|        if [ x = xxen ]; then insmod xzio; insmod lzopio; fi                                  |
|        insmod part_msdos                                                                     |
|        insmod ext2                                                                           |
|        linux   /image-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178/boot/vml\|
|inuz-4.9.0-9-2-amd64 root=/dev/sda4 rw console=tty0 console=ttyS0,115200n8 \                  |
|quiet                  net.ifnames=0 biosdevname=0                 loop=ima\                  |
|ge-Edgecore-SONiC_20200116_051623_ec201904_128/fs.squashfs loopfstype=squas\                  |v
+----------------------------------------------------------------------------------------------+

     Minimum Emacs-like screen editing is supported. TAB lists
     completions. Press Ctrl-x or F10 to boot, Ctrl-c or F2 for
     a command-line or ESC to discard edits and return to the GRUB menu.
```

Step 2: Add init=/bin/bash as shown below,

```
 +----------------------------------------------------------------------------------------------+
|        insmod gzio                                                                           |^
|        if [ x = xxen ]; then insmod xzio; insmod lzopio; fi                                  |
|        insmod part_msdos                                                                     |
|        insmod ext2                                                                           |
|        linux   /image-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178/boot/vml\|
|inuz-4.9.0-9-2-amd64 root=/dev/sda4 rw console=tty0 console=ttyS0,115200n8 \                  |
|init=/bin/bash quiet                  net.ifnames=0 biosdevname=0          \                  |
|       loop=image-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178/fs.squashfs l\|
|oopfstype=squashfs                                    apparmor=1 securit\                     |
|y=apparmor varlog_size=4096 usbcore.autosuspend=-1 tg3.short_preamble=1 tg3\                  |
|.bcm5718s_reset=1                                                                             |
|        echo    'Loading SONiC-OS OS initial ramdisk ...'                                     |v
+----------------------------------------------------------------------------------------------+

     Minimum Emacs-like screen editing is supported. TAB lists
     completions. Press Ctrl-x or F10 to boot, Ctrl-c or F2 for
     a command-line or ESC to discard edits and return to the GRUB menu.
```

Note: This insertion will be restored after rebooting. Therefore, it's not necessary to edit grub option again.

Step 3: Press Ctrl+x to boot

```
Loading SONiC-OS OS kernel ...
Loading SONiC-OS OS initial ramdisk ...
[ 2.708038] irq 16: nobody cared (try booting with the "irqpoll" option)
[ 2.714804] handlers:
[ 2.717210] [<ffffffff94c427b0>] serial8250_interrupt
[ 2.722539] Disabling IRQ #16
[ 3.647833] ixgbe 0000:08:00.1: HW Init failed: -17
bash: cannot set terminal process group (-1): Inappropriate ioctl for device
bash: no job control in this shell
root@(none):/#
```

Step 4: Now you're able to modify the password by linux command. Username root as example.

```
root@(none):/# passwd root
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
```

Step 5: Reboot the switch by linux command or Power cycle the switch

```
root@(none):/# reboot -f
```

Reference: SONiC Github https://github.com/Azure/SONiC/blob/master/doc/SONiC-User-Manual.md#23-how-to-reset-password

SONiC overview

What is SONiC?

https://azure.github.io/SONiC/

SONiC Configuration

SONiC supports 2 methods as below,

1. SONiC CLI (Command line interface)
2. config_db.json file(Edit /etc/sonic/config_db.json)

SONiC setting relation

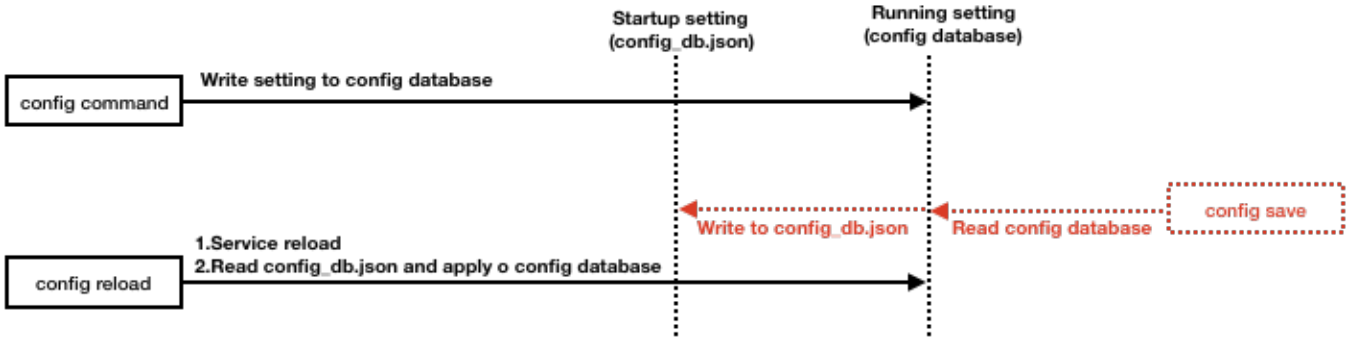

SONiC Architecture

https://github.com/Azure/SONiC/wiki/Architecture

SONiC frr routing relation

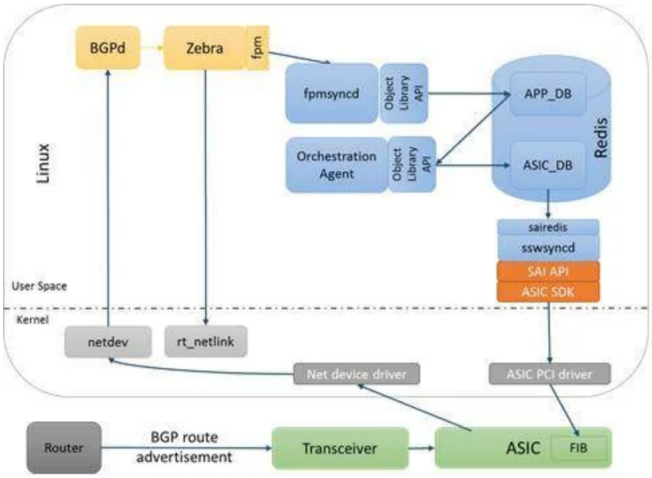

Username and Password

Username and Password

Default username is admin and password is YourPaSsWoRd

How to modify password in SONiC?

- Linux command: passwd
- admin@sonic:~$ passwd -h
  Usage: passwd [options] [LOGIN]

  Options:
  -a, --all report password status on all accounts
  -d, --delete delete the password for the named account
  -e, --expire force expire the password for the named account
  -h, --help display this help message and exit
  -k, --keep-tokens change password only if expired
  -i, --inactive INACTIVE set password inactive after expiration
  to INACTIVE
  -l, --lock lock the password of the named account
  -n, --mindays MIN_DAYS set minimum number of days before password
  change to MIN_DAYS
  -q, --quiet quiet mode
  -r, --repository REPOSITORY change password in REPOSITORY repository
  -R, --root CHROOT_DIR directory to chroot into
  -S, --status report password status on the named account
  -u, --unlock unlock the password of the named account
  -w, --warndays WARN_DAYS set expiration warning days to WARN_DAYS
  -x, --maxdays MAX_DAYS set maximum number of days before password
  change to MAX_DAYS

  admin@sonic:~$

For example:

```
admin@sonic:~$ sudo passwd admin
Enter new UNIX password: xxxx << enter new password here
Retype new UNIX password: xxxx << enter new password here
passwd: password updated successfully
```

SUDO

Default sudo password

Default, when you access command by sudo group, system won't request password.

```
admin@sonic:~$ sudo sed -n "50p" /etc/sudoers
%sudo ALL=(ALL:ALL) NOPASSWD: ALL
```

How to change the interactive mode of sudo group in SONiC?

If you want to request password every time when you run sudo group command. You can modify it as following.

```
admin@sonic:~$ sudo visudo
%sudo ALL=(ALL:ALL) ALL
```

Verify JSON format of /etc/sonic/config_db.json in SONiC

Example: Verify JSON format

Reference model:

- Switch model name: AS7726-32X,AS7816-64X,AS7326-56X
- Edgecore SONiC version:
  - SONiC.Edgecore-SONiC_20200116_051623_ec201904_128(AS7816-64X),
  - SONiC.Edgecore-SONiC_20200103_052306_ec201904_108(AS7726-32X,AS7326-56X)

Command: sonic-cfggen

```
#sonic-cfggen -j /etc/sonic/config_db.json --print-data
```

-j JSON : indicate json file name

--print-data : print all data

Restriction:

- "sonic-cfggen" is for JSON format verificaton only. It's NOT for parameters verification.

Example:

Wrong JSON format on the /etc/sonic/config_db.json

```
 1 {
 2   "DEVICE_METADATA": {
 3       "localhost": {
 4           "bgp_asn": "4200000002
 5           "docker_routing_config_mode": "split"
 6           "hostname": "AS7816-64X",
 7           "hwsku": "Accton-AS7816-64X",
 8           "mac": "68:21:5f:a7:b1:10",
 9           "platform": "x86_64-accton_as7816_64x-r0",
 10          "type": "LeafRouter",
 11      }
 12      }
 13  },
 omitted...
}
```

Note: wrong format with red print

Step 1. Use "sonic-cfggen" command to verify /etc/sonic/config_db.json

```
admin@AS7816-64X:~$ sonic-cfggen -j /etc/sonic/config_db.json --print-data
Traceback (most recent call last):
 File "/usr/local/bin/sonic-cfggen", line 280, in <module>
 main()
 File "/usr/local/bin/sonic-cfggen", line 226, in main
 deep_update(data, FormatConverter.to_deserialized(json.load(stream)))
 File "/usr/lib/python2.7/json/__init__.py", line 291, in load
 **kw)
 File "/usr/lib/python2.7/json/__init__.py", line 339, in loads
 return _default_decoder.decode(s)
 File "/usr/lib/python2.7/json/decoder.py", line 364, in decode
 obj, end = self.raw_decode(s, idx=_w(s, 0).end())
 File "/usr/lib/python2.7/json/decoder.py", line 380, in raw_decode
 obj, end = self.scan_once(s, idx)
ValueError: Invalid control character at: line 4 column 36 (char 85)
```

Note: Get error. It indicates which line of the file is incorrect.

Step 2. Refer to JSON format and correct the file.

Step 3. Use "sonic-cfggen" command to verify until no errors.

```
admin@AS7816-64X:~$ sonic-cfggen -j /etc/sonic/config_db.json --print-data
{
 "DEVICE_METADATA": {
 "localhost": {
 "bgp_asn": "4200000002",
 "docker_routing_config_mode": "split",
 "hostname": "AS7816-64X",
 "hwsku": "Accton-AS7816-64X",
 "mac": "68:21:5f:a7:b1:10",
 "platform": "x86_64-accton_as7816_64x-r0",
 "type": "LeafRouter"
 }
 },
omitted...
}
```

Warm Reboot & Fast Reboot

Example: Warm Reboot

Reference model:

- Switch model name: AS7726-32X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178

Note: Warm reboot is to be able in-service restart and upgrade SONiC software without impacting the data plane.

https://github.com/Azure/SONiC/blob/master/doc/warm-reboot/SONiC_Warmboot.md#cold-restart-fallback

Note: /etc/sonic/config_db.json won't be reloaded.

Procedure :

Step 1. Upgrade SONiC software. (Refer to this article)

```
admin@sonic:~$ sudo sonic_installer install http://192.168.1.3/SONiC.Edgecore-
SONiC_20200827_110345_ec201911_2020-aug_enhanced_178.bin -y
```

Step 2. Warm-reboot to apply the new software version without impacting the data plane.

```
admin@sonic:~$  sudo warm-reboot

Warning: Stopping telemetry.service, but it can still be activated by:
telemetry.timer
Warning: docker.service changed on disk. Run 'systemctl daemon-reload' to reload units.
[ 465.253571] kexec_core: Starting new kernel
[ 0.175853] DMAR-IR: Failed to copy IR table for dmar0 from previous kernel
[ 6.035563] systemd[1]: [/etc/systemd/system/procdockerstatsd.service:9] Failed to parse service restart
specifier, ignoring: Always
[ 6.495522] rc.local[518]: + sonic-cfggen -y /etc/sonic/sonic_version.yml -v build_version
[ 6.661165] kdump-tools[516]: Starting kdump-tools: no crashkernel= parameter in the kernel cmdline ... failed!
[ 7.729980] rc.local[518]: + SONIC_VERSION=Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
[ 7.748462] rc.local[518]: + FIRST_BOOT_FILE=/host/image-Edgecore-SONiC_20200827_110345_ec201911_2020-
aug_enhanced_178/platform/firsttime
[ 7.772443] rc.local[518]: + logger SONiC version Edgecore-SONiC_20200827_110345_ec201911_2020-
aug_enhanced_178 starting up...
[ 7.800752] rc.local[518]: + [ ! -e /host/machine.conf ]
[ 7.816478] rc.local[518]: + . /host/machine.conf
[ 7.832472] rc.local[518]: + onie_arch=x86_64
[ 7.848425] rc.local[518]: + onie_bin=
[ 7.860455] rc.local[518]: + onie_boot_reason=install
[ 7.876437] rc.local[518]: + onie_build_date=2018-02-08T21:53+0800
[ 7.892424] rc.local[518]: + onie_build_machine=accton_as7726_32x
[ 7.908420] rc.local[518]: + onie_build_platform=x86_64-accton_as7726_32x-r0
[ 7.928438] rc.local[518]: + onie_cli_static_parms=
[ 7.944537] rc.local[518]: + onie_cli_static_url=http://188.188.99.1/ecsonic/Edgecore-
SONiC_20200827_110345_ec201911_2020-aug_enhanced_178.bin
[ 7.968547] rc.local[518]: + onie_config_version=1
[ 7.984437] rc.local[518]: + onie_dev=/dev/sda2
[ 8.000377] rc.local[518]: + onie_exec_url=http://188.188.99.1/ecsonic/Edgecore-
SONiC_20200827_110345_ec201911_2020-aug_enhanced_178.bin
[ 8.028579] rc.local[518]: + onie_firmware=auto
[ 8.048539] rc.local[518]: + onie_grub_image_name=grubx64.efi
[ 8.064809] rc.local[518]: + onie_initrd_tmp=/
[ 8.080919] rc.local[518]: + onie_installer=/var/tmp/installer
[ 8.101087] rc.local[518]: + onie_kernel_version=4.9.57
[ 8.116758] rc.local[518]: + onie_machine=accton_as7726_32x
[ 8.132647] rc.local[518]: + onie_machine_rev=0
[ 8.148679] rc.local[518]: + onie_partition_type=gpt
[ 8.164781] rc.local[518]: + onie_platform=x86_64-accton_as7726_32x-r0
[ 8.180460] rc.local[518]: + onie_root_dir=/mnt/onie-boot/onie
[ 8.196483] rc.local[518]: + onie_skip_ethmgmt_macs=no
[ 8.212787] rc.local[518]: + onie_switch_asic=bcm
[ 8.228562] rc.local[518]: + onie_vendor_id=259
[ 8.244418] rc.local[518]: + onie_version=2017.11.00.05
[ 8.260399] rc.local[518]: + program_console_speed
[ 8.291047] rc.local[518]: + cat /proc/cmdline
[ 8.311316] rc.local[518]: + grep -Eo console=ttyS[0-9]+,[0-9]+
[ 8.329711] rc.local[518]: + cut -d , -f2
[ 8.344571] rc.local[518]: + speed=115200
[ 8.356399] rc.local[518]: + [ -z 115200 ]
[ 8.368396] rc.local[518]: + CONSOLE_SPEED=115200
[ 8.384401] rc.local[518]: + sed -i s|\-\-keep\-baud .* %I| 115200 %I|g /lib/systemd/system/serial-getty@.
service
```

```
[ 8.404415] rc.local[518]: + systemctl daemon-reload
[ 8.420422] rc.local[518]: + [ -f /host/image-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
/platform/firsttime ]
[ 8.444406] rc.local[518]: + echo First boot detected. Performing first boot tasks...
[ 8.464849] rc.local[518]: First boot detected. Performing first boot tasks...
[ 8.484483] rc.local[518]: + [ -n ]
[ 8.496571] rc.local[518]: + [ -n x86_64-accton_as7726_32x-r0 ]
[ 8.512395] rc.local[518]: + platform=x86_64-accton_as7726_32x-r0
[ 8.528887] rc.local[518]: + [ -d /host/old_config ]
[ 8.544852] rc.local[518]: + mv -f /host/old_config /etc/sonic/
[ 8.560533] rc.local[518]: + rm -rf /etc/sonic/old_config/old_config
[ 8.576581] rc.local[518]: + touch /tmp/pending_config_migration
[ 8.593109] rc.local[518]: + touch /tmp/notify_firstboot_to_platform
[ 8.609068] rc.local[518]: + [ -d /host/image-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
/platform/x86_64-accton_as7726_32x-r0 ]
[ 8.632913] rc.local[518]: + dpkg -i /host/image-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
/platform/x86_64-accton_as7726_32x-r0/sonic-platform-accton-as7726-32x_1.1_amd64.deb
[ 8.662929] rc.local[518]: Selecting previously unselected package sonic-platform-accton-as7726-32x.
[ 8.685029] rc.local[518]: (Reading database ... 27986 files and directories currently installed.)
[ 8.705084] rc.local[518]: Preparing to unpack .../sonic-platform-accton-as7726-32x_1.1_amd64.deb ...
[ 8.724933] rc.local[518]: Unpacking sonic-platform-accton-as7726-32x (1.1) ...
[ 8.745034] rc.local[518]: Setting up sonic-platform-accton-as7726-32x (1.1) ...
[ 8.989943] rc.local[518]: Created symlink /etc/systemd/system/opennsl-modules.service.wants/as7726-32x-
platform-handle_mac.service  /lib/systemd/system/as7726-32x-platform-handle_mac.service.
[ 9.230029] rc.local[518]: Created symlink /etc/systemd/system/multi-user.target.wants/as7726-32x-platform-
init.service  /lib/systemd/system/as7726-32x-platform-init.service.
[ 9.536387] rc.local[518]: Created symlink /etc/systemd/system/multi-user.target.wants/as7726-32x-platform-
monitor-fan.service  /lib/systemd/system/as7726-32x-platform-monitor-fan.service.
[ 9.876695] rc.local[518]: Created symlink /etc/systemd/system/multi-user.target.wants/as7726-32x-platform-
monitor-psu.service  /lib/systemd/system/as7726-32x-platform-monitor-psu.service.
[ 10.182216] rc.local[518]: Created symlink /etc/systemd/system/multi-user.target.wants/as7726-32x-platform-
monitor.service  /lib/systemd/system/as7726-32x-platform-monitor.service.
[ OK ] Started Accton AS7726-32X Platform initialization service.
[ OK ] Started Accton AS7726-32X Platform Monitoring FAN service.
[ OK ] Started Accton AS7726-32X Platform Monitoring PSU service.
[ OK ] Started Accton AS7726-32X Platform Monitoring service.
[ 27.640112] rc.local[518]: + sync
[ 27.666331] rc.local[518]: + [ -n x86_64-accton_as7726_32x-r0 ]
[ OK ] Started /etc/rc.local Compatibility.
[ 27.684754] rc.local[518]: + [ -n ]
Starting Database container...
[ 27.712464] rc.local[518]: + mkdir -p /var/platform
[ OK ] Started Getty on tty1.
[ 27.748670] rc.local[518]: + firsttime_exit
[ OK ] Started Serial Getty on ttyS0.
[ 27.776564] rc.local[518]: + rm -rf /host/image-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
/platform/firsttime
[ OK ] Reached target Login Prompts.
[ 27.816588] rc.local[518]: + exit 0


Debian GNU/Linux 9 sonic ttyS0

sonic login:
```

Example: Fast Reboot

Reference model:

- Switch model name: AS7726-32X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178

Note: Fast-reboot is for updating the control plane, it could let the switch reboot up quickly (<= 25 seconds), so it impacts the data plane with minimum disruption.

https://github.com/Azure/SONiC/wiki/Fast-Reboot

Restriction:

Since the fast-reboot may take 25 seconds to boot the system, there are some restrictions.

- Fast-Reboot should use stale FIB (forwarding information base) information while control plane reboots
- Fast-Reboot must support at least 2000 hosts connected to SONiC VLAN interfaces
- Fast-Reboot must support at least 6000 ipv4 BGP routes and 3000 ipv6 /64 BGP routes
- LACP mode must be in SLOW mode for all LAG interfaces on a SONiC device

Procedure :

```
admin@sonic:~$ sudo fast-reboot

Warning: Stopping telemetry.service, but it can still be activated by:
telemetry.timer
Warning: docker.service changed on disk. Run 'systemctl daemon-reload' to reload units.
[ 393.491515] kexec_core: Starting new kernel
[ 0.176460] DMAR-IR: Failed to copy IR table for dmar0 from previous kernel
[ 5.964796] systemd[1]: [/etc/systemd/system/procdockerstatsd.service:9] Failed to parse service restart
specifier, ignoring: Always
[ 6.270328] rc.local[480]: + sonic-cfggen -y /etc/sonic/sonic_version.yml -v build_version
[ 6.740018] kdump-tools[482]: Starting kdump-tools: no crashkernel= parameter in the kernel cmdline ... failed!
[ 8.099533] rc.local[480]: + SONIC_VERSION=Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
[ 8.120516] rc.local[480]: + FIRST_BOOT_FILE=/host/image-Edgecore-SONiC_20200827_110345_ec201911_2020-
aug_enhanced_178/platform/firsttime
[ 8.144592] rc.local[480]: + logger SONiC version Edgecore-SONiC_20200827_110345_ec201911_2020-
aug_enhanced_178 starting up...
[ 8.168463] rc.local[480]: + [ ! -e /host/machine.conf ]
[ 8.184482] rc.local[480]: + . /host/machine.conf
[ 8.200522] rc.local[480]: + onie_arch=x86_64
[ 8.216455] rc.local[480]: + onie_bin=
[ 8.228525] rc.local[480]: + onie_boot_reason=install
[ 8.244416] rc.local[480]: + onie_build_date=2018-02-08T21:53+0800
[ 8.260544] rc.local[480]: + onie_build_machine=accton_as7726_32x
[ 8.280459] rc.local[480]: + onie_build_platform=x86_64-accton_as7726_32x-r0
[ 8.304252] rc.local[480]: + onie_cli_static_parms=
[ 8.324273] rc.local[480]: + onie_cli_static_url=http://188.188.99.1/ecsonic/Edgecore-
SONiC_20200827_110345_ec201911_2020-aug_enhanced_178.bin
[ 8.352241] rc.local[480]: + onie_config_version=1
[ 8.368245] rc.local[480]: + onie_dev=/dev/sda2
[ 8.388255] rc.local[480]: + onie_exec_url=http://188.188.99.1/ecsonic/Edgecore-
SONiC_20200827_110345_ec201911_2020-aug_enhanced_178.bin
[ 8.412229] rc.local[480]: + onie_firmware=auto
[ 8.428264] rc.local[480]: + onie_grub_image_name=grubx64.efi
[ 8.444250] rc.local[480]: + onie_initrd_tmp=/
[ 8.460249] rc.local[480]: + onie_installer=/var/tmp/installer
[ 8.476240] rc.local[480]: + onie_kernel_version=4.9.57
[ 8.492347] rc.local[480]: + onie_machine=accton_as7726_32x
[ 8.508602] rc.local[480]: + onie_machine_rev=0
[ 8.524474] rc.local[480]: + onie_partition_type=gpt
[ 8.540410] rc.local[480]: + onie_platform=x86_64-accton_as7726_32x-r0
[ 8.556396] rc.local[480]: + onie_root_dir=/mnt/onie-boot/onie
[ 8.572466] rc.local[480]: + onie_skip_ethmgmt_macs=no
[ 8.588490] rc.local[480]: + onie_switch_asic=bcm
[ 8.604323] rc.local[480]: + onie_vendor_id=259
[ 8.620399] rc.local[480]: + onie_version=2017.11.00.05
[ 8.636413] rc.local[480]: + program_console_speed
[ 8.653650] rc.local[480]: + cat /proc/cmdline
[ 8.674874] rc.local[480]: + grep -Eo console=ttyS[0-9]+,[0-9]+
[ 8.693641] rc.local[480]: + cut -d , -f2
[ 8.709598] rc.local[480]: + speed=115200
[ 8.728309] rc.local[480]: + [ -z 115200 ]
[ 8.740294] rc.local[480]: + CONSOLE_SPEED=115200
[ 8.756265] rc.local[480]: + sed -i s|\-\-keep\-baud .* %I| 115200 %I|g /lib/systemd/system/serial-getty@.
service
[ 8.776290] rc.local[480]: + systemctl daemon-reload
[ 8.796402] rc.local[480]: + [ -f /host/image-Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178
/platform/firsttime ]
[ 8.820264] rc.local[480]: + exit 0

Debian GNU/Linux 9 sonic ttyS0

sonic login:
```

Basic Configuration

ACL (Access Control List)

Example: ACL

Restriction:

1. There is no SONiC command to add ACL table and rules.
2. LAG's member port shall not be added to the ACL Tables, or will be considered as invalid configuration and return fail.
3. LAG ACL configurations will be automatically applied to all the LAG members.
4. 1K ACL rules for L3 ACL table (TBD)
5. 256 ACL rules for mirror (TBD)

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Procedure :

Step 1: Create an ACL table on config_db.json

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{
    ...
    "ACL_TABLE": {
        "ACL_ETH0": {
            "policy_desc": "drop_1.0",
            "ports": [
                "Ethernet0"
            ],
            "stage": "ingress",
            "type": "L3"
        }
    },
    ...
}
```

Note:

- ACL_ETH0 is name of this ACL
- Key "policy_desc" is for description. Value of policy_desc is JSON string "drop_1.0"
- Value of ports is JSON array consist of name of physical interface, LAG or VLAN interface. e.g

```
"ports": ["Ethernet0", "Ethernet1"]
```

```
"ports": ["Vlan10"]
```

- Key "stage" is for ACL direction. SONiC ACL supports only ingress direction. The value is only "ingress"
- If you do NOT configure "stage", it means it will match all of the packets. (TBD)
- Key "types" is the type of ACL. SONiC supports 3 types: L3, L3V6 and MIRROR. (About the MIRROR, please refer to this article.)

Step 2: Reload config and check ACL table by SONiC command.

```
admin@sonic:~$ sudo config reload -y
...omitted
admin@sonic:~$ show acl table
Name      Type    Binding    Description
--------  ------  ---------  -------------
ACL_ETH0  L3      Ethernet0  drop_1.0
```

Step 3: Create ACL rules on config_db.json

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{
    ...
    "ACL_RULE": {
        "ACL_ETH0|ACE_FORWARD": {
            "PACKET_ACTION": "FORWARD",
            "PRIORITY": "1",
            "IP_TYPE": "ANY"
     },
        "ACL_ETH0|ACE_DROP": {
            "PACKET_ACTION": "DROP",
            "PRIORITY": "2",
            "SRC_IP": "192.168.1.10/32",
            "L4_SRC_PORT": "53"
        }
    },
    ...
}
```

Note:

- Key "ACL_ETH0|ACE_FORWARD" and "ACL_ETH0|ACE_DROP" are the names of rules of ACL ACL_ETH0.
- If key "types" of ACL TABLE is L3 or L3V6
  - Key in ACL rule is "PACKET_ACTION"
  - The value of "PACKET_ACTION" is FORWARD or DROP.
- If key "types" of ACL table is MIRROR,
  - Key in ACL rule is "MIRROR_ACTION"
  - The value of "MIRROR_ACTION" is the name of mirror session
- The number of priority is higher, it means priority is high.
  Take the above example, the priority 2 will match first.
- Here are the values for "IP_TYPE":

ANY	Filter IPv4, IPv6, Ether type
IP	Filter IPv4, IPv6
NON_IP	Filter Ether type only
IPV4ANY	Filter IPv4 only
NON_IPv4	Filter IPv6, Ether type
IPV6ANY	Filter IPv6 only
NON_IPv6	Filter IPv4, Ether type
ARP	Filter ARP request, reply
ARP_REQUEST	Filter ARP request only
ARP_REPLY	Filter ARP reply only

  Caution: NON_IP, NON_IPv4, NON_IPv6 could not work now.
- Other keys(parameters):

1. DST_IP example: "DST_IP": "192.168.1.10/32"
2. SRC_IPV6 example: "SRC_IPV6": "2001::db:1"
   Caution: The type of the ACL table should be modified to "L3V6"
3. DST_IPV6 example: "DST_IPV6": "2001::db:2"
   Caution: The type of the ACL table should be modified to "L3V6"
4. ETHER_TYPE example: "ETHER_TYPE": "0x842" or "ETHER_TYPE": "2114"
5. L4_SRC_PORT example: "L4_SRC_PORT": "53"
6. L4_DST_PORT example: "L4_DST_PORT": "53"
7. IP_PROTOCOL example: "IP_PROTOCOL": "1"
8. L4_SRC_PORT_RANGE example: "L4_SRC_PORT_RANGE": "1028-4096"
9. L4_DST_PORT_RANGE example: "L4_DST_PORT_RANGE": "1028-4096"
10. ICMP_TYPE example: "ICMP_TYPE": "0"
11. ICMPV6_TYPE example: "ICMPV6_TYPE": "128"
    Caution: The type of the ACL table should be modified to "L3V6"
12. TCP_FLAGS example: "TCP_FLAGS": "16/255"
    Note:
    FIN = 0x01 "TCP_FLAGS": "1/255"
    SYN = 0x02 "TCP_FLAGS": "2/255"
    RST = 0x04 "TCP_FLAGS": "4/255"
    PSH = 0x08 "TCP_FLAGS": "8/255"
    ACK = 0x10 "TCP_FLAGS": "16/255"
    URG = 0x20 "TCP_FLAGS": "32/255"
    ECE = 0x40 "TCP_FLAGS": "64/255"
    CWR = 0x80 "TCP_FLAGS": "128/255"

Step 4. Check ACL rules by SONiC command

```
admin@sonic:~$ sudo config reload -y
...omitted
admin@sonic:~$ show acl rule

Table      Rule          Priority    Action    Match
---------  ------------  ----------  --------  ----------------------
ACL_ETH0   ACE_DROP      2           DROP      L4_SRC_PORT: 53
                                               SRC_IP: 192.168.1.10/32
ACL_ETH0   ACE_FORWARD   1           FORWARD   IP_TYPE: ANY
```

DHCP Relay

Example: DHCP Relay

Restriction:

1. There is no SONiC command for DHCP Relay

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Topology :

DHCP server belongs to VLAN 10 and connects to the switch Ethernet4.

DHCP client belongs to VLAN 20 and connects to the switch Ethernet2.

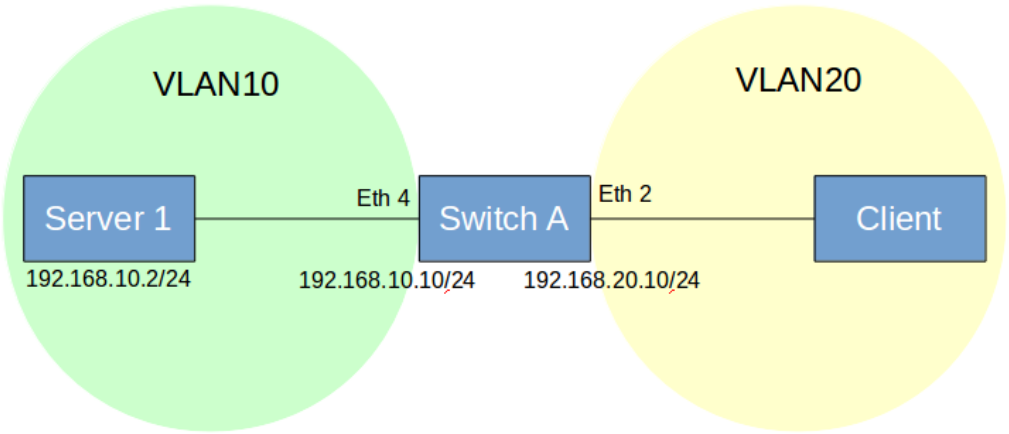

Procedure :

Step 1. Configure VLAN and IP address properly (Refer to this article)

Create VLAN10 with member Ethernet4 and binding the IP to VLAN10.

```
admin@sonic:~$ sudo config vlan add 10
admin@sonic:~$ sudo config vlan member add -u 10 Ethernet4
admin@sonic:~$ sudo config interface ip add Vlan10 192.168.10.10/24
```

Create VLAN20 with member Ethernet2 and binding the IP to VLAN20.

```
admin@sonic:~$ sudo config vlan add 20
admin@sonic:~$ sudo config vlan member add -u 20 Ethernet2
admin@sonic:~$ sudo config interface ip add Vlan20 192.168.20.10/24
```

Step 2. Save above setting to config_db.json

```
admin@sonic:~$ sudo config save -y
```

Step 3. Edit the file /etc/sonic/config_db.json

Add DHCP relay server IP address to VLAN 20 which DHCP client belongs to

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{
    omitted...
    "VLAN": {
        "Vlan10": {
            "members": [
                "Ethernet4"
            ],
            "vlanid": "10"
        },
        "Vlan20": {
            "dhcp_servers": [
                "192.168.1.10"
            ],
            "members": [
                "Ethernet2"
            ],
            "vlanid": "20"
        }
    },
    "VLAN_INTERFACE": {
        "Vlan10|192.168.10.10/24": {},
        "Vlan20|192.168.20.10/24": {}
    },
    "VLAN_MEMBER": {
        "Vlan10|Ethernet4": {
            "tagging_mode": "untagged"
        },
        "Vlan20|Ethernet2": {
            "tagging_mode": "untagged"
        }
    },
    omitted...
}
```

Step 4.config reload or power cycle the switch

```
admin@sonic:~$ sudo config reload -y
```

Step 5. Check the setting information by SONiC command.

```
admin@sonic:~$ show vlan brief
+-----------+------------------+-----------+----------------+-----------------------+
|   VLAN ID | IP Address       | Ports     | Port Tagging   | DHCP Helper Address   |
+===========+==================+===========+================+=======================+
|        10 | 192.168.10.10/24 | Ethernet4 | untagged       |                       |
+-----------+------------------+-----------+----------------+-----------------------+
|        20 | 192.168.20.10/24 | Ethernet2 | untagged       | 192.168.10.2          |
+-----------+------------------+-----------+----------------+-----------------------+
```

Result: VLAN20 DHCP client could obtain IP address from VLAN20 DHCP Server through the switch.

DNS

Reference model:

- Switch model name: AS5835-54X,AS7326-56X,AS7726-32X,AS7816-64X
- Edgecore SONiC version: All version

Procedure:

Step 1. Add nameserver ip (8.8.8.8 as example) to /etc/resolve.conf

```
admin@sonic:~$ sudo vi /etc/resolv.conf
```

```
# This file is managed by man:systemd-resolved(8). Do not edit.
#
# This is a dynamic resolv.conf file for connecting local clients directly to
# all known uplink DNS servers. This file lists all configured search domains.
#
# Third party programs must not access this file directly, but only through the
# symlink at /etc/resolv.conf. To manage man:resolv.conf(5) in a different way,
# replace this symlink by a static file or a different symlink.
#
# See man:systemd-resolved.service(8) for details about the supported modes of
# operation for /etc/resolv.conf.

#nameserver 10.5.0.2
nameserver 8.8.8.8
search ap-south-1.compute.internal
```

Caution:

- Make sure switch IP address and connectivity works well. (refer to Management IP address)

Result: The switch resolves domain name properly.

```
admin@sonic:~$ ping www.google.com
PING www.google.com (216.58.200.228) 56(84) bytes of data.
64 bytes from tsa03s01-in-f228.1e100.net (216.58.200.228): icmp_seq=1 ttl=55 time=2.66 ms
64 bytes from tsa03s01-in-f228.1e100.net (216.58.200.228): icmp_seq=2 ttl=55 time=2.68 ms
64 bytes from tsa03s01-in-f228.1e100.net (216.58.200.228): icmp_seq=3 ttl=55 time=2.65 ms
^C
--- www.google.com ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 2.656/2.668/2.682/0.043 ms
admin@sonic:~$
```

Everflow mirror session

Restriction:

- SONiC Everflow is similar to Cisco ERSPAN.
- SONiC does NOT support Local SPAN and Remote SPAN. (Definitions of SPAN/RSPAN /ERSPAN, please refer to Cisco Types of SPAN)
- Everflow supports mirror only ingress traffic of source interface. Because SONiC ACL supports only ingress direction(refer to this article).
- There is no SONiC command for add ACL Table for type MIRROR. (refer to this article)

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Example 1. Basic Everflow configuration

Procedure :

Step 1. Configure Switch IP address (refer to this article)

i.g 192.168.1.1/24 on Ethernet0

```
admin@sonic:~$ sudo config interface ip add Ethernet0 192.168.1.1/24
```

Step 2. Create a mirror session to specify source and destination IP address.

This command is used to add or remove mirroring sessions. Mirror session is identified by "session_name".

```
config mirror_session add [OPTIONS] <session_name> <src_ip> <dst_ip> <dscp> <ttl> [gre_type] [queue]
```

Note: session_name, src_ip, dst_ip, dscp, ttl are required for adding a new session. are_type and queue are optional.

```
admin@sonic:~$ sudo config mirror_session add ts1_everflow 192.168.1.1 192.168.3.35 63 255
```

Caution: src_ip is switch IP address. (refer to Step 1. Configure switch IP address)

Step 3. Check mirror_session by SONiC commands

```
admin@sonic:~$ show mirror_session
Name Status SRC IP DST IP GRE DSCP TTL Queue
------------ -------- ----------- ------------ ----- ------ ----- -------
ts1_everflow error 192.168.1.1 192.168.3.35 63 255
```

Caution/known issue: status of mirror session always display error.

This is a known issue and it is a display issue only. It will be fixed in new version.

https://github.com/Azure/sonic-buildimage/issues/3188

Step 4. Save above settings to /etc/sonic/config_db.json, since we will edit /etc/sonic/config_db.json on the following steps.

```
admin@sonic:~$ sudo config save -y
```

Step 5. Create ACL table and rules for traffic classification. (How to add ACL table and rules? Please refer to this article)

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{
    omitted...
    "ACL_TABLE": {
        "ACL_Mirror": {
            "policy_desc": "mirror",
            "type": "MIRROR",
            "stage": "ingress",
            "ports": ["Ethernet1"]
        }
    },
    "ACL_RULE": {
        "ACL_Mirror|ACE_Mirror": {
            "PRIORITY": "55",
            "IP_TYPE": "ipv4any",
            "MIRROR_ACTION": "ts1_everflow"
        }
    },
    omitted...
}
```

Note:

- The value of ACL type must be MIRROR
- The value of MIRROR_ACTION is the name of mirror session

Steps 6. Reload config or Reboot the switch

```
admin@sonic:~$ sudo config reload -y
```

Step 7. Check ACL table and rules

```
admin@sonic:~$ show acl table
Name Type Binding Description
---------- ------ --------- -------------
ACL_Mirror MIRROR Ethernet1 mirror
admin@sonic:~$ show acl rule
Table Rule Priority Action Match
---------- ---------- ---------- -------------------- ----------------
ACL_Mirror ACE_Mirror 55 MIRROR: ts1_everflow IP_TYPE: ipv4any
```

Caution: Have to make sure IP connectivity works properly between the switch and remote server. Otherwise,

Steps 8. Check the arp of monitor server is learned on arp table.

```
admin@sonic:~$ show arp
Address MacAddress Iface Vlan
------------ ----------------- ---------- ------
192.168.3.35 8c:ea:1b:30:da:4a Ethernet16 -
Total number of entries 1
```

Caution: Before switch learn the arp of monitor device to arp table, the ERSPAN won't work.

ACL Table example: Classify LLDP(0x88CC) and LACP(0x8809) packets

"ETHER_TYPE": "0x88cc" for LLDP

"ETHER_TYPE": "0x8809" for LACP

/etc/sonic/config_db.json

```
{
    omitted...
    "ACL_TABLE": {
        "ACL_Mirror": {
            "policy_desc": "mirror",
            "type": "MIRROR",
            "stage": "ingress",
            "ports": ["Ethernet8"]
        }
    },
    "ACL_RULE": {
        "ACL_Mirror|ACE_Mirror1": {
            "PRIORITY": "55",
            "ETHER_TYPE": "0x88cc",
            "MIRROR_ACTION": "ts1_everflow"
        },
        "ACL_Mirror|ACE_Mirror2": {
            "PRIORITY": "54",
            "ETHER_TYPE": "0x8809",
            "MIRROR_ACTION": "ts1_everflow"
        },
    },
    omitted...
}
```

LLDP (Link Layer Discovery Protocol)

Example:

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Restriction:

1. There is no SONiC command to disable LLDP. (Workaround: "sudo systemctl stop lldp.service". However, LLDP service restart after rebooting)
    - Start/Stop LLDP service: sudo systemctl [start|stop] lldp
    - Enable/Disable LLDP service by default: sudo systemctl [enable|disable] lldp

Default setting:

1. LLDP is enabled
2. LLDP PDU interval : 30 seconds
3. Support LLDP extended MIB (lldpremtable, lldplocporttable, lldpremmanaddrtable, lldplocmanaddrtable, lldpLocalSystemData). Refer to SNMP configuration

Procedures:

Step 1. check LLDP table

```
admin@sonic:~$ show lldp table
Capability codes: (R) Router, (B) Bridge, (O) Other
LocalPort    RemoteDevice    RemotePortID    Capability    RemotePortDescr
-----------  --------------  --------------  ------------  -----------------
Ethernet24   sonic           fortyGigE52     BR            Ethernet60
------------------------------------------------
```

Step 2. check LLDP neighbor

```
admin@sonic:~$ show lldp neighbors Ethernet24
-------------------------------------------------------------------------------
LLDP neighbors:
-------------------------------------------------------------------------------
Interface:    Ethernet24, via: LLDP, RID: 5, Time: 0 day, 00:04:51
  Chassis:
    ChassisID:    mac a8:2b:b5:9d:c0:bb
    SysName:      sonic
    SysDescr:     Debian GNU/Linux 9 (stretch) Linux 4.9.0-9-2-amd64 #1 SMP Debian 4.9.168-1+deb9u3 (2015-12-
19) x86_64
    TTL:          120
    MgmtIP:       10.1.0.1
    MgmtIP:       fe80::aa2b:b5ff:fe9d:c0bb
    Capability:   Bridge, on
    Capability:   Router, on
    Capability:   Wlan, off
    Capability:   Station, off
  Port:
    PortID:       local fortyGigE52
    PortDescr:    Ethernet60
-------------------------------------------------------------------------------
```

MAC address table and aging time

Example:

Restriction:

1. Aging time is not configurable

Reference model:

- Switch model name: AS7726-32X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200103_052306_ec201904_108

Default setting:

1. MAC aging time is 600s by default

Note:

https://github.com/Azure/sonic-utilities/blob/master/doc/Command-Reference.md#fdb

This command displays the MAC (FDB) entries either in full or partial as given below.

```
admin@sonic:~$ show mac
No. Vlan MacAddress Port Type
----- ------ ------------ ------ ------
Total number of entries 0
```

1. show mac - displays the full table
2. show mac -v - displays the MACs learnt on the particular VLAN ID.
3. show mac -p - displays the MACs learnt on the particular port.

Clear the MAC (FBD) table

```
admin@sonic:~$ sonic-clear fdb all
FDB entries are cleared.
```

Appendix:

Check MAC aging time by bcmshell

```
admin@sonic:~$ bcmcmd 'age'
age
Current age timer is 600.
```

NTP (Network Time Protocol)

Example: Sync the time with remote NTP server.

Reference model:

- Switch model name: AS5835-54X
- Edgecore SONiC version:

Restriction:

1. SONiC 201904 version does NOT support NTP config command https://github.com/Azure/sonic-utilities/blob/master/doc/Command-Reference.md#ntp-config-commands
2.
```
config ntp add <ip_address>
```

Procedure:

Step 1. Configure IP address to network communication (Refer to Management and front port IPv4/IPv6 Address)

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{
    ...omitted
    "MGMT_INTERFACE": {
        "eth0|192.168.10.200/24": {
            "gwaddr": "192.168.10.254"
        }
    },
    "MGMT_PORT": {
        "eth0": {
            "admin_status": "up",
            "alias": "eth0"
        }
    },
    ...omitted
}
```

Step 2. Add the NTP server.

```
admin@sonic:/etc/sonic$ sudo vi config_db.json

{
    ...omitted
    "NTP_SERVER": {
        "1.debian.pool.ntp.org": {},
        "time.google.com": {},
        "118.163.170.6": {}
    },
    ...omitted
}
```

Note: We could only configure the NTP servers via config_db.json on SONiC version 2019040 (Refer to Restriction #1)

Step 3. Apply the config_db.json.

```
admin@sonic:~$ sudo config reload -y
```

Step 4. Add the DNS server. (Refer to this article)

```
admin@sonic:~$ sudo vi /etc/resolv.conf
nameserver 8.8.8.8
```

Step 5-1. List all the timezone information.

```
admin@sonic:~$ sudo timedatectl list-timezones
```

Note: Use arrow key to roll up and down, press "Space" for next page, "q" for quit.

Step 5-2. Modify to your timezone.

```
admin@sonic:~$ sudo timedatectl set-timezone Asia/Taipei
```

Step 6. Check the system time by linux command

```
admin@sonic:~$ date
Tue Apr 21 16:15:05 CST 2020
```

Step 7. Check NTP status

Result: Success

```
admin@sonic:~$ date
Tue Apr 21 16:18:09 CST 2020
admin@sonic:~$ show ntp
   remote           refid      st t when poll reach   delay   offset  jitter
==============================================================================
*118.163.170.6   218.73.139.35   2 u  11  64     1    2.571  -4.708   3.085
+103.226.213.30  .PPS.           1 u  11  64     1    3.210  -4.401   3.040
+216.239.35.8    .GOOG.          1 u  10  64     1    5.475  -5.020   3.215
```

Result: Failed (SONiC version 201904)

```
admin@sonic:~$ show ntp
ntpq: read: Connection refused
admin@sonic:~$
```

Known issue: SONiC NTP does not work well on version 201904. The following steps are workaround solution.

Analysis:

Compare current time and system time. If the difference time is larger than 1600 seconds, thus NTP is malfunction.

e.g

Current date & time: 2020-04-28 13:14:30

Check hardware clock (date and time is from BIOS).

```
admin@sonic:~$ sudo hwclock
2017-04-28 05:12:06.952322+0000
```

The difference of above time is around 1 year. therefore NTP gets failed.

Workaround solution: on SONiC version 20191904:

If the time doesn't sync to NTP server, please try the following commands.

Step 1. Stop NTP service

```
admin@sonic:~$ sudo systemctl stop ntp
```

Step 2. Sync date & time from NTP server manually.

```
admin@sonic:~$ sudo ntpd -qg
```

Note: This command takes few seconds, if it stuck, please use "ctrl + c" to quit then ping to NTP server and make sure it could work.

Step 3. Write system time to hardware clock (BIOS)

```
admin@sonic:~$ sudo hwclock -w
```

Step 4. Start NTP service

```
admin@sonic:~$ sudo systemctl start ntp
```

SNMP (Simple Network Management Protocol)

Example:

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Restriction:

1. There's no SONiC command to modify SNMP settings. Edit the file /etc/sonic/snmp.yml instead.
2. Does not support SNMP trap

Default setting:

1. SNMP is enabled
2. Support "Read" only for SNMPv2
3. SNMP community default is public.
4. Support list:

RFC 1213	MIB-II
RFC 2737	Physical Table MIB
RFC 2863	Interfaces MIB
RFC 3433	Sensor Table MIB
RFC 4363	dot1qTpFdbPort in Q-BRIDGE-MIB
IEEE 802.1 AB	LLDP-MIB

Procedure :
Steps 1. Edit the file /etc/sonic/snmp.yml that's SNMP config file.

```
admin@sonic:~$ vi /etc/sonic/snmp.yml
snmp_rocommunity: public
snmp_location: public
```

Steps 2. config reload or power cycle the switch

```
admin@sonic:~$ sudo config reload -y
```

TACACS+ for user authentication

Example:

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Default Setting:

```
admin@sonic:~$ show tacacs
TACPLUS global auth_type pap (default)
TACPLUS global timeout 5 (default)
TACPLUS global passkey <EMPTY_STRING> (default)
```

By default, AAA users local database for authentication.

```
admin@sonic:~$ show aaa
AAA authentication login local (default)
AAA authentication failthrough False (default)
```

Procedure :

Step 1. Set the management IP on the switch (Refer to this article)

Step 2. Add the TACACS Server host to the switch

```
admin@sonic:~$ sudo config tacacs add 192.168.1.10
```

Step 3. Set the TACACS authentication key (testing123 as example)

```
admin@sonic:~$ sudo config tacacs passkey testing123
```

Step 4. check the TACACS server settings

```
admin@sonic:~$ show tacacs
TACPLUS global auth_type pap (default)
TACPLUS global timeout 5 (default)
TACPLUS global passkey testing123

TACPLUS_SERVER address 192.168.1.10
 priority 1
 tcp_port 49
```

Step 5. Use tacacs+ database for user authentication

```
admin@sonic:~$ sudo config aaa authentication login tacacs+
```

Step 6. check aaa settings

```
admin@sonic:~$ show aaa
AAA authentication login tacacs+
AAA authentication failthrough False (default)
```

Step 7. Save the setting to config_db.json

```
admin@sonic:~$ sudo config save -y
```

Result: User named "test" login to the SONiC by TACACS authentication

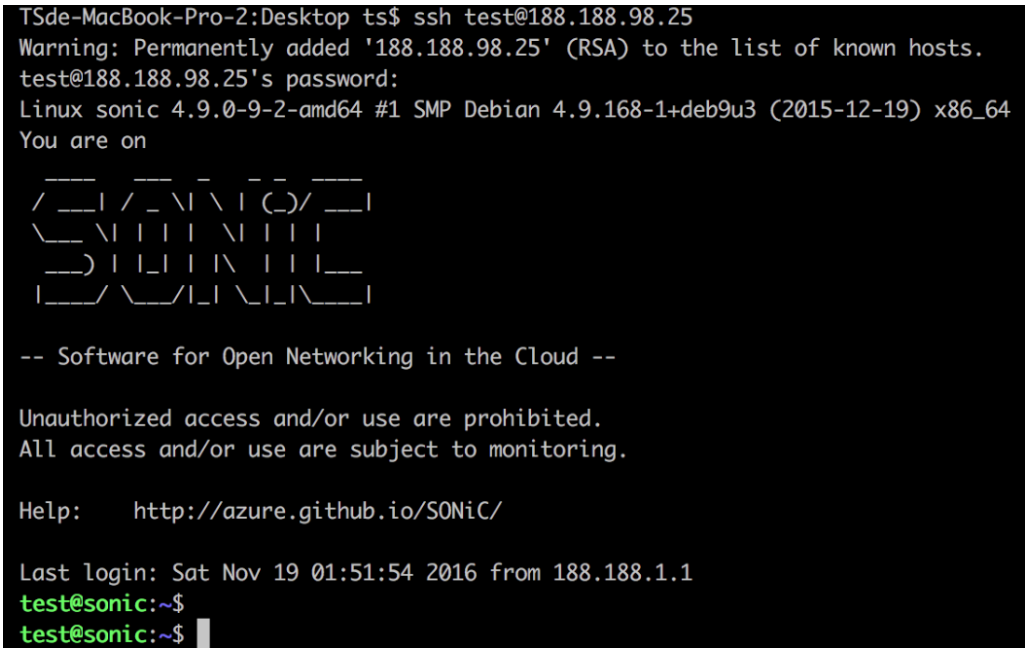

Appendix: TACACS+ server configuration

Step 1: Generate the login password by "tac_pwd"

```
[ Server Setting ] :
Steps 1. Generate the login password by "tac_pwd"
Server:~# tac_pwd
Password to be encrypted: test
BKe8b/ZgWAQ92
Server:~# tac_pwd
Password to be encrypted: admin
9HYczqUTI2Aoo
```

Step 2: TACACS Server configuration.

```
Server:~# vi /etc/tacacs+/tac_plus.conf
accounting file = /var/log/tac_plus.acct
key = testing123
user = DEFAULT {
login = PAM
service = ppp protocol = ip {}
}
group = network_admin {
    default service = permit
    service = exec {
        priv-lvl = 15
    }
    cmd = show {
        permit .*
    }
}
user = test {
    login = des BKe8b/ZgWAQ92
    pap = des BKe8b/ZgWAQ92
    member = network_admin
}
user = admin {
    login = des 9HYczqUTI2Aoo
    pap = des 9HYczqUTI2Aoo
    member = network_admin
}
```

VLAN & Inter-VLAN Routing

Example: Create VLAN and set untagged or tagged VLAN member on the interface.
Topology :

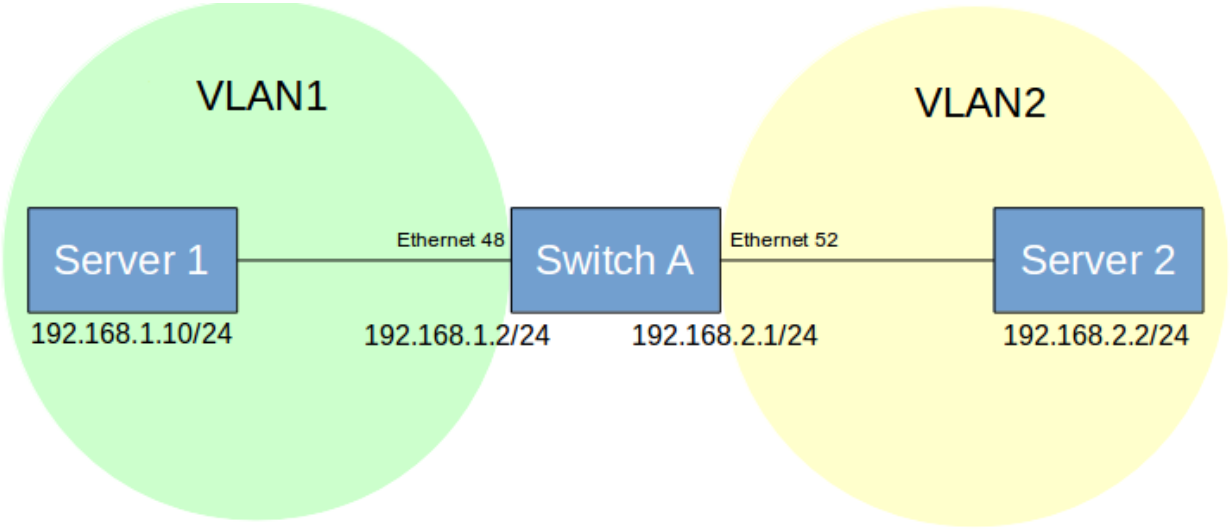

Procedure :

Steps 1. Check IP interface before adding IP address to VLAN interface.

Caution: IP address has to be removed from physical interface

```
admin@sonic:~$ show ip interfaces
Interface    IPv4 address/mask    Admin/Oper
-----------  -------------------  ------------
Ethernet0    192.168.1.1/24       up/up
docker0      240.127.1.1/24       up/down
eth0         188.188.98.25/16     up/up
lo           127.0.0.1/8          up/up
             10.1.0.1/32
admin@sonic:~$ sudo config interface ip remove Ethernet0 192.168.1.1/24
```

i.e Since there is IP address on physical interface "Ethernet0", we have to remove IP address from it.

Steps 2. Create VLAN1 and VLAN2 on the switch

```
admin@sonic:~$ sudo config vlan add 1
admin@sonic:~$ sudo config vlan add 2
```

Steps 3. Let the Ethernet48 as VLAN1 untag member.

```
admin@sonic:~$ sudo config vlan member add -u 1 Ethernet48
```

Steps 4. Let the Ethernet52 as VLAN2 tag member.

```
admin@sonic:~$ sudo config vlan member add 2 Ethernet52
```

Caution: please make sure the port name is correct. Otherwise it does not work.

Since there's no warning messages for wrong command.

i.e use "*show vlan brief*" to double check above settings.

```
admin@sonic:~$ show vlan brief
+-----------+----------------+------------+----------------+-----------------------+
|   VLAN ID | IP Address     | Ports      | Port Tagging   | DHCP Helper Address   |
+===========+================+============+================+=======================+
|         1 |                | Ethernet48 | untagged       |                       |
+-----------+----------------+------------+----------------+-----------------------+
|         2 |                | Ethernet52 | tagged         |                       |
+-----------+----------------+------------+----------------+-----------------------+
```

Example: Inter-VLAN routing

Steps 1. Configure IP addresses on VLAN1 and VLAN2

```
admin@sonic:~$ sudo config interface ip add Vlan1 192.168.1.2/24
admin@sonic:~$ sudo config interface ip add Vlan2 192.168.2.1/24
```

Caution: Vlan1 and Vlan2 are the Names of VLAN interface. So, "V" must be capital.

Steps 2. Check the VLAN setting and IP addresses

```
admin@sonic:~$ show vlan brief
+-----------+----------------+------------+----------------+-----------------------+
|   VLAN ID | IP Address     | Ports      | Port Tagging   | DHCP Helper Address   |
+===========+================+============+================+=======================+
|         1 | 192.168.1.2/24 | Ethernet48 | untagged       |                       |
+-----------+----------------+------------+----------------+-----------------------+
|         2 | 192.168.2.1/24 | Ethernet52 | tagged         |                       |
+-----------+----------------+------------+----------------+-----------------------+

admin@sonic:~$ show ip int
Interface    IPv4 address/mask    Admin/Oper
-----------  -------------------  ------------
Vlan1        192.168.1.2/24       up/up
Vlan2        192.168.2.1/24       up/up
docker0      240.127.1.1/24       up/down
eth0         188.188.98.30/16     up/up
lo           127.0.0.1/8          up/up
             10.1.0.1/32
```

Steps 3. Check the native VLAN (PVID)

```
admin@sonic:~$ sudo bridge vlan
port        vlan ids
docker0          1 PVID Egress Untagged

Bridge           1
         2

dummy          1 PVID Egress Untagged

Ethernet48           1 PVID Egress Untagged

Ethernet52           1 PVID Egress Untagged
         2
```

Steps 4. Save the setting to config_db.json

```
admin@sonic:~$ sudo config save -y
```

Steps 5. Check VLAN setting on config_db.json

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{

    ...
    "VLAN": {
        "Vlan1": {
            "members": [
                "Ethernet48",
            ],
            "vlanid": "1"
        },
        "Vlan2": {
            "members": [
                "Ethernet52",
            ],
            "vlanid": "2"
        },
    },
    "VLAN_MEMBER": {
        "Vlan1|Ethernet48": {
            "tagging_mode": "untagged"
        },
        "Vlan2|Ethernet52": {
            "tagging_mode": "tagged"
        }
    },
     "VLAN_INTERFACE": {
        "Vlan1|192.168.1.2/24": {},
        "Vlan2|192.168.2.1/24": {}
    },
    ...
}
```

Debug and Troubleshooting

Hardware, Software and System Information

Example:

- Hardware information
- Software information
- System information

Reference model:

- Switch model name: AS5835-54X,AS7326-56X,AS7726-32X,AS7816-64X
- Edgecore SONiC version: all versions

Example: Hardware information

show environment This command displays the platform environmentals, such as voltages, temperatures and fan speeds.

```
admin@sonic:~$ show environment
```

show platform psustatus This command displays PSU status information.

```
admin@sonic:~$ show platform psustatus
PSU    Status
-----  --------
PSU 1  OK
PSU 2  NOT OK
```

Note:

OK: input power present

NOT OK: no input power

show platform summary This command displays a summary of the device hardware platform

```
admin@sonic:~$ show platform summary
Platform: x86_64-accton_as7816_64x-r0
HwSKU: Accton-AS7816-64X
ASIC: broadcom
```

show platform syseeprom This command displays system EEPROM information such as MAC address, serial number, platform name

```
admin@sonic:~$ show platform syseeprom
TlvInfo Header:
   Id String:    TlvInfo
   Version:      1
   Total Length: 172
TLV Name             Code Len Value
-------------------- ---- --- -----
Manufacture Date     0x25  19 08/27/2018 11:19:49
Label Revision       0x27   4 R0AA
Manufacturer         0x2B   6 Accton
Manufacture Country  0x2C   2 TW
Vendor Name          0x2D   8 Edgecore
Product Name         0x21  18 7816-64X-O-AC-F-JN
Part Number          0x22  13 FP3AT7664000A
Serial Number        0x23  11 AAA1834AAAO
Base MAC Address     0x24   6 B8:6A:97:19:C2:46
MAC Addresses        0x2A   2 300
Platform Name        0x28  27 x86_64-accton_as7816_64x-r0
ONIE Version         0x29  13 2018.05.00.08
Diag Version         0x2E  11 01.03.00.08
CRC-32               0xFE   4 0x8EA1132B
(checksum valid)
```

Example: Software information

show version This command displays the currently SONiC image version as well as Docker image versions.

```
admin@sonic:~$ show version

SONiC Software Version: SONiC.Edgecore-SONiC_20200103_052306_ec201904_108
Distribution: Debian 9.11
Kernel: 4.9.0-9-2-amd64
Build commit: 9f46f888
Build date: Fri Jan 3 07:37:30 UTC 2020
Built by: ubuntu@ip-10-5-1-128

Platform: x86_64-accton_as7726_32x-r0
HwSKU: Accton-AS7726-32X
ASIC: broadcom
Serial Number: 772632X1846012
Uptime: 23:27:52 up 6:11, 1 user, load average: 2.18, 2.05, 1.94

Docker images:
REPOSITORY TAG IMAGE ID SIZE
docker-syncd-brcm Edgecore-SONiC_20200103_052306_ec201904_108 2d593cf6e660 429MB
docker-syncd-brcm latest 2d593cf6e660 429MB
docker-snmp-sv2 Edgecore-SONiC_20200103_052306_ec201904_108 4e64f714880b 312MB
docker-snmp-sv2 latest 4e64f714880b 312MB
docker-dhcp-relay Edgecore-SONiC_20200103_052306_ec201904_108 fb33515c109f 287MB
docker-dhcp-relay latest fb33515c109f 287MB
docker-database Edgecore-SONiC_20200103_052306_ec201904_108 b89aac84342c 279MB
docker-database latest b89aac84342c 279MB
docker-lldp-sv2 Edgecore-SONiC_20200103_052306_ec201904_108 cc2b21ff545d 298MB
docker-lldp-sv2 latest cc2b21ff545d 298MB
docker-orchagent Edgecore-SONiC_20200103_052306_ec201904_108 16c0ee198929 318MB
docker-orchagent latest 16c0ee198929 318MB
docker-teamd Edgecore-SONiC_20200103_052306_ec201904_108 5b118bf91623 300MB
docker-teamd latest 5b118bf91623 300MB
docker-sonic-telemetry Edgecore-SONiC_20200103_052306_ec201904_108 a0c682a1b520 301MB
docker-sonic-telemetry latest a0c682a1b520 301MB
docker-router-advertiser Edgecore-SONiC_20200103_052306_ec201904_108 fc586a727f57 278MB
docker-router-advertiser latest fc586a727f57 278MB
docker-platform-monitor Edgecore-SONiC_20200103_052306_ec201904_108 48deb7abd8e4 325MB
docker-platform-monitor latest 48deb7abd8e4 325MB
docker-fpm-frr Edgecore-SONiC_20200103_052306_ec201904_108 9ae92394036f 318MB
docker-fpm-frr latest 9ae92394036f 318MB
```

show processes cpu

show processes memory

show processes summary

These commands display the current CPU & Memory usage by process and summary information about all process.

```
admin@sonic:~$ show processes cpu
admin@sonic:~$ show processes memory
admin@sonic:~$ show processes summary
```

show service This command displays the state of all the SONiC processes running inside a docker container. This helps to identify the status of SONiC's critical processes.

```
admin@sonic:~$ show services
```

show system-memory This command displays the system-wide memory utilization information

```
admin@sonic:~$ show system-memory
 total used free shared buff/cache available
Mem: 15951 1255 13382 19 1312 14373
Swap: 0 0 0
```

Example: System information

show uptime This command displays the current system uptime

```
admin@sonic:~$ show uptime
shup 6 hours, 30 minutes
```

show clock This command displays the current date and time configured on the system

```
admin@sonic:~$ show clock
Thu Nov 3 23:46:55 UTC 2016
```

show user This command displays a list of users currently logged in to the device

```
admin@sonic:~$ show user
admin ttyS0 2016-11-03 23:46
```

show logging This command displays all the currently stored log messages.

```
admin@sonic:~$ show logging
```

Note: Refer to this article for detail

show reboot-cause This command displays the cause of the previous reboot

```
admin@sonic:~$ show reboot-cause
User issued 'reboot' command [User: admin, Time: Wed Nov  9 10:14:59 UTC 2016]
```

```
admin@sonic:~$ show reboot-cause
Unexpected reboot
```

Note: Unexpected reboot: power outage

Note: log is saved on /host/reboot-cause/.

```
admin@sonic:~$ ls -l /host/reboot-cause/
total 8
-rw-r--r-- 1 root root 79 Aug 6 16:22 previous-reboot-cause.txt
-rw-r--r-- 1 root root 18 Aug 6 16:23 reboot-cause.txt
```

Syslog

Example: system log

Restriction:

1. There is no SONiC command for system log

Reference model:

- Switch model name: AS7726-32X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200103_052306_ec201904_108

Note:

This command displays all the currently stored log messages. All the latest processes and corresponding transactions are stored in the "syslog" file. This file is saved in the path /var/log and can be viewed by giving the command sudo cat syslog as this requires root login.

```
admin@sonic:~$ show logging
```

Optionally, you can follow the log live as entries are written to it by specifying the -f or --follow flag

```
admin@sonic:~$ show logging --follow
```

Example: remote syslog server

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Restriction:

1. There's no SONiC command for remote syslog server.

Procedure :

Step 1. Edit /etc/sonic/config_db.json instead of SONiC command.

Add syslog server and its IP address.

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{

    ...
    "SYSLOG_SERVER": {
        "192.168.1.3": {}
    },
    ...
}
```

Step 2. config reload or power cycle the switch

```
admin@sonic:~$ sudo config reload -y
```

Caution:

- Make sure switch IP address and connectivity works well. (refer to Management IP address)

Result: Switch sends syslog traps (all severity levels) to syslog server.

No.	Time	Source	Destination	Protocol	Length	Info
115	1.857263	188.188.98.2	188.188.36.8	Syslog	148	Nov 10 11:29:58.435508 sonic ERR sshd[12235]: e
116	1.858543	188.188.98.2	188.188.36.8	Syslog	150	Nov 10 11:29:58.436823 sonic ERR sshd[12235]: e
219	4.138847	188.188.98.2	188.188.36.8	Syslog	150	Nov 10 11:30:00.717079 sonic INFO sshd[12235]:
220	4.140552	188.188.98.2	188.188.36.8	Syslog	154	Nov 10 11:30:00.718690 sonic INFO sshd[12235]:
287	5.168307	188.188.98.2	188.188.36.8	Syslog	153	Nov 10 11:30:01.746408 sonic INFO CRON[12253]:
288	5.169507	188.188.98.2	188.188.36.8	Syslog	205	Nov 10 11:30:01.747574 sonic INFO CRON[12254]:
289	5.199678	188.188.98.2	188.188.36.8	Syslog	142	Nov 10 11:30:01.777863 sonic INFO CRON[12253]:

Time	IP Address	Msg T...	Message
Aug 12 10:46:09	188.188.98.2	user.info	Nov 10 11:30:01.777863 sonic INFO CRON[12253]: pam_unix(cron:session): session closed for user root
Aug 12 10:46:09	188.188.98.2	user.info	/dev/null 2>&1; /usr/sbin/logrotate /etc/logrotate.conf > /dev/null 2>&1)
Aug 12 10:46:09	188.188.98.2	user.info	Nov 10 11:30:01.746408 sonic INFO CRON[12253]: pam_unix(cron:session): session opened for user root by (uid=0)
Aug 12 10:46:08	188.188.98.2	user.info	Nov 10 11:30:00.718690 sonic INFO sshd[12235]: pam_unix(sshd:session): session opened for user admin by (uid=0)
Aug 12 10:46:08	188.188.98.2	user.info	Nov 10 11:30:00.717079 sonic INFO sshd[12235]: Accepted password for admin from 188.188.1.1 port 49528 ssh2
Aug 12 10:46:06	188.188.98.2	user.info	Nov 10 11:29:58.436823 sonic ERR sshd[12235]: error: Could not load host key: /etc/ssh/ssh_host_ed25519_key
Aug 12 10:46:06	188.188.98.2	user.info	Nov 10 11:29:58.435508 sonic ERR sshd[12235]: error: Could not load host key: /etc/ssh/ssh_host_ecdsa_key

Example: Adjust severity level for remote syslog server

Restriction:

1. There's no SONiC command

Default setting: all severity levels

```
admin@sonic:~$ sudo vi /usr/share/sonic/templates/rsyslog.conf.j2
omitted...

#Set remote syslog server
{% for server in SYSLOG_SERVER %}
*.* @{{ server }}:514;SONiCFileFormat
{% endfor %}

omitted...
```

Step 1. Edit the file /usr/share/sonic/templates/rsyslog.conf.j2

Warning and below. (Level 0 ~ 4)

```
admin@sonic:~$ sudo vi /usr/share/sonic/templates/rsyslog.conf.j2
{ % for server in SYSLOG_SERVER % }
*.warning @{{ server }}:514;SONiCFileFormat
{ % endfor % }
```

Only warning level. (Only Level 4)

```
admin@sonic:~$ sudo vi /usr/share/sonic/templates/rsyslog.conf.j2
{% for server in SYSLOG_SERVER %}
*.=warning @{{ server }}:514;SONiCFileFormat
{% endfor %}
```

All severity levels except warning. (Level 0 ~ 7 except Level 4)

```
admin@sonic:~$ sudo vi /usr/share/sonic/templates/rsyslog.conf.j2
{% for server in SYSLOG_SERVER %}
*.debug;*.!=warning @{{ server }}:514;SONiCFileFormat
{% endfor %}
```

Note:

Here's the Severity level by the standard.

Value	Severity	Keyword
0	Emergency	emerg
1	Alert	alert

2	Critical	crit
3	Error	err
4	Warning	warning
5	Notice	notice
6	Informational	info
7	Debug	debug

Step 2: Restart syslog service

```
admin@sonic:~$ sudo systemctl restart rsyslog-config
```

Tech-Support Dump

Example:

Reference model:

- Switch model name: AS5835-54X,AS7326-56X,AS7726-32X,AS7816-64X
- Edgecore SONiC version: All version

Procedure :

Generate a dump of troubleshooting data. An archive file containing the dump will be saved to the device /var/dump which can be sent to the Edgecore team for troubleshooting.

```
admin@sonic:~$ show techsupport
omitted...
removed '/var/dump/sonic_dump_sonic_20161103_182033/log/telemetry.log.20.gz'
mkdir: created directory '/var/dump/sonic_dump_sonic_20161103_182033/core'
sonic_dump_sonic_20161103_182033/core/syncd.1478686198.28.core.gz
removed '/var/dump/sonic_dump_sonic_20161103_182033/core/syncd.1478686198.28.core.gz'
removed directory '/var/dump/sonic_dump_sonic_20161103_182033/log'
removed directory '/var/dump/sonic_dump_sonic_20161103_182033/dump'
removed directory '/var/dump/sonic_dump_sonic_20161103_182033/core'
removed '/var/dump/sonic_dump_sonic_20161103_182033/etc'
removed directory '/var/dump/sonic_dump_sonic_20161103_182033'
/var/dump/sonic_dump_sonic_20161103_182033.tar.gz
admin@sonic:~$
```

The dump file is saved to /var/dump

```
admin@sonic:~$ ls -lh /var/dump/
total 28M
-rw-r--r-- 1 root root 14M Nov 3 18:21 sonic_dump_sonic_20161103_182033.tar.gz
-rw-r--r-- 1 root root 14M Nov 3 18:22 sonic_dump_sonic_20161103_182225.tar.gz
```

Optionally, you can generate the troubleshooting data since given date by specifying --since flag

For example:

```
admin@sonic:~$ sudo show techsupport --since="yesterday"
```

Interface, Cables, Optics and Transceivers

AS5835-54X 100G Port Setting

Due to AS5835-54X hardware design, it requires additional settings for 100G ports.

Reference: Basic Port Speed Configuration

Example: Let 100G port link up on port interface 49 and 50

Caution:

1. It's necessarily to re-configure it again if you reboot the switch.
2. This issue has been **enhanced** on SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178 version so you **DON'T HAVE** to configure 100G port setting on SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178 version and later .

Reference model:

- Switch model name: AS5835-54X
- SONiC Software Version: SONiC.Edgecore-SONiC_20200722_070543_ec201911_141

Procedures:

Step 1. Make sure license is loaded properly.

```
admin@sonic:~$ cat /host/license/license.lic
aos-license/1.0
Name: Edgecore-TS
CPU-MAC-Address: 04-F8-F8-8D-81-E1
Project-Number: 5835-54X-O
Accept-Mode: *
License-Number: 1280b0a4-da90-4742-b112-f53ba344c99f
License-Issue-Date: Mon Aug 3 10:45:27 2020
License-Valid-Start-Date: Mon Aug 3 00:00:00 2020
License-Valid-End-Date: Fri Aug 2 00:00:00 2030
```

Note : License is not required on SONiC 201911.2 2020 Aug version and later.

Step 2. Check the status of "QSFP+_MOD_RST" by getting under */sys/bus/i2c/devices/3-0062* file. The default value should be **0x00** which means **Reset QSFP+ module by SW**.

```
root@sonic:/sys/bus/i2c/devices/3-0062# i2cget -f -y 3 0x62 0x15
0x00
```

Step 3. Check the link status of the ports which are connected. Below example takes loop connection between port 49 and port 50. Link status are down due to Reset QSFP+ mode by default.

```
admin@sonic:~$ show int status
  Interface        Lanes    Speed    MTU          Alias    Vlan    Oper    Admin    Type    Asym PFC
-----------  -----------  -------  -----  -------------  ------  ------  -------  ------  ----------
  Ethernet0            1      10G   9100       tenGigE1  routed    down       up     N/A         N/A
  Ethernet1            2      10G   9100       tenGigE2  routed    down       up     N/A         N/A
...
 Ethernet47           76      10G   9100      tenGigE48  routed    down       up     N/A         N/A
 Ethernet48  37,38,39,40     100G   9100  hundredGigE49  routed    down       up     N/A         N/A

 Ethernet52  29,30,31,32     100G   9100  hundredGigE50  routed    down       up     N/A         N/A
```

Step 4. Configure "QSFP+_MOD_RST" of port 49 and port 50 to normal mode.

```
root@sonic:/sys/bus/i2c/devices/3-0062# echo "1" > module_reset_49
root@sonic:/sys/bus/i2c/devices/3-0062# echo "1" > module_reset_50
```

Table below is QSFP+_MOD_RST bit configuration

Bit7	Bit6	Bit5	Bit4	Bit3	Bit2	Bit1	Bit0
Reserved	Reserved	QSFP+ Port54	QSFP+ Port53	QSFP+ Port52	QSFP+ Port51	QSFP+ Port50	QSFP+ Port49

QSFP+_MOD_RST status for QSFP+ Port 49 to Port 54
0: Reset QSFP+ module by SW.
1: Normal status
Note : Depend on transceiver supplier support this function or not.

Step 5. Check the status of QSFP+_MOD_RST and link status.

```
root@sonic:/sys/bus/i2c/devices/3-0062# i2cget -f -y 3 0x62 0x15
0x03
```

0x03 means bit0 (Port49) and bit1 (Port50) has been set to 1 (normal status).

```
root@sonic:/sys/bus/i2c/devices/3-0062# show int status
  Interface       Lanes    Speed    MTU          Alias    Vlan    Oper    Admin            Type    Asym PFC
-----------  ----------  -------  -----  -------------  ------  ------  -------  --------------  ----------
  Ethernet0           1      10G   9100      tenGigE1  routed    down      up             N/A         N/A
  Ethernet1           2      10G   9100      tenGigE2  routed    down      up             N/A         N/A
...
 Ethernet47          76      10G   9100     tenGigE48  routed    down      up             N/A         N/A
 Ethernet48  37,38,39,40     100G   9100  hundredGigE49  routed      up      up             N/A         N/A

 Ethernet52  29,30,31,32     100G   9100  hundredGigE50  routed      up      up             N/A         N/A
 Ethernet56  33,34,35,36     100G   9100  hundredGigE51  routed    down      up             N/A         N/A
```

AS7326-56X change port speed 25G to 10G

Reference: Basic Port Speed Configuration

AS7326-56X Hardware Restriction:

- Four 25G interfaces are recognized as a group, and all interface members in the group SHALL operate at the same speed. In other words, when user configures only one interface, three other interfaces in the same group will be set to the same interface speed. For interface group list, please refer to **Interface Group Table** below.
- A group consists of 4 x 25G interfaces are separated from 100G (i.e 4 lanes).

Example: change port speed on interface 1, 2, 3, 6

Caution:

1. Interface 1, 2, 3, 6 are a group on AS7326-56X due to hardware restriction.
2. Interface numbering starts from 1 on front panel. However, it starts from 0 on SONiC. e.g Interface 1 is Ethernet0 on SONiC.

Reference model:

- Switch model name: AS7326-56X
- SONiC Software Version: SONiC.Edgecore-SONiC_20200507_052107_ec201911_74

Procedures:

Step 1. Configure 25G interface (SFP) to 10G interface. Let's take Ethernet0 as an example here.

```
admin@sonic:~$ sudo config interface speed Ethernet0 10000
```

Caution:

1. Please make sure **port name** (i.e Ethernet0) is correct. Otherwise it does not work.
2. Due to hardware restriction which is mentioned above, configuring Ethernet0 to 10G will automatically configure Ethernet1, Ethernet2, Ethernet5 which are in the same group as well.

Step 2. Save the configuration to config_db.json

```
admin@sonic:~$ sudo config save -y
```

Step 3. Check interface status from bcmshell.

```
admin@sonic:~$ bcmcmd 'ps'
ps
ena/ speed/ link auto STP lrn inter max cut loop
port link Lns duplex scan neg? state pause discrd ops face frame thru? back
xe0( 3) up 1 10G FD SW No Forward None F XFI 9122 No
xe1( 2) up 1 10G FD SW No Forward None F XFI 9122 No
xe2( 4) up 1 10G FD SW No Forward None F XFI 9122 No
xe3( 8) down 1 25G FD SW No Forward None F CR 9122 No
xe4( 7) down 1 25G FD SW No Forward None F CR 9122 No
xe5( 1) up 1 10G FD SW No Forward None F XFI 9122 No
xe6( 5) down 1 25G FD SW No Forward None F CR 9122 No
xe7( 16) down 1 25G FD SW No Forward None F CR 9122 No
xe8( 6) down 1 25G FD SW No Forward None F CR 9122 No
xe9( 14) down 1 25G FD SW No Forward None F CR 9122 No
xe10( 13) down 1 25G FD SW No Forward None F CR 9122 No
xe11( 15) down 1 25G FD SW No Forward None F CR 9122 No
xe12( 23) down 1 25G FD SW No Forward None F CR 9122 No
xe13( 22) down 1 25G FD SW No Forward None F CR 9122 No
xe14( 24) down 1 25G FD SW No Forward None F CR 9122 No
xe15( 32) down 1 25G FD SW No Forward None F CR 9122 No
xe16( 31) down 1 25G FD SW No Forward None F CR 9122 No
xe17( 21) down 1 25G FD SW No Forward None F CR 9122 No
xe18( 29) down 1 25G FD SW No Forward None F CR 9122 No
xe19( 36) down 1 25G FD SW No Forward None F CR 9122 No
xe20( 30) down 1 25G FD SW No Forward None F CR 9122 No
xe21( 34) down 1 25G FD SW No Forward None F CR 9122 No
xe22( 33) down 1 25G FD SW No Forward None F CR 9122 No
xe23( 35) down 1 25G FD SW No Forward None F CR 9122 No
xe24( 43) down 1 25G FD SW No Forward None F CR 9122 No
xe25( 42) down 1 25G FD SW No Forward None F CR 9122 No
xe26( 44) down 1 25G FD SW No Forward None F CR 9122 No
xe27( 52) down 1 25G FD SW No Forward None F CR 9122 No
xe28( 51) down 1 25G FD SW No Forward None F CR 9122 No
xe29( 41) down 1 25G FD SW No Forward None F CR 9122 No
xe30( 49) down 1 25G FD SW No Forward None F CR 9122 No
xe31( 60) down 1 25G FD SW No Forward None F CR 9122 No
xe32( 50) down 1 25G FD SW No Forward None F CR 9122 No
xe33( 58) down 1 25G FD SW No Forward None F CR 9122 No
xe34( 57) down 1 25G FD SW No Forward None F CR 9122 No
xe35( 59) down 1 25G FD SW No Forward None F CR 9122 No
xe36( 62) down 1 25G FD SW No Forward None F CR 9122 No
xe37( 63) down 1 25G FD SW No Forward None F CR 9122 No
xe38( 64) down 1 25G FD SW No Forward None F CR 9122 No
xe39( 67) down 1 25G FD SW No Forward None F CR 9122 No
xe40( 68) down 1 25G FD SW No Forward None F CR 9122 No
xe41( 61) down 1 25G FD SW No Forward None F CR 9122 No
xe42( 70) down 1 25G FD SW No Forward None F CR 9122 No
xe43( 71) down 1 25G FD SW No Forward None F CR 9122 No
xe44( 69) down 1 25G FD SW No Forward None F CR 9122 No
xe45( 73) down 1 25G FD SW No Forward None F CR 9122 No
xe46( 74) down 1 25G FD SW No Forward None F CR 9122 No
xe47( 72) down 1 25G FD SW No Forward None F CR 9122 No
ce0( 79) down 4 100G FD SW No Forward None F CAUI4 9122 No
ce1( 87) down 4 100G FD SW No Forward None F CAUI4 9122 No
ce2( 95) down 4 100G FD SW No Forward None F CAUI4 9122 No
ce3( 99) down 4 100G FD SW No Forward None F CAUI4 9122 No
ce4(107) down 4 100G FD SW No Forward None F CAUI4 9122 No
ce5(115) down 4 100G FD SW No Forward None F CAUI4 9122 No
ce6(123) down 4 100G FD SW No Forward None F CAUI4 9122 No
ce7(127) down 4 100G FD SW No Forward None F CAUI4 9122 No
drivshell>
```

xe0 = Port 1 = Ethernet0
xe1 = Port 2 = Ethernet1
xe2 = Port 3 = Ethernet2
xe3 = Port 4 = Ethernet3
xe4 = Port 5 = Ethernet4
xe5 = Port 6 = Ethernet5

.... so on

Caution (known issue) :

As mentioned above, configuring one port (take an example here when we only configure one port Ethernet0 to 10G) will automatically configure three other ports in the same interface group. But, 'show interface status' command shows 10G port speed only on Ethernet0. This is known display issue and we will fix this issue in the next release. So, please check interface status from bcmshell instead which mentioned on step 3.

```
admin@sonic:~$ show interface status
Interface Lanes Speed MTU Alias Vlan Oper Admin Type Asym PFC
----------- --------------- ------- ----- ---------------- ------ ------ ------- -------------- ----------
Ethernet0 3 10G 9100 twentyfiveGigE1 routed up up SFP/SFP+/SFP28 N/A
Ethernet1 2 25G 9100 twentyfiveGigE2 routed up up SFP/SFP+/SFP28 N/A
Ethernet2 4 25G 9100 twentyfiveGigE3 routed up up SFP/SFP+/SFP28 N/A
Ethernet3 8 25G 9100 twentyfiveGigE4 routed down up N/A N/A
Ethernet4 7 25G 9100 twentyfiveGigE5 routed down up N/A N/A
Ethernet5 1 25G 9100 twentyfiveGigE6 routed up up SFP/SFP+/SFP28 N/A
Ethernet6 5 25G 9100 twentyfiveGigE7 routed down up N/A N/A
Ethernet7 16 25G 9100 twentyfiveGigE8 routed down up N/A N/A
Ethernet8 6 25G 9100 twentyfiveGigE9 routed down up N/A N/A
Ethernet9 14 25G 9100 twentyfiveGigE10 routed down up N/A N/A
Ethernet10 13 25G 9100 twentyfiveGigE11 routed down up N/A N/A
```

Appendix:

Interface Group Table

Group	Front Panel Interface	SONiC Ethernet Interface
1	1	Ethernet 0
	2	Ethernet 1
	3	Ethernet 2
	6	Ethernet 5
2	4	Ethernet 3
	5	Ethernet 4
	7	Ethernet 6
	9	Ethernet 8
3	8	Ethernet 7
	10	Ethernet 9
	11	Ethernet 10
	12	Ethernet 11
4	13	Ethernet 12
	14	Ethernet 13
	15	Ethernet 14
	18	Ethernet 17
5	16	Ethernet 15
	17	Ethernet 16
	19	Ethernet 18
	21	Ethernet 20
6	20	Ethernet 19
	22	Ethernet 21
	23	Ethernet 22
	24	Ethernet 23
7	25	Ethernet 24
	26	Ethernet 25
	27	Ethernet 26
	30	Ethernet 29

8	28	Ethernet 27
	29	Ethernet 28
	31	Ethernet 30
	33	Ethernet 32
9	32	Ethernet 31
	34	Ethernet 33
	35	Ethernet 34
	36	Ethernet 35
10	37	Ethernet 36
	38	Ethernet 37
	39	Ethernet 38
	42	Ethernet 41
11	40	Ethernet 39
	41	Ethernet 40
	43	Ethernet 42
	45	Ethernet 44
12	44	Ethernet 43
	46	Ethernet 45
	47	Ethernet 46
	48	Ethernet 47

AS7816-64X breakout cable 100G to 10G/25G

Example: Breakout cable 100G to 10G on ROW1

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200507_052107_ec201911_74

Default setting:

1. 64x100G interfaces

Restriction:

1. 4x25G or 4x10G are in unit of row, as shown in the following figure. Due to chip restriction, it is NOT allowed to configure some ports thus you SHALL to configure the whole row (16 ports per row) which contains the ports you desire to modify.
2. Hardware restriction: 64x100G or 128x25G or 128x10G

Note:

1. The purpose of procedure below is to guide you to configure breakout on row 1. The same method is applied to configure breakout cable on row 2, 3 and 4.
2. This breakout method only works on ecsonic version of 201911 (including) later. For older version, please refer to this tutorial.
3. Basically, you can modify config_db.json file by yourself to configure breakout but in order to make it easier in configuring, we provide four port_config.ini files for each row to obtain a new configured config_db.json. Please follow the procedure below

Procedure:

Step 1: Download port_config.ini file based on the row you want to breakout.

- port_config.ini download list
    1. port_config_row1.ini
    2. port_config_row2.ini
    3. port_config_row3.ini
    4. port_config_row4.ini

Step 2: Upload port_config.ini to SONiC

```
$sudo scp root@188.188.99.1:/var/lib/tftpboot/test/port_config_row1.ini ~/
```

Note: Copy the file "port_config_row1.ini" from a remote host (188.188.99.1) to the local host.

Step 3: Re-generate config_db.json by sonic command (Refer to this article) and reload.

```
admin@sonic:~$ sudo sonic-cfggen -H -p ~/port_config_row1.ini --preset t1 -k Accton-AS7816-64X > ~
/port_config_row1.log
admin@sonic:~$ sudo mv port_config_row1.log /etc/sonic/config_db.json
admin@sonic:~$ sudo config reload -y
```

Step 4: Checking the port status.

```
  admin@sonic:~$ show interfaces status
  Interface            Lanes    Speed    MTU          Alias    Vlan    Oper    Admin             Type    Asym
PFC
-----------  ---------------  -------  -----  -------------  ------  ------  -------  ---------------
----------
  Ethernet0               73      10G   9100         swp001  routed    down       up              N/A         N
/A
  Ethernet1               74      10G   9100         swp002  routed    down       up              N/A         N
/A
  Ethernet2               75      10G   9100         swp003  routed    down       up              N/A         N
/A
  Ethernet3               76      10G   9100         swp004  routed    down       up              N/A         N
/A
  Ethernet4      65,66,67,68     100G   9100   hundredGigE2  routed    down       up              N/A         N
/A
  Ethernet8               81      10G   9100         swp801  routed    down       up              N/A         N
/A
  Ethernet9               82      10G   9100         swp802  routed    down       up  QSFP28 or later         N
/A
 Ethernet10               83      10G   9100         swp803  routed    down       up              N/A         N
/A
 Ethernet11               84      10G   9100         swp804  routed    down       up   QSFP+ or later         N
/A
 Ethernet12      89,90,91,92     100G   9100   hundredGigE4  routed    down       up              N/A         N
/A
 Ethernet16              105      10G   9100        swp1601  routed    down       up              N/A         N
/A
 Ethernet17              106      10G   9100        swp1602  routed    down       up              N/A         N
/A
 Ethernet18              107      10G   9100        swp1603  routed    down       up              N/A         N
/A
 Ethernet19              108      10G   9100        swp1604  routed    down       up              N/A         N
/A
 Ethernet20     97,98,99,100     100G   9100   hundredGigE6  routed    down       up              N/A         N
/A
 ..
 omitted
```

Counter(Port counter and L3 RIF counter)

Examples :

- Port counter
- L3 RIF counter

Example: Port counter

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Procedure:

```
admin@sonic:~$ show interface counter

      IFACE    STATE    RX_OK    RX_BPS    RX_UTIL    RX_ERR    RX_DRP    RX_OVR    TX_OK     TX_BPS
TX_UTIL    TX_ERR    TX_DRP    TX_OVR
-----------  -------  -------  --------  ---------  --------  --------  --------  -------  ---------
---------  --------  --------  --------
  Ethernet0        U        3  1.04 B/s      0.00%         0         0         0       22  12.02 B/s
0.00%         0         0         0
  Ethernet4        D        0  0.00 B/s      0.00%         0         0         0        0   0.00 B/s
0.00%         0         0         0
  Ethernet8        U        3  1.04 B/s      0.00%         0         0         0       22  12.02 B/s
0.00%         0         0         0
 Ethernet12        D        0  0.00 B/s      0.00%         0         0         0        0   0.00 B/s
0.00%         0         0         0
 Ethernet16        U       12  8.88 B/s      0.00%         0         3         0       22  12.05 B/s
0.00%         0         0         0
 Ethernet20        D        0  0.00 B/s      0.00%         0         0         0        0   0.00 B/s
0.00%         0         0         0
 Ethernet24        D        0  0.00 B/s      0.00%         0         0         0        0   0.00 B/s
0.00%         0         0         0
 Ethernet28        D        0  0.00 B/s      0.00%         0         0         0        0   0.00 B/s
0.00%         0         0         0
 .
 .
 .
 omitted
```

Example: L3 RIF counter (NS, not support)

Restriction: Due to chipset restriction. Edgecore SONiC does NOT support L3 RIF counter.

Reference model:

- Switch model name:
  - AS5812-54X
  - AS5835-56X
  - AS7326-56X
  - AS7726-32X
  - AS7816-64X
- Edgecore SONiC version: ALL

Procedure:

```
admin@sonic:~$ show interface counter rif
      IFACE    RX_OK    RX_BPS    RX_PPS    RX_ERR    TX_OK    TX_BPS    TX_PPS    TX_ERR
-----------  -------  --------  --------  --------  -------  --------  --------  --------
  Ethernet0      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
  Ethernet1      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
  Ethernet2      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
  Ethernet3      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
  Ethernet4      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
  Ethernet8      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
  Ethernet9      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
 Ethernet10      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
 Ethernet11      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
 Ethernet12      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
 Ethernet16      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
 Ethernet17      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
 Ethernet18      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
 Ethernet19      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
 Ethernet20      N/A       N/A       N/A       N/A      N/A       N/A       N/A       N/A
```

```
  Ethernet24      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet25      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet26      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet27      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet28      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet32      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet33      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet34      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet35      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet36      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet40      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet41      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet42      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet43      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet44      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet48      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet49      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet50      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet51      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet52      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet56      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet57      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet58      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet59      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet60      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet64      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet65      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet66      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet67      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet68      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet72      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet73      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet74      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet75      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet76      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet80      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet81      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet82      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet83      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet84      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet88      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet89      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet90      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet91      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet92      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet96      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet97      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet98      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
  Ethernet99      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet100      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet104      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet105      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet106      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet107      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet108      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet112      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet113      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet114      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet115      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet116      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet120      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet121      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet122      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet123      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet124      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet128      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet132      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet136      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet140      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet144      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet148      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet152      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet156      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet160      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet164      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet168      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet172      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet176      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet180      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet184      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet188      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet192      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet196      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
 Ethernet200      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
```

```
Ethernet204      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet208      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet212      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet216      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet220      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet224      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet228      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet232      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet236      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet240      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet244      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet248      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
Ethernet252      N/A       N/A       N/A       N/A       N/A       N/A       N/A       N/A
```

Dynamic Port Breakout

Reference model:

- Switch model name: AS7326-56X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178

Known issue :

1. Port breakout fails if portchannel exists. Please remove portchannel before doing breakout.
2. After doing breakout, MTU size of the first interface will be disappeared. Doing config reload after save the configuration will correct the MTU.
3. In SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178 version, dynamic port breakout fails after using default configuration file generated by sonic-cfggen and restore.sh.

Example : Dynamic Port Breakout

Procedure:

Step 1. Check IP interface and portchannel status before doing port breakout. In this example, we may take Ethernet48 to be a breakout port.

Caution: IP address has to be removed from physical interface and VLAN interface. This port has also to be removed from Vlan member. It's now allowed to do port breakout if the switch has portchannel setting.

Remove IP address on Ethernet48

```
admin@sonic:~$ show ip interfaces
Interface    IPv4 address/mask    Admin/Oper
-----------  -------------------  ------------
Ethernet48   192.168.1.1/24       up/up
docker0      240.127.1.1/24       up/down
eth0         188.188.98.25/16     up/up
lo           127.0.0.1/8          up/up
             10.1.0.1/32
admin@sonic:~$ sudo config interface ip remove Ethernet48 192.168.1.1/24
```

Since there is IP address on physical interface "Ethernet48", we have to remove IP address from it.

Remove IP address on Vlan Interface

```
admin@sonic:~$ show vlan brief
+-----------+---------------+------------+----------------+-----------------------+
|   VLAN ID | IP Address    | Ports      | Port Tagging   | DHCP Helper Address   |
+===========+===============+============+================+=======================+
|       100 | 192.168.2.2/24 | Ethernet52 | untagged       |                       |    --> Let's take
Ethernet52 to be a breakout port
+-----------+---------------+------------+----------------+-----------------------+

admin@sonic:~$ show ip int
Interface    IPv4 address/mask    Admin/Oper
-----------  -------------------  ------------
Vlan100      192.168.2.2/24       up/up
docker0      240.127.1.1/24       up/down
eth0         188.188.98.30/16     up/up
lo           127.0.0.1/8          up/up
             10.1.0.1/32
admin@sonic:~$ sudo config interface ip remove Vlan100 192.168.2.2/24
admin@sonic:~$ sudo config vlan member del 100 Ethernet52
```

Since there is IP address on "Vlan100" interface, we have to remove IP address from it. Don't forget to delete this port from Vlan member. In this Vlan example, remove Ethernet52 from Vlan100.

Remove IP address on Vlan Interface

```
admin@sonic:~$ show interface portchannel
Flags: A - active, I - inactive, Up - up, Dw - Down, N/A - not available,
S - selected, D - deselected, * - not synced
No. Team Dev Protocol Ports
----- ----------- ----------- ------------
PortChannel LACP(A)(Dw) Ethernet0(D)


admin@sonic:~$ sudo config portchannel del 1
```

Delete ALL portchannel. This is known issue that **portchannel existence** may let breakout fail.

Step 2. Check breakout option availability and current breakout mode.

Show breakout status

```
admin@sonic:~$ show interfaces breakout
{
        ...omitted...
 "Ethernet48": {
      "index": "49,49,49,49",
      "default_brkout_mode": "1x100G[40G]",
      "child ports": "Ethernet48",
      "breakout_modes": "1x100G[40G], 2x50G, 4x25G, 4x10G",                          --> Available breakout mode
option
      "child port speeds": "100G",
      "Current Breakout Mode": "1x100G[40G]",
      "lanes": "77,78,79,80",
      "alias_at_lanes": "Eth49/1, Eth49/2, Eth49/3, Eth49/4"
  },
.......omitted...
}


admin@sonic:~$ show interfaces breakout current-mode Ethernet48
+-------------+-------------------------+
| Interface   | Current Breakout Mode   |
+=============+=========================+
| Ethernet48  | 1x100G[40G]             |                                          --> current mode of
Ethernet48
+-------------+-------------------------+
```

Step 3. Configure dynamic port breakout

Dynamic Port Breakout

```
admin@sonic:~$ sudo config interface breakout Ethernet48 4x25G            --> Breakout Ethernet48 from 1x100G
to 4x25G

Do you want to Breakout the port, continue? [y/N]: y

Running Breakout Mode : 4x25G
Target Breakout Mode : 1x100G[40G]

Ports to be deleted :
{
"Ethernet50": "25000",
"Ethernet51": "25000",
"Ethernet49": "25000",
"Ethernet48": "25000"
}
Ports to be added :
{
"Ethernet48": "100000"
}

After running Logic to limit the impact

Final list of ports to be deleted :
{
"Ethernet50": "25000",
"Ethernet51": "25000",
"Ethernet49": "25000",
"Ethernet48": "25000"
}
Final list of ports to be added :
{
"Ethernet48": "100000"
}
Loaded below Yang Models
['sonic-acl', 'sonic-extension', 'sonic-interface', 'sonic-loopback-interface', 'sonic-port', 'sonic-portchannel', 'sonic-types',
'sonic-vlan']
Note: Below table(s) have no YANG models:
DEVICE_METADATA, FLEX_COUNTER_TABLE, BREAKOUT_CFG, VERSIONS, TELEMETRY,
Below Config can not be verified, It may cause harm to the system
{
"BREAKOUT_CFG": {
"Ethernet48": {
"brkout_mode": "4x25G"
}
}
}
Do you wish to Continue? [y/N]: y
Breakout process got successfully completed.
Please note loaded setting will be lost after system reboot. To preserve setting, run `config save`.
admin@sonic:~$
```

Step 4. Show breakout current mode status and interface status

Breakout Status

```
admin@sonic:~$ show interfaces breakout current-mode Ethernet48
+-------------+-------------------------+
| Interface | Current Breakout Mode     |
+=============+=========================+
| Ethernet48 | 4x25G                    |
+-------------+-------------------------+

admin@sonic:~$ show interfaces status
Interface Lanes Speed MTU Alias Vlan Oper Admin Type Asym PFC
----------- --------------- ------- ----- ------- ------ ------ ------- ------ ----------
...... omitted
Ethernet48 77 25G  N/A Eth49/1 routed up up N/A N/A                              --> Do config reload after
save it to get MTU corrected
Ethernet49 78 25G 9100 Eth49/2 routed up up N/A N/A
Ethernet50 79 25G 9100 Eth49/3 routed up up N/A N/A
Ethernet51 80 25G 9100 Eth49/4 routed up up N/A N/A

...... omitted
```

Step 5. Save configuration to startup configuration. Please do config reload to get first MTU corrected (Step 4).

Dynamic Port Breakout

```
admin@sonic:~$ sudo config save -y
Running command: /usr/local/bin/sonic-cfggen -d --print-data > /tmp/configWO5wNK.db


admin@sonic:~$ sudo config reload -y



admin@sonic:~$ show interfaces status
Interface Lanes Speed MTU Alias Vlan Oper Admin Type Asym PFC
----------- --------------- ------- ----- ------- ------ ------ ------- ------ ----------
...... omitted
Ethernet48 77 25G 9100 Eth49/1 routed up up N/A N/A
Ethernet49 78 25G 9100 Eth49/2 routed up up N/A N/A
Ethernet50 79 25G 9100 Eth49/3 routed up up N/A N/A
Ethernet51 80 25G 9100 Eth49/4 routed up up N/A N/A

...... omitted
```

Step 6. Undo breakout port.

Dynamic Port Breakout

```
admin@sonic:~$ sudo config interface breakout Ethernet48 1x100G[40G]         --> Undo Ethernet48 from 4x25G to
1x100G

Do you want to Breakout the port, continue? [y/N]: y

Running Breakout Mode : 4x25G
Target Breakout Mode : 1x100G[40G]

Ports to be deleted :
{
"Ethernet50": "25000",
"Ethernet51": "25000",
"Ethernet49": "25000",
"Ethernet48": "25000"
}
Ports to be added :
{
"Ethernet48": "100000"
}

After running Logic to limit the impact

Final list of ports to be deleted :
{
"Ethernet50": "25000",
"Ethernet51": "25000",
"Ethernet49": "25000",
"Ethernet48": "25000"
}
Final list of ports to be added :
{
"Ethernet48": "100000"
}
Loaded below Yang Models
['sonic-acl', 'sonic-extension', 'sonic-interface', 'sonic-loopback-interface', 'sonic-port', 'sonic-portchannel', 'sonic-types',
'sonic-vlan']
Note: Below table(s) have no YANG models:
DEVICE_METADATA, FLEX_COUNTER_TABLE, BREAKOUT_CFG, VERSIONS, TELEMETRY,
Below Config can not be verified, It may cause harm to the system
{
"BREAKOUT_CFG": {
"Ethernet48": {
"brkout_mode": "4x25G"
}
}
}
Do you wish to Continue? [y/N]: y
Breakout process got successfully completed.
Please note loaded setting will be lost after system reboot. To preserve setting, run `config save`.
admin@sonic:~$
```

Step 7. Show current breakout mode and interface status.

Breakout Status

```
admin@sonic:~$ show interfaces breakout current-mode Ethernet48
+-------------+-------------------------+
| Interface | Current Breakout Mode     |
+=============+=========================+
| Ethernet48 | 1x100G[40G]              |
+-------------+-------------------------+

admin@sonic:~$ show interfaces status
Interface Lanes Speed MTU Alias Vlan Oper Admin Type Asym PFC
----------- --------------- ------- ----- ------- ------ ------ ------- ------ ----------
...... omitted
Ethernet48 77,78,79,80 100G N/A Eth49/1 routed down up N/A N/A                    --> Do config reload after
save it to get MTU corrected
Ethernet52 85,86,87,88 100G 9100 Eth50/1 routed down up N/A N/A
Ethernet56 93,94,95,96 100G 9100 Eth51/1 routed down up N/A N/A
Ethernet60 97,98,99,100 100G 9100 Eth52/1 routed down up N/A N/A
...... omitted
```

LAG (Link Aggregation)

Example: Establish LAG (Link Aggregation)

Restriction:

1. Both community SONiC and ecSONiC support only LACP. (Does NOT support static link aggregation)
2. LACP is from open source project - **teamd** in SONiC. It follows IEEE802.3ad and supports LACP only for front physical ports and breakout ports.
3. Before LACP establish, the LAG blocks all packets. (LACP fallback = False)
4. SONiC LACP rate(LACP PDU interval) is 30 seconds(slow mode) in default configuration. And it's NOT configurable.
5. SONiC Fast-Reboot(Fast-Reload) requires LACP slow mode for all LAG interface. (https://github.com/Azure/SONiC/wiki/Fast-Reboot)

Reference model:

- Switch model name: AS5835-54X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200507_052107_ec201911_74

Topology:

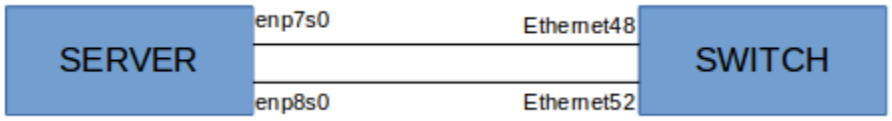

Procedure :

Steps 1. Checking the port configuration

Here's the command to check the port configuration.

```
admin@sonic:~$ show vlan brief
admin@sonic:~$ show ip interfaces
admin@sonic:~$ show interfaces status
```

Caution: have to remove VLAN and IP address from the LAG member ports before LAG Establish. Refer to the following articles.

1. VLAN
2. IP Address

Steps 2. Create a PortChannel interface (LACP fallback = False)

```
admin@sonic:~$ sudo config portchannel add PortChannel1
```

(Optional) LACP fallback = True configuration:

```
admin@sonic:~$ sudo config portchannel add PortChannel1 --fallback=true
```

Caution:

1. LACP Fallback feature is not verified on AS5812-54X
2. If you'd like to change Fallback status from False to True and vice versa, please delete the port channel and create it again with Fallback setting. Otherwise, the fallback setting won't take effect.

Steps 3. Add member ports to PortChannel interface

```
admin@sonic:~$ sudo config portchannel member add PortChannel1 Ethernet48

admin@sonic:~$ sudo config portchannel member add PortChannel1 Ethernet52
```

Steps 4. Save the setting to config_db.json

```
admin@sonic:~$ sudo config save -y
```

Steps 5. Check the appended setting.

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{
    ...
   "PORTCHANNEL": {
        "PortChannel1": {
            "admin_status": "up",
            "mtu": "9100"
        }
    },
    "PORTCHANNEL_MEMBER": {
        "PortChannel1|Ethernet48": {},
        "PortChannel1|Ethernet52": {}
    },
    ...
}
```

Steps 6. Check the setting status on PortChannel.

```
admin@sonic:~$ show interfaces portchannel
Flags: A - active, I - inactive, Up - up, Dw - Down, N/A - not available,
       S - selected, D - deselected, * - not synced
  No.  Team Dev      Protocol     Ports
-----  ------------  -----------  --------------------------
    1  PortChannel1  LACP(A)(Up)  Ethernet52(S) Ethernet48(S)
```

If you see there are no port members (N/A), please return to Step 1.

```
admin@sonic:~$ show interfaces portchannel
Flags: A - active, I - inactive, Up - up, Dw - Down, N/A - not available,
       S - selected, D - deselected, * - not synced
  No.  Team Dev      Protocol     Ports
-----  ------------  -----------  -------
    1  PortChannel1  LACP(A)(Dw)  N/A
```

Steps 7. Check the interface status

```
admin@sonic:~$ show interfaces status | grep "Ethernet48\|Ethernet52\|PortChannel"
  Ethernet48      77,78,79,80     100G   9100     hundredGigE49  PortChannel1      up       up  QSFP28 or
later         N/A
  Ethernet52      85,86,87,88     100G   9100     hundredGigE50  PortChannel1      up       up  QSFP28 or
later         N/A
PortChannel1             N/A     200G   9100               N/A        routed      up       up             N
/A         N/A
```

MTU (Maximum Transmission Unit)

Example:

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Restriction:

1. There's no SONiC command for setting MTU.
2. Software restriction: Maximum of MTU is TBC

Default setting:

1. MTU is 9100 by default

Procedure :
Steps 1. Edit the file config_db.json

Modify the value of MTU per interface (Example: MTU of Ethernet0 is 4000)

```
admin@sonic:~$ sudo vi /etc/sonic/config_db.json
{
omitted...
    "PORT": {
        "Ethernet0": {
            "admin_status": "up",
            "alias": "hundredGigE1",
            "index": "1",
            "lanes": "73,74,75,76",
            "mtu": "4000",
            "speed": "100000"
        },
        "Ethernet4": {
            "admin_status": "up",
            "alias": "hundredGigE2",
            "index": "2",
            "lanes": "65,66,67,68",
            "mtu": "9100",
            "speed": "100000"
        },
omitted...
}
```

Steps 2. config reload or power cycle the switch

```
admin@sonic:~$ sudo config reload -y
```

Steps 3. Check MTU parameter

by SONiC show command.

```
admin@sonic:~$ show interface status | head -n 3
  Interface            Lanes    Speed    MTU           Alias    Vlan    Oper    Admin             Type    Asym
PFC
-----------  ---------------  -------  -----  -------------  ------  ------  -------  ---------------
----------
  Ethernet0      73,74,75,76     100G   4000   hundredGigE1  routed    down       up              N/A         N
/A
```

by Linux IP command.

```
admin@sonic:~$ ip link show |grep Ethernet0
160: Ethernet0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 4000 qdisc pfifo_fast state DOWN mode DEFAULT group
default qlen 1000
```

Appendix:

Check MTU by bcmshell

```
admin@sonic:~$ bcmcmd 'ps'
ps
                ena/        speed/ link auto    STP                 lrn  inter   max   cut   loop
           port  link  Lns   duplex scan neg?   state   pause  discrd ops   face frame  thru?  back
       ce0(  1)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
       ce1(  2)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
       ce2(  3)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
       ce3(  4)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
       ce4(  5)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
       ce5(  6)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
       ce6(  7)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
       ce7(  8)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
       ce8(  9)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
       ce9( 10)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
      ce10( 11)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
      ce11( 12)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
      ce12( 13)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
      ce13( 14)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
      ce14( 15)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
      ce15( 16)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
      ce16( 34)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
      ce17( 35)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
      ce18( 36)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  4022   No
      ce19( 37)  down   4  100G  FD   SW  No   Forward          None   FA  CAUI4  9122   No
..
omitted
drivshell>
```

Notes.

"Ethernet0" maps to the bcm port ce18(36). Have to check config.bcm on /usr/share/sonic/device/x86_64-accton_as7816_64x-r0/Accton-AS7816-64X/.

e.g. Ethernet0 is ce18 (36)

```
admin@sonic:/usr/share/sonic/device/x86_64-accton_as7816_64x-r0/Accton-AS7816-64X$ cat sai.profile
SAI_INIT_CONFIG_FILE=/usr/share/sonic/hwsku/th2-as7816-64x100G.config.bcm

admin@sonic:/usr/share/sonic/device/x86_64-accton_as7816_64x-r0/Accton-AS7816-64X$ cat th2-as7816-64x100G.
config.bcm

omitted...
#Port0
#FC18
portmap_36=73:100
phy_chain_rx_lane_map_physical{73.0}=0x3210
phy_chain_tx_lane_map_physical{73.0}=0x3021
phy_chain_rx_polarity_flip_physical{73.0}=0x0
phy_chain_rx_polarity_flip_physical{74.0}=0x0
phy_chain_rx_polarity_flip_physical{75.0}=0x0
phy_chain_rx_polarity_flip_physical{76.0}=0x1
phy_chain_tx_polarity_flip_physical{73.0}=0x0
phy_chain_tx_polarity_flip_physical{74.0}=0x0
phy_chain_tx_polarity_flip_physical{75.0}=0x1
phy_chain_tx_polarity_flip_physical{76.0}=0x0
omitted...
```

Port Speed configuration

Examples:

1. Configure 10G interface (SFP+) to be 1G interface
2. Configure 100G interface (QSFP+) to be 40G interface

Restriction:

1. Require additional settings for Port Type changes. i.g
   - Breakout Cable QSFP+ to 4 x SFP+/SFP28
   - AS7326-56X 25G to 10G (refer to this article)
   - AS7816-64X 100G to 10G/25G (refer to this article)
2. **nown issue:** It does not return failed if configure wrong speed on interface

   e.g: configure1000Gon the interface

```
admin@sonic:~$ sudo config interface speed Ethernet2 1000000

  admin@sonic:~$ show interfaces status | head -n 5
  Interface            Lanes    Speed    MTU          Alias    Vlan    Oper    Admin
Type    Asym PFC
-----------  ---------------  -------  -----  -------------  ------  ------  -------  ---------------
----------
  Ethernet0                73      10G   9100         swp001   trunk      up       up   QSFP+ or
later         N/A
  Ethernet1                74      10G   9100         swp002   trunk      up       up   QSFP+ or
later         N/A
  Ethernet2                75    1000G   9100         swp003  routed    down       up              N
/A         N/A
```

Reference model:

- Switch model name: AS5835-54X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20191003_122002_ec201904_175

Procedure:

Steps 1. Configure 10G interface (SFP+) to be 1G interface.

Ethernet0 and Ethernet1 as an example.

```
admin@sonic:~$ sudo config interface speed Ethernet0 1000
admin@sonic:~$ sudo config interface speed Ethernet1 1000
```

Caution:

1. Please make sure **port name** (i.e Ethernet0) is correct. Otherwise it does not work.
2. Make sure **Tx_Disable** is **Disabled** on 10G interface. Otherwise, the 10G port can NOT link UP even speed configuration is correct. (refer to this article)
3.
```
admin@sonic:~$ echo 0 | sudo tee /sys/bus/i2c/devices/3-0061/module_tx_disable_1
admin@sonic:~$ echo 0 | sudo tee /sys/bus/i2c/devices/3-0061/module_tx_disable_2
```

Steps 2. Configure 100G interface (QSFP) to be 40G interface.

Ethernet48 and Ethernet52 as an example.

```
admin@sonic:~$ sudo config interface speed Ethernet48 40000
admin@sonic:~$ sudo config interface speed Ethernet52 40000
```

Caution: please make sure **port name** is correct. Otherwise it does not work.

Steps 3. Save the configuration to config_db.json

```
admin@sonic:~$ sudo config save -y
```

Steps 4. Check the port speed setting on config_db.json

```
admin@sonic:~$ sudo vi /etc/sonic/config_db.json
{
   omitted...
   "PORT": {
      "Ethernet0": {
         "admin_status": "up",
         "alias": "tenGigE1",
         "index": "1",
         "lanes": "1",
         "mtu": "9100",
         "speed": "1000"
      },
      "Ethernet1": {
         "admin_status": "up",
         "alias": "tenGigE2",
         "index": "2",
         "lanes": "2",
         "mtu": "9100",
         "speed": "1000"
      },
      omitted...
      "Ethernet48": {
         "admin_status": "up",
         "alias": "hundredGigE49",
         "index": "49",
         "lanes": "37,38,39,40",
         "mtu": "9100",
         "speed": "40000"
      },
      "Ethernet52": {
         "admin_status": "up",
         "alias": "hundredGigE50",
         "index": "53",
         "lanes": "29,30,31,32",
         "mtu": "9100",
         "speed": "40000"
      },
   omitted...
}
```

Steps 5. Check the port speed by SONiC command

```
admin@sonic:~$ show interface status | grep 'Ethernet0 \|Ethernet1 \|Ethernet48\|Ethernet52'
 Ethernet0 1 1G 9100 tenGigE1 routed up up SFP/SFP+/SFP28 N/A
 Ethernet1 2 1G 9100 tenGigE2 routed up up SFP/SFP+/SFP28 N/A
 Ethernet48 37,38,39,40 40G 9100 hundredGigE49 routed up up QSFP+ or later N/A
 Ethernet52 29,30,31,32 40G 9100 hundredGigE50 routed up up QSFP+ or later N/A
```

Appendix:

Check port speed by bcmshell

```
admin@sonic:~$ bcmcmd 'ps'
 ps
 ena/ speed/ link auto STP lrn inter max cut loop
 port link Lns duplex scan neg? state pause discrd ops face frame thru? back
 ge0( 1) up 1 1G FD SW No Forward TX RX None F SGMII 9122 No
 ge1( 2) up 1 1G FD SW No Forward TX RX None F SGMII 9122 No
 xe0( 3) down 1 10G FD SW No Forward None F XFI 9122 No
 xe1( 4) down 1 10G FD SW No Forward None F XFI 9122 No
 xe2( 5) down 1 10G FD SW No Forward None F XFI 9122 No
 xe3( 6) down 1 10G FD SW No Forward None F XFI 9122 No
 xe4( 7) down 1 10G FD SW No Forward None F XFI 9122 No
 xe5( 8) down 1 10G FD SW No Forward None F XFI 9122 No
 xe6( 9) down 1 10G FD SW No Forward None F XFI 9122 No
 xe7( 10) down 1 10G FD SW No Forward None F XFI 9122 No
 xe8( 11) down 1 10G FD SW No Forward None F XFI 9122 No
 xe9( 12) down 1 10G FD SW No Forward None F XFI 9122 No
 xe10( 13) down 1 10G FD SW No Forward None F XFI 9122 No
 xe11( 14) down 1 10G FD SW No Forward None F XFI 9122 No
 xe12( 15) down 1 10G FD SW No Forward None F XFI 9122 No
 xe13( 16) down 1 10G FD SW No Forward None F XFI 9122 No
 xe14( 17) down 1 10G FD SW No Forward None F XFI 9122 No
 xe15( 18) down 1 10G FD SW No Forward None F XFI 9122 No
 xe16( 19) down 1 10G FD SW No Forward None F XFI 9122 No
 xe17( 20) down 1 10G FD SW No Forward None F XFI 9122 No
 xe18( 21) down 1 10G FD SW No Forward None F XFI 9122 No
 xe19( 22) down 1 10G FD SW No Forward None F XFI 9122 No
 xe20( 23) down 1 10G FD SW No Forward None F XFI 9122 No
 xe21( 24) down 1 10G FD SW No Forward None F XFI 9122 No
 xe22( 39) down 1 10G FD SW No Forward None F XFI 9122 No
 xe23( 40) down 1 10G FD SW No Forward None F XFI 9122 No
 xe24( 41) down 1 10G FD SW No Forward None F XFI 9122 No
 xe25( 42) down 1 10G FD SW No Forward None F XFI 9122 No
 xe26( 43) down 1 10G FD SW No Forward None F XFI 9122 No
 xe27( 44) down 1 10G FD SW No Forward None F XFI 9122 No
 xe28( 45) down 1 10G FD SW No Forward None F XFI 9122 No
 xe29( 46) down 1 10G FD SW No Forward None F XFI 9122 No
 xe30( 47) down 1 10G FD SW No Forward None F XFI 9122 No
 xe31( 48) down 1 10G FD SW No Forward None F XFI 9122 No
 xe32( 49) down 1 10G FD SW No Forward None F XFI 9122 No
 xe33( 50) down 1 10G FD SW No Forward None F XFI 9122 No
 xe34( 51) down 1 10G FD SW No Forward None F XFI 9122 No
 xe35( 52) down 1 10G FD SW No Forward None F XFI 9122 No
 xe36( 53) down 1 10G FD SW No Forward None F XFI 9122 No
 xe37( 54) down 1 10G FD SW No Forward None F XFI 9122 No
 xe38( 55) down 1 10G FD SW No Forward None F XFI 9122 No
 xe39( 56) down 1 10G FD SW No Forward None F XFI 9122 No
 xe40( 57) down 1 10G FD SW No Forward None F XFI 9122 No
 xe41( 58) down 1 10G FD SW No Forward None F XFI 9122 No
 xe42( 59) down 1 10G FD SW No Forward None F XFI 9122 No
 xe43( 60) down 1 10G FD SW No Forward None F XFI 9122 No
 xe44( 61) down 1 10G FD SW No Forward None F XFI 9122 No
 xe45( 62) down 1 10G FD SW No Forward None F XFI 9122 No
 xe46( 27) up 4 40G FD SW No Forward TX RX None F XLAUI 9122 No
 xe47( 25) up 4 40G FD SW No Forward TX RX None F XLAUI 9122 No
 ce0( 26) down 4 100G FD SW No Forward None F CAUI4 9122 No
 ce1( 35) down 4 100G FD SW No Forward None F CAUI4 9122 No
 ce2( 34) down 4 100G FD SW No Forward None F CAUI4 9122 No
 ce3( 33) down 4 100G FD SW No Forward None F CAUI4 9122 No
 xe48( 64) down 1 10G FD SW No Forward None FA XFI 9412 No

drivshell>
```

Transceiver status and DDMI

Example:

Reference model:

- Switch model name: AS5835-54X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200616_052146_ec201911_104

Default setting:

- SFP/SFP+/SFP28 transmit power is disable (Refer to this article)
- QSFP/QSFP+/QSFP28 is high power mode.

Example: Presence of SFP transceiver

```
admin@sonic:~$ show interfaces transceiver presence
Port Presence
---------- -----------
Ethernet0 Present
Ethernet1 Present
Ethernet2 Not present
Ethernet3 Not present
omitted...
```

Example: Status of low-power mode of QSFP/QSFP+/QSFP28 transceiver

```
admin@sonic:~$ show interface transceiver lpmode
Port Low-power Mode
---------- ----------------
Ethernet0 Off
Ethernet1 Off
Ethernet2 Off
omitted...
```

Example: Information stored on the EEPROM on SFP transceiver

```
admin@sonic:~$ show interface transceiver eeprom
Ethernet0: SFP EEPROM detected
Connector: LC
Encoding: 64B/66B
Extended Identifier: GBIC/SFP defined by twowire interface ID
Extended RateSelect Compliance: Unspecified
Identifier: SFP/SFP+/SFP28
LengthOM3(UnitsOf10m): 30
Nominal Bit Rate(100Mbs): 103
Specification compliance:
10GEthernetComplianceCode: 10G Base-SR
Vendor Date Code(YYYY-MM-DD Lot): 2018-02-02
Vendor Name: Edgecore
Vendor OUI: 70-72-cf
Vendor PN: ET5402-SR
Vendor Rev: 01
Vendor SN: J11805000385

Ethernet1: SFP EEPROM Not detected

omitted..
```

Optionally, you can get DDMI (Digital Diagnostics Monitoring) by specifying the -d or --dom flag.

```
admin@sonic:~$ show interface transceiver eeprom Ethernet0 -d

Ethernet0: SFP EEPROM detected

Connector: LC

Encoding: 64B/66B

Extended Identifier: GBIC/SFP defined by twowire interface ID

Extended RateSelect Compliance: Unspecified

Identifier: SFP/SFP+/SFP28
```

```
LengthOM3(UnitsOf10m): 30
Nominal Bit Rate(100Mbs): 103
Specification compliance:
10GEthernetComplianceCode: 10G Base-SR
Vendor Date Code(YYYY-MM-DD Lot): 2018-02-02
Vendor Name: Edgecore
Vendor OUI: 70-72-cf
Vendor PN: ET5402-SR
Vendor Rev: 01
Vendor SN: J11805000409
MonitorData:
RXPower: -2.6922dBm
TXBias: 6.2400mA
TXPower: -3.4208dBm
Temperature: 28.7188C
Vcc: 3.2976Volts
ThresholdData:
TempHighAlarm : OffC
TempHighWarning: OffC
TempLowAlarm : OffC
TempLowWarning : OffC
VccHighAlarm : OffVolts
VccHighWarning : OffVolts
VccLowAlarm : OffVolts
VccLowWarning : OffVolts
RxPowerHighAlarm : OffdBm
RxPowerHighWarning: OffdBm
RxPowerLowAlarm : OffdBm
RxPowerLowWarning : OffdBm
TxBiasHighAlarm : OffmA
TxBiasHighWarning : OffmA
TxBiasLowAlarm : OffmA
TxBiasLowWarning : OffmA
TxPowerHighAlarm : OffdBm
TxPowerHighWarning: OffdBm
TxPowerLowAlarm : OffdBm
TxPowerLowWarning : OffdBm
```

Note.DDMI thresholddata isn't display correctly, this issue is submit to Jira (SONIC-876) and tracking on zendesk(#9060)

- [x] SONIC-876 - [TS] Trnasceiver thresholdData BACKLOG

QoS (Quality of Service)

Quality of Service (QoS)

1. Queue counter
2. CoS and Queue mapping
3. DSCP and Queue mapping (TBD)
4. QoS Scheduling (TBD)

Example: Queue Counter

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Procedure:

Steps 1. Clear the queue counter

```
admin@sonic:~$ sonic-clear queuecounters
```

Steps 2. Check Ethernet0 queue counter

```
admin@sonic:~$ show queue counters Ethernet0
Last cached time was 2016-11-07 00:54:58.623429
     Port    TxQ    Counter/pkts    Counter/bytes    Drop/pkts    Drop/bytes
---------  -----  --------------  ---------------  -----------  ------------
Ethernet0    UC0               0                0            0             0
Ethernet0    UC1               0                0            0             0
Ethernet0    UC2               0                0            0             0
Ethernet0    UC3               0                0            0             0
Ethernet0    UC4               0                0            0             0
Ethernet0    UC5               0                0            0             0
Ethernet0    UC6               0                0            0             0
Ethernet0    UC7               0                0            0             0
Ethernet0    UC8             N/A              N/A          N/A           N/A
Ethernet0    UC9             N/A              N/A          N/A           N/A
Ethernet0   MC10               0                0            0             0
Ethernet0   MC11               0                0            0             0
Ethernet0   MC12               0                0            0             0
Ethernet0   MC13               0                0            0             0
Ethernet0   MC14               0                0            0             0
Ethernet0   MC15               0                0            0             0
Ethernet0   MC16               0                0            0             0
Ethernet0   MC17               0                0            0             0
Ethernet0   MC18             N/A              N/A          N/A           N/A
Ethernet0   MC19             N/A              N/A          N/A           N/A
```

Notes:

1. 8 dedicated queues (i.e UC0, UC1, ..., UC7) for unicast.
2. Other 8 dedicated queues (i.e MC10, MC11, ..., UC17) for multicast.
3. UC8, UC9, UC18, UC19 are not available.

Example: CoS and Queue mapping

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200116_051623_ec201904_128

Known issue:

AS7816-64X breakout cable doesn't support CoS mapping, no matter which CoS priority of packets you sent, they all belong to Queue 0.

Default setting:

1. Port priority is 0 by default. i.g Untagged packet (without CoS priority)
2. CoS and Queue mapping table

CoS	Queue (UC0 ~ UC7)
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7

SONiC RoCEv1 Tuning Tutorial

Reference model:

- Switch model name: AS7726-32X

Topology:

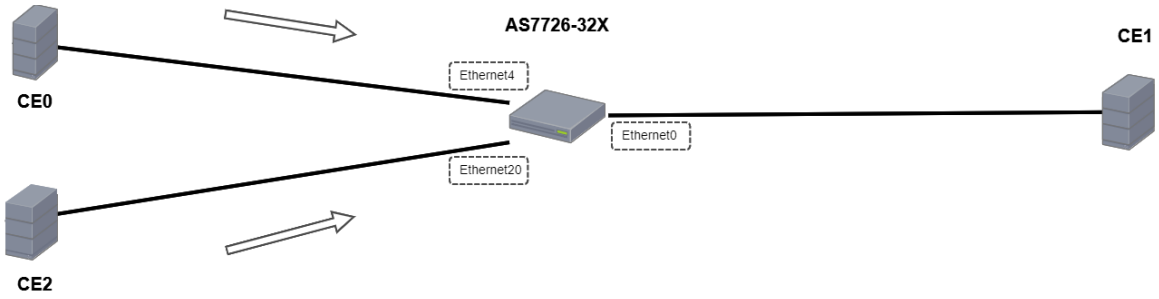

Tuning Procedure:

Step 1. Measure bandwidth performance and latency of each traffic with Perftest tool. In this topology example, measure bandwidth performance of CE0 → CE1 and CE2 → CE1. From this measurement, you may get the big picture of current real performance and how to tune it later. Please check your traffic PFC queue number. It's necessary for configuring port QOS setting on next step.

Step 2. Configure TC to priority group map setting. Set up TC-to-Priority group map by adding below group map configuration into config_db.json. By default, we will use PTP to classify traffic to queue.

TC to priority goup mapping

```
"TC_TO_PRIORITY_GROUP_MAP": {
    "AZURE": {                                                    → group map profile name
        "0": "0",
        "1": "0",
        "2": "0",
        "3": "3",
        "4": "4",
        "5": "0",
        "6": "0",
        "7": "7"
    }
},
```

Step 3. Enable PFC on queue. In this example, we will enable queue 3 of PFC on Ethernet0, Ethernet4, Ethernet20 respectively.

PORT QOS Setting

```
"PORT_QOS_MAP": {
    "Ethernet0": {
        "pfc_enable": "3",                                                  → Enable PFC 3
        "tc_to_pg_map": "[TC_TO_PRIORITY_GROUP_MAP|AZURE]"          → Bind TC-to-Priority group map profile to Ethernet0
    },
    "Ethernet4": {
        "pfc_enable": "3",                                                  → Enable PFC 3
        "tc_to_pg_map": "[TC_TO_PRIORITY_GROUP_MAP|AZURE]"          → Bind TC-to-Priority group map profile to Ethernet4
    },
    "Ethernet20": {
        "pfc_enable": "3",                                                  → Enable PFC 3
        "tc_to_pg_map": "[TC_TO_PRIORITY_GROUP_MAP|AZURE]"          → Bind TC-to-Priority group map profile to Ethernet20
    }
}
```

Step 4. Buffer pool for each model as attached on Appendix 1 below. If you can't find it, please request this buffer pool from us by submit a ticket through support system https://support.edge-core.com/hc/en-us/requests/new. Add this buffer pool setting to config_db.json file. In this example, AS7726-32X buffer pool setting as shown below.

Buffer Pool setting

```
"BUFFER_POOL": {                                                    → This is global buffer pool setting which is used for all ports and queues
    "ingress_lossless_pool": {
        "mode": "static",
        "size": "10875072",
        "type": "ingress",
        "xoff": "4194112"
    }
},
```

Step 5. Configure and tune buffer setting profile. Configuration below is based on our best tuning configuration of AS7726-32X. The details of how to calculate this configuration are on below.

Buffer Profile setting

```
"BUFFER_PROFILE": {

    "pg_300m_profile": {                                                                    → This buffer setting profile will be bound to certain or/and all ports and queues (Step 5)

        "pool": "[BUFFER_POOL|ingress_lossless_pool]",

         "size": "69632",

         "static_th": "0",

         "xoff": "13056",

         "xon": "512"

    }

},
```

The picture above is buffer design of **one ingress-queue on a port**. Xon/ Xoff uses headroom. When the profile.size is full, it will use headroom space for packet buffering. Once profile.xoff buffer is full (buffer line touch the upper limit of profile.xoff in green block), it sends 802.1Qbb pause frame to partner and requests to stop transmission. And when the buffer is less then profile.xon, it sends request to partner to resume the traffic transmission. AS7726-32X has 32 ports. In condition of 1:1 traffic, it will reach the most congested traffic with 16 ports traffic transmission to single port.

In order to achieve the **best bandwitdh performance along with low latency and lossless packet**, we have to calculate some buffer parameters.

1. MTU_cell = MTU/ 256 * 256 → Cell size : 256 (**This cell size is for AS7726-32X. Please check appendix 2 below to match cell size of your switch model.**). Rounding up the division result of RoCE MTU size by 256.
2. profile.size = 16 * MTU_cell → Tuning profile.size to maximum size. In this example, 16 is the max ports number which would let AS7726-32X reach the most congested traffic.
3. profile.xoff = 3 * MTU_cell → Tuning profile.xoff size to best performance. In this example, 3 is the best start constant.
4. profile.xon = 2 * MTU_cell

Please check the MTU size of your NIC card. In this example, we use the maximum size of RoCE's packets is 4174 bytes. Using formula above, calculate the parameter setting for this example,

MTU_cell = 4174 / 256 *256 = 4352

profile.size = 4352 * 16 = 69632

profile.xoff = 4352 * 3 = 13056

profile.xon = 256 * 2 = 512

Tips in tuning : if there is packet drop, try to tune profile size. If bandwidth result is low, try to tune profile.xoff larger.

Step 6. Bind buffer profile (Step 4) to port and queue number then check the result. Compare performance and latency result after tuning and before tuning (Step 1). Tune again by follow **tips in tuning** on step 4 till you get a best result which means the bandwidth is better than before but doesn't increase latency.

Buffer priority group binding

```
"BUFFER_PG": {

"Ethernet20|3": {                                                                    → Port|P
FC Queue number

"profile": "[BUFFER_PROFILE|pg_300m_profile]"                                        → Bind pg_300m_profile
buffer profile to Ethernet20 Queue3

},

"Ethernet4|3": {

"profile": "[BUFFER_PROFILE|pg_300m_profile]"

},

"Ethernet0|3": {

"profile": "[BUFFER_PROFILE|pg_300m_profile]"

}

},
```

Reference:

Test Result Comparison Table:

Test Item	Single RoCE QP		Two Tuned RoCE QP
	CE2 -> CE1	CE0 -> CE1	CE2&CE0->CE1
Bandwidth (Mbps)	76.51	83.32	79.13
Latency (µs)	2.046	2.126	2.048

Note :

1. QP : Queue Pair
2. Tuning two RoCE QPs with lossless packet can achieve bandwidth with almost approaching the average bandwidth of two single RoCE QPs bandwidth and latency.
3. Bandwidth and latency results are average of 5 times test.

Appendix 1:

Buffer Pool

Switch Model	Buffer Pool
AS4630-54PE	*To be provided later
AS5835-54X	*To be provided later

AS7326-56X	<pre>"BUFFER_POOL": {
	    "ingress_lossless_pool": {
	        "mode": "dynamic",
	        "size": "12766208",
	        "type": "ingress",
	        "xoff": "196608"
	    }
	},</pre>

AS7726-32X	<pre>"BUFFER_POOL": {
	    "ingress_lossless_pool": {
	        "mode": "dynamic",
	        "size": "12766208",
	        "type": "ingress",
	        "xoff": "196608"
	    }
	},</pre>

AS7816-64X	<pre>"BUFFER_POOL": {
	    "ingress_lossless_pool": {
	        "mode": "dynamic",
	        "size": "12766208",
	        "type": "ingress",
	        "xoff": "4625920"
	    }
	},</pre>

AS9716-32D	<pre>"BUFFER_POOL": {
	    "ingress_lossless_pool": {
	        "mode": "dynamic",
	        "size": "59001152",
	        "type": "ingress",
	        "xoff": "7428992"
	    }
	},</pre>

AS8000	*To be provided later

Appendix 2:

Cell Size Table

Switch Model	Chipset Model	Chipset Version	Packet Buffer Cell Size
AS4630-54PE	Helix 5	BCM56371	256 bytes
AS5835-54X	Trident 3	BCM56771	256 bytes
AS7326-56X	Trident 3	BCM56873	256 bytes
AS7726-32X	Trident 3	BCM56870	256 bytes
AS7816-64X	Tomahawk 2	BCM56970	208 bytes
AS9716-32D	Tomahawk 3	BCM56980	254 bytes
AS8000	Tomahawk 3	BCM56980	254 bytes

Routing

BGP Step 1 - Establish BGP Session

Example: Establish BGP session

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC-OS-Edgecore-SONiC_20200116_051623_ec201904_128

Procedure:

Step 1. FRR config initialization (refer to this article)

Step 2. Configure IP address properly for BGP session (refer to this article)

Step 3. Enter Vty shell.

```
admin@sonic:~$ vtysh

Hello, this is FRRouting (version 7.0.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic#
```

Note: *vtysh* provides a combined frontend to all FRR daemons in a single combined session.

Step 4. Assign the Router ID

```
Switch#1
sonic# configure terminal
sonic(config)# router-id 1.1.1.1

Switch#2
sonic# configure terminal
sonic(config)# router-id 2.2.2.2
```

Note: Router ID is used to identify routers in AS(autonomous systems).

Caution: Router ID must be unique for each BGP router in the network. Otherwise, it's failed to establish BGP session.

Step 5. Configuring BGP Neighbor Session (EBGP or IBGP)

```
EBGP: EBGP is peering between two different AS.
Switch#1's setting.
sonic# configure terminal sonic(config)# router bgp 65101 --> Assign the AS number
sonic(config-router)# neighbor 10.10.10.3 remote-as 65102 --> Setting the neighbor's IP and AS number

Switch#2's setting.
sonic# configure terminal
sonic(config)# router bgp 65102
sonic(config-router)# neighbor 10.10.10.2 remote-as 65101

IBGP: IBGP is peering between same AS.
Switch#1's setting.
sonic# configure terminal sonic(config)# router bgp 65101
sonic(config-router)# neighbor 10.10.10.3 remote-as 65101

Switch#2's setting.
sonic# configure terminal
sonic(config)# router bgp 65101
sonic(config-router)# neighbor 10.10.10.2 remote-as 65101
```

Note:
AS (autonomous systems) is a unique identifier and it used to exchange routing information with other AS.
Its available pool of 16-bit AS numbers.
1 ~ 64511 reserve to the public, you need to apply for it form IANA (Internet Assigned Numbers Authority)
64512~ 65534 reserve for private.
Since AS numbers are not enough, IANA extended the AS Number field to 32 bits in size.
131072 - 4199999999 reserve for public and 4200000000-4294967294 reserve for private.

Step 6. Checking the neighbor status.

On the AS7816-64X

```
sonic# show bgp summary

IPv4 Unicast Summary:
BGP router identifier 1.1.1.1, local AS number 65101 vrf-id 0
BGP table version 0
RIB entries 0, using 0 bytes of memory
Peers 1, using 21 KiB of memory

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down State/PfxRcd
10.10.10.3      4      65102       4       4        0    0    0 00:01:57            0

Total number of neighbors 1

sonic# show bgp neighbors
BGP neighbor is 10.10.10.3, remote AS 65102, local AS 65101, external link
Hostname: sonic
  BGP version 4, remote router ID 2.2.2.2, local router ID 1.1.1.1
  BGP state = Established, up for 00:02:02
  Last read 00:00:02, Last write 00:00:02
  Hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    4 Byte AS: advertised and received
    AddPath:
      IPv4 Unicast: RX advertised IPv4 Unicast and received
    Route refresh: advertised and received(old & new)
    Address Family IPv4 Unicast: advertised and received
    Hostname Capability: advertised (name: sonic,domain name: n/a) received (name: sonic,domain name: n/a)
    Graceful Restart Capabilty: advertised and received
      Remote Restart timer is 120 seconds
      Address families by peer:
        none
  Graceful restart information:
    End-of-RIB send: IPv4 Unicast
    End-of-RIB received: IPv4 Unicast
  Message statistics:
    Inq depth is 0
    Outq depth is 0
                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                1          1
    Keepalives:             3          3
    Route Refresh:          0          0
    Capability:             0          0
    Total:                  5          5
  Minimum time between advertisement runs is 0 seconds

 For address family: IPv4 Unicast
  Update group 2, subgroup 2
  Packet Queue length 0
  Community attribute sent to this neighbor(all)
  0 accepted prefixes

  Connections established 1; dropped 0
  Last reset never
Local host: 10.10.10.2, Local port: 179
Foreign host: 10.10.10.3, Foreign port: 48044
Nexthop: 10.10.10.2
Nexthop global: fe80::ba6a:97ff:fe19:c246
Nexthop local: fe80::ba6a:97ff:fe19:c246
BGP connection: shared network
BGP Connect Retry Timer in Seconds: 120
Read thread: on  Write thread: on
```

Step 7. Save the routing setting.

```
sonic# write
Note: this version of vtysh never writes vtysh.conf
Building Configuration...
Configuration saved to /etc/frr/zebra.conf
Configuration saved to /etc/frr/bgpd.conf
Configuration saved to /etc/frr/staticd.conf
```

BGP Step 2.1 Redistribute routes to BGP process (Optional)

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC-OS-Edgecore-SONiC_20200116_051623_ec201904_128

Procedure: Enter Vty shell and BGP ASN

```
sonic# configure terminal
sonic(config)# router bgp 65101
sonic(config-router)# address-family ipv4 unicast  // Enter address-family ipv4
sonic(config-router-af)# redistribute connected
sonic(config-router-af)# redistribute kernel
sonic(config-router-af)# redistribute static
sonic(config-router)# end
```

Note: This configuration example says that redistribute connected route, kernel route and static route to BGP process

Address-family: Multiprotocol extensions enable BGP to carry routing information for multiple network layer protocols. BGP supports an Address Family Identifier (AFI) for IPv4 and IPv6.

```
sonic(config-router)# address-family
ipv4 Address Family
ipv6 Address Family
l2vpn Address Family [TBD]
vpnv4 Address Family [TBD]
vpnv6 Address Family [TBD]
```

Note: Route redistribution is used to learn from one routing protocol (static, connected route, and OSPF ....) available to a different routing protocol.

```
sonic(config-router-af)# redistribute
  babel       Babel routing protocol (Babel) [TBD]
  connected   Connected routes (directly attached subnet or host)
  eigrp       Enhanced Interior Gateway Routing Protocol (EIGRP) [TBD]
  isis        Intermediate System to Intermediate System (IS-IS) [TBD]
  kernel      Kernel routes (not installed via the zebra RIB)
  nhrp        Next Hop Resolution Protocol (NHRP) [TBD]
  openfabric  OpenFabric Routing Protocol [TBD]
  ospf        Open Shortest Path First (OSPFv2) [TBD]
  rip         Routing Information Protocol (RIP) [TBD]
  sharp       Super Happy Advanced Routing Protocol (sharpd) [TBD]
  static      Statically configured routes
  table       Non-main Kernel Routing Table [TBD]
  vnc         Virtual Network Control (VNC) [TBD]
  vnc-direct  VNC direct (not via zebra) routes [TBD]
```

Note: Kernel routes and static routes are different, please refer to this article.

Note: Currently, SONiC FRRouting supports redistribute connected, kernel and static only.

BGP Step 2 - Add the announcement network

Example: Add the announcement network

Reference model:

- Switch model name: AS7816-64X
- Edgecore SONiC version: SONiC-OS-Edgecore-SONiC_20200116_051623_ec201904_128

Procedure: Enter Vty shell and BGP ASN

```
Switch#1
sonic# configure terminal
sonic(config)# router bgp 65101
sonic(config-router)# network 192.168.1.0 mask 255.255.255.0
sonic(config-router)# end

Switch#2
sonic# configure terminal
sonic(config)# router bgp 65102
sonic(config-router)# network 192.168.2.0 mask 255.255.255.0
sonic(config-router)# end
```

Note: This configuration example says that network 192.168.1.0/24 and 192.168.2.0/24 will be announced to all neighbors.

Note: (Optional)Redistribute routing entires such as local network, static route (refer to this article)

Step 1. Check network result

```
sonic# show ip bgp ipv4
BGP table version is 2, local router ID is 1.1.1.1, vrf id 0
Default local pref 100, local AS 65101
Status codes:  s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes:  i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 192.168.1.0/24   0.0.0.0                  0         32768 i
*> 192.168.2.0/24   10.10.10.3               0             0 65102 i

Displayed  2 routes and 2 total paths
```

Note: Network 192.168.1.0/24 Next Hop 0.0.0.0 this is local network

Step 2. Save the setting.

```
sonic# write
Note: this version of vtysh never writes vtysh.conf
Building Configuration...
Configuration saved to /etc/frr/zebra.conf
Configuration saved to /etc/frr/bgpd.conf
Configuration saved to /etc/frr/staticd.conf
```

BGP Unnumbered

Reference model:

- Switch model name: AS7726-32X, AS7326-56X
- SONiC Software Version: Edgecore-SONiC_20200722_070543_ec201911_141

Example: Establish BGP unnumbered session

Procedure:

Step 1. Initialize FRR configuration. Please refer to FRRouting and config initialization tutorial.

Step 2. Configure IP address properly for BGP session. Please refer to Management and front port IPv4/IPv6 Address tutorial.

Step 3. Enter Vty shell.

```
admin@sonic:~$ vtysh

Hello, this is FRRouting (version 7.0.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic#
```

Note: *vtysh* provides a combined frontend to all FRR daemons in a single combined session.

Step 4. Configuring BGP Neighbor Session

```
Switch#1's setting.
sonic# configure terminal

sonic(config)# router bgp 65000 --> Assign AS number
sonic(config-router)# bgp router-id 1.1.1.1 --> Assign router ID
sonic(config-router)# neighbor TEST peer-group --> Assign peer group tag
sonic(config-router)# neighbor TEST remote-as external --> Assign peer group tag
sonic(config-router)# neighbor Ethernet0 interface peer-group TEST --> Assign interface port to peer group tag



Switch#2's setting.

sonic# configure terminal

sonic(config)# router bgp 65001 --> Assign AS number
sonic(config-router)# bgp router-id 2.2.2.2 --> Assign router ID
sonic(config-router)# neighbor TEST peer-group --> Assign peer group tag
sonic(config-router)# neighbor TEST remote-as external --> Assign peer group tag
sonic(config-router)# neighbor Ethernet56 interface peer-group TEST --> Assign interface port to peer group tag
```

Note:

1. AS (autonomous systems) is a unique identifier and it used to exchange routing information with other AS.
   Its available pool of 16-bit AS numbers.
   1 ~ 64511 reserve to the public, you need to apply for it form IANA (Internet Assigned Numbers Authority)
   64512~ 65534 reserve for private.
   Since AS numbers are not enough, IANA extended the AS Number field to 32 bits in size.
   131072 - 4199999999 reserve for public and 4200000000-4294967294 reserve for private.
2. Router ID is used to identify routers in AS(autonomous systems).
   Caution: Router ID must be unique for each BGP router in the network. Otherwise, it's failed to establish BGP session.

Step 5. Checking the BGP connection and neighbor status.

On the AS7726-32X (Switch1)

```
sonic# show bgp summary

IPv4 Unicast Summary:
BGP router identifier 1.1.1.1, local AS number 65000 vrf-id 0
BGP table version 0
RIB entries 0, using 0 bytes of memory
Peers 1, using 20 KiB of memory
Peer groups 1, using 64 bytes of memory

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down State/PfxRcd
Ethernet0       4      65001      24      24        0    0    0 00:21:31            0

Total number of neighbors 1


sonic# show bgp neighbors
BGP neighbor on Ethernet0: fe80::6a21:5fff:fedc:3ac8, remote AS 65001, local AS 65000, external link
Hostname: sonic
  Member of peer-group TEST for session parameters
  BGP version 4, remote router ID 2.2.2.2, local router ID 1.1.1.1
  BGP state = Established, up for 00:23:08
  Last read 00:00:08, Last write 00:00:08
  Hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    4 Byte AS: advertised and received
    AddPath:
      IPv4 Unicast: RX advertised IPv4 Unicast and received
    Route refresh: advertised and received(old & new)
    Address Family IPv4 Unicast: advertised and received
    Hostname Capability: advertised (name: sonic,domain name: n/a) received (name: sonic,domain name: n/a)
    Graceful Restart Capabilty: advertised and received
      Remote Restart timer is 120 seconds
      Address families by peer:
        none
  Graceful restart information:
    End-of-RIB send: IPv4 Unicast
    End-of-RIB received: IPv4 Unicast
  Message statistics:
    Inq depth is 0
    Outq depth is 0
                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                1          1
    Keepalives:            24         24
    Route Refresh:          0          0
    Capability:             0          0
    Total:                 26         26
  Minimum time between advertisement runs is 0 seconds

 For address family: IPv4 Unicast
  Update group 1, subgroup 1
  Packet Queue length 0
  Community attribute sent to this neighbor(all)
  0 accepted prefixes

  Connections established 1; dropped 0
  Last reset 00:32:39, Waiting for NHT
Local host: fe80::6f8:f8ff:fe9b:d0d8, Local port: 58716
Foreign host: fe80::6a21:5fff:fedc:3ac8, Foreign port: 179
Nexthop: 1.1.1.1
Nexthop global: ffe80::6f8:f8ff:fe9b:d0d8
Nexthop local: fe80::6f8:f8ff:fe9b:d0d8
BGP connection: shared network
BGP Connect Retry Timer in Seconds: 120
Read thread: on  Write thread: on  FD used: 22
```

On the AS7326-56X (Switch2)

```
sonic# show bgp summary

IPv4 Unicast Summary:
BGP router identifier 2.2.2.2, local AS number 65001 vrf-id 0
BGP table version 0
RIB entries 0, using 0 bytes of memory
Peers 1, using 20 KiB of memory
Peer groups 1, using 64 bytes of memory

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down State/PfxRcd
Ethernet56      4      65000      46      47        0    0    0 00:43:48            0

Total number of neighbors 1


sonic# show bgp neighbors
BGP neighbor on Ethernet56: fe80::6f8:f8ff:fe9b:d0d8, remote AS 65000, local AS 65001, external link
Hostname: sonic
  Member of peer-group TEST for session parameters
  BGP version 4, remote router ID 1.1.1.1, local router ID 2.2.2.2
  BGP state = Established, up for 00:45:08
  Last read 00:00:52, Last write 00:00:52
  Hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    4 Byte AS: advertised and received
    AddPath:
      IPv4 Unicast: RX advertised IPv4 Unicast and received
    Route refresh: advertised and received(old & new)
    Address Family IPv4 Unicast: advertised and received
    Hostname Capability: advertised (name: sonic,domain name: n/a) received (name: sonic,domain name: n/a)
    Graceful Restart Capabilty: advertised and received
      Remote Restart timer is 120 seconds
      Address families by peer:
        none
  Graceful restart information:
    End-of-RIB send: IPv4 Unicast
    End-of-RIB received: IPv4 Unicast
  Message statistics:
    Inq depth is 0
    Outq depth is 0
                         Sent       Rcvd
    Opens:                  2          1
    Notifications:          0          0
    Updates:                1          1
    Keepalives:            46         46
    Route Refresh:          0          0
    Capability:             0          0
    Total:                 49         48
  Minimum time between advertisement runs is 0 seconds

 For address family: IPv4 Unicast
  Update group 1, subgroup 1
  Packet Queue length 0
  Community attribute sent to this neighbor(all)
  0 accepted prefixes

  Connections established 1; dropped 0
  Last reset 00:56:15, No AFI/SAFI activated for peer
Local host: fe80::6a21:5fff:fedc:3ac8, Local port: 179
Foreign host: fe80::6f8:f8ff:fe9b:d0d8, Foreign port: 58716
Nexthop: 2.2.2.2
Nexthop global: fe80::6a21:5fff:fedc:3ac8
Nexthop local: fe80::6a21:5fff:fedc:3ac8
BGP connection: shared network
BGP Connect Retry Timer in Seconds: 120
Read thread: on  Write thread: on  FD used: 22
```

Step 6. Save the routing setting.

```
sonic# write
Note: this version of vtysh never writes vtysh.conf
Building Configuration...
Configuration saved to /etc/frr/zebra.conf
Configuration saved to /etc/frr/bgpd.conf
Configuration saved to /etc/frr/staticd.conf
```

FRRouting and config initialization

Reference model:

- Switch model name: AS5835-54X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200507_052107_ec201911_74

Restriction:

1. [**Mandatory**] FRR is an routing protocol suite for Linux platform. It has other config files instead of /etc/sonic/config_db.json. However, default config_db.json is incorrect and incomplete. Please refer to the following procedures to initialize FRRouting and config
2. FRR does not support Integrated configuration mode. Both Community SONiC and ecSONiC couldn't integrate the individual files(/etc/sonic/frr/bgpd.conf, staticd.conf, zebra.conf) to the single file. (frr.conf), since frr being handled via supervisord and not by watchfrr. https://github.com/Azure/sonic-buildimage/pull/3486

Note:

FRR configuration folder: /etc/sonic/frr

```
bgpd.conf staticd.conf zebra.conf
```

FRR Integrated configuration mode:

```
root@sonic:/etc/sonic/frr# cat vtysh.conf
no service integrated-vtysh-config
```

Caution: Integrated configuration mode is configurable. However, it won't take effect due to software restriction.

Initialize FRR and config (for Restriction #1)

Step 1: Edit /etc/sonic/config_db.json

- Add // to comment out this object "BGP_NEIGHBOR": { }
- Add "docker_routing_config_mode": "split" to "DEVICE_METADATA" as shown below,

```
{
    "//BGP_NEIGHBOR": {
        "10.0.0.1": {
            "asn": "65200",
            "holdtime": "180",
            "keepalive": "60",
            "local_addr": "10.0.0.0",
            "name": "ARISTA01T2",
            "nhopself": 0,
            "rrclient": 0
        },
        omitted...
    },
    omitted...
    "DEVICE_METADATA": {
        "localhost": {
        "bgp_asn": "65100",
        "hostname": "sonic",
        "hwsku": "Accton-AS5835-54X",
        "mac": "80:a2:35:d2:47:b5",
        "platform": "x86_64-accton_as5835_54x-r0",
        "type": "LeafRouter",
        "docker_routing_config_mode": "split"
        }
    },
   omitted...
}
```

Step 2. Delete 3 files bgpd.conf, staticd.conf and zebra.conf on /etc/sonic/frr/.

Note: Need Root authority

Caution: Do NOT delete /etc/sonic/frr/vtysh.conf

- Method 1: delete 3 files one by one

```
admin@sonic:~$ sudo ls -ls /etc/sonic/frr
total 28
16 -rw-r--r-- 1 300 300 16001 Aug 17 00:17 bgpd.conf
4 -rw-r--r-- 1 300 300 352 Aug 17 00:17 staticd.conf
4 -rw-r----- 1 300 300 35 Aug 17 00:17 vtysh.conf
4 -rw-r--r-- 1 300 300 2345 Aug 17 00:17 zebra.conf

admin@sonic:~$ sudo rm /etc/sonic/frr/bgpd.conf
admin@sonic:~$ sudo rm /etc/sonic/frr/staticd.conf
admin@sonic:~$ sudo rm /etc/sonic/frr/zebra.conf

admin@sonic:~$ sudo ls -ls /etc/sonic/frr
total 4
4 -rw-r----- 1 300 300 35 Aug 17 00:17 vtysh.conf
admin@sonic:~$
```

- Method 2: delete all files and create vtysh.conf manually

```
admin@sonic:~$ sudo ls -ls /etc/sonic/frr
total 16
4 -rw-r----- 1 300 300 150 May 20 08:36 bgpd.conf
4 -rw-r----- 1 300 300 148 May 20 08:36 staticd.conf
4 -rw-r--r-- 1 root root 35 May 20 08:36 vtysh.conf
4 -rw-r----- 1 300 300 158 May 20 08:36 zebra.conf
admin@sonic:~$ sudo rm /etc/sonic/frr/*
admin@sonic:~$ sudo ls -ls /etc/sonic/frr
total 0
admin@sonic:~$ echo "no service integrated-vtysh-config" | sudo tee /etc/sonic/frr/vtysh.conf
admin@sonic:~$ sudo ls -ls /etc/sonic/frr
total 4
4 -rw-r--r-- 1 root root 35 May 20 08:37 vtysh.conf
admin@sonic:~$
```

Note: regnerate vtysh.confg by this command "echo "no service integrated-vtysh-config" | sudo tee /etc/sonic/frr/vtysh.conf"

Step 3. Reload config or power cycle the switch

```
admin@sonic:~$ sudo config reload -y
```

Step 4. Verify the result

- Only vtysh.conf exists on /etc/sonic/frr/.

```
admin@sonic:~$ sudo ls -l /etc/sonic/frr
total 4
-rw-r----- 1 300 300 35 Aug 17 00:17 vtysh.conf
admin@sonic:~$
```

Caution: If the result is failed, please go back to Step 1 and redo these steps.

MCLAG

MC-LAG (L2 MC-LAG)

Reference model:

- Switch model name: AS7326-56X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200722_070543_ec201911_141

Restriction:

1. There is no SONiC command to add MC-LAG
2. There can be only one MC-LAG domain configured in the system.

Example 1 : Establish L2 MC-LAG

Topology:

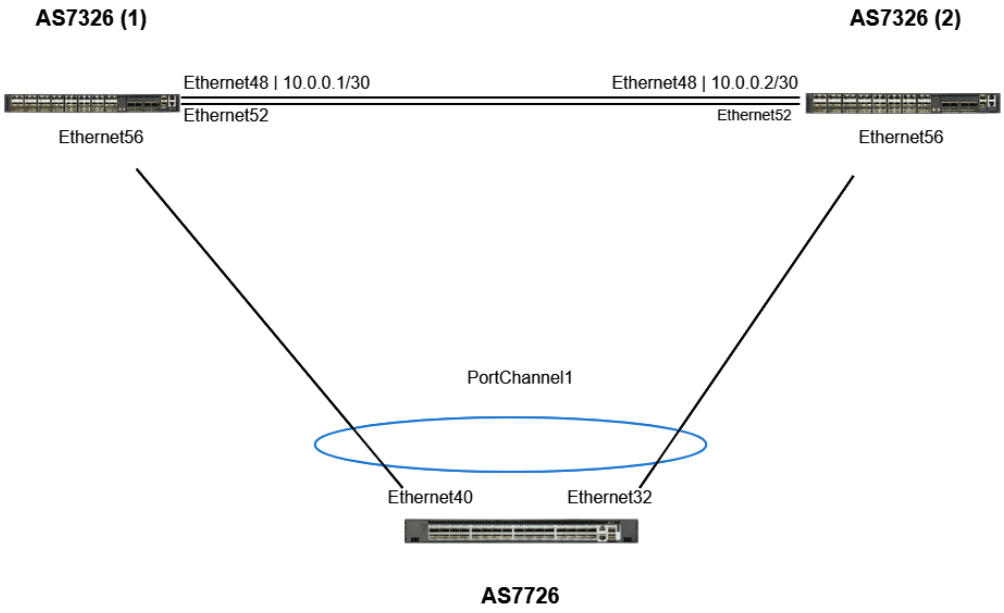

Procedure :

Step 1: Create port-channel on all of the switches. Please refer to LAG (Link-Aggregation) tutorial.

Step 2: Create and add the VLAN to all of the Interfaces. Please refer to VLAN & Inter-VLAN Routing tutorial.

For example: Add VLAN 10 to PortChannel1, and Ethernet52.

Step 3: Add the IP address on Ethernet48. Please refer to Management and front port IPv4/IPv6 Address tutorial.

Step 4: Create an MC-LAG domain on config_db.json

```
admin@sonic:/etc/sonic$ sudo vi config_db.json
{
    ...
    "MC_LAG": {
        "1": {
            "local_ip": "10.0.0.1",
            "peer_ip": "10.0.0.2",
            "peer_link": "Ethernet52",
            "mclag_interface": "PortChannel1"
        }
    },
    ...
}
```

Note:

"1" is the ID of the MC_LAG domain ID.

Key "local_ip" is source IPv4 address to be used for MC-LAG session. This is local address.

Key "peer_ip" is peer IPv4 address to be used for MC-LAG session

Key "peer_link" is mandatory for L2 config. this links is used to carry data traffic when MC-LAG interface is down.

Key "mclag_interface" is MC-LAG member port, it needs to be portchannels.

Step 5: Reload config and check MC-LAG status.

AS7326 (1)

```
admin@7326:~$ sudo config reload -y
...omitted


admin@7326:~$ mclagdctl dump state
The MCLAG's keepalive is: OK
Domain id: 1
Local Ip: 10.0.0.1
Peer Ip: 10.0.0.2
Peer Link Interface: Ethernet52
Peer Link Mac: 80:a2:35:4f:4f:40
Role: Active
MCLAG Interface: PortChannel1
Loglevel: NOTICE
```

AS7326 (2)

```
admin@7326:~$ sudo config reload -y
...omitted


admin@7326:~$ mclagdctl dump state
The MCLAG's keepalive is: OK
Domain id: 1
Local Ip: 10.0.0.2
Peer Ip: 10.0.0.1
Peer Link Interface: Ethernet52
Peer Link Mac: 80:a2:35:4e:a2:08
Role: Standby
MCLAG Interface: PortChannel1
Loglevel: NOTICE
```

EVPN - VxLAN

1. EPVN L2 VxLAN
2. EVPN L3 VxLAN
    a. Asymmetric
    b. Symmetric

Asymmetric EVPN IRB

Reference model:

- Switch model name: AS7326-56X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178

Known issue:

1. Creating VLAN-VNI mapping table (Procedure Step 5) **will cause the container crash**.
   Temporary workaround is by **adding IP address on VLAN** (Procedure step 2).
2. The VxLAN's operation status **will always in "oper_down"** even the VxLAN work normally.

Example : Asymmetric EVPN IRB

Topology:

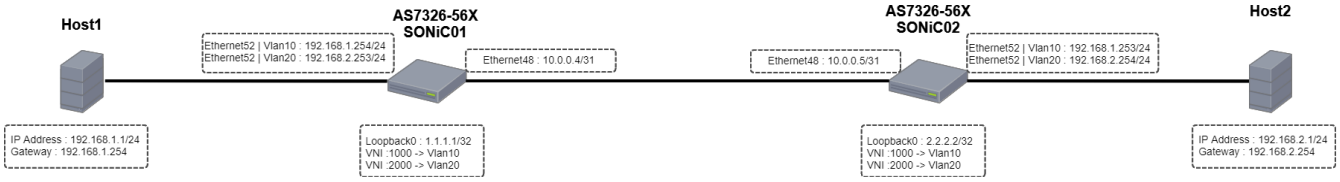

Note :

- BGP, EVPN and VxLAN related articles are available in here, Routing (BGP) , EVPN & VxLAN.
- All VNI is belong to L2 VNI.

Procedure:

Step 1. Setup VLAN 10 and VLAN 20. Add Ethernet52 of both switches to VLAN 10 and VLAN 20. Please refer to VLAN & Inter-VLAN Routing article.

Step 2. Configure IP address to VLAN 10 and VLAN 20. In this example as shown on topology diagram above, configure the IP address to VLAN 10 and VLAN 20 of both switches. There's known issue as listed above which requires IP address on VLAN. Please refer to VLAN & Inter-VLAN Routing article.

Step 3. Configure IP address to both Ethernet48 of both switches.

SONiC01

```
admin@SONIC01:~$sudo config interface ip add Ethernet48 10.0.0.4/31
```

SONiC02

```
admin@SONIC02:~$sudo config interface ip add Ethernet48 10.0.0.5/31
```

Step 4: Configure IP address to Loopback0 of both switches.

SONiC01

```
admin@SONIC01:~$ sudo config interface ip remove Loopback0 10.1.0.1/32
admin@SONIC01:~$ sudo config interface ip add Loopback0 1.1.1.1/32
```

SONiC02

```
admin@SONIC02:~$ sudo config interface ip remove Loopback0 10.1.0.1/32
admin@SONIC02:~$ sudo config interface ip add Loopback0 2.2.2.2/32
```

Step 5. Create VxLAN

SONiC01

```
admin@SONIC01:~$sudo config vxlan add vtep 1.1.1.1                          → configuring VTEP_name (vtep) and its IP address

admin@SONIC01:~$sudo config vxlan evpn_nvo add nvo vtep                     → create nvo_name (nvo) and bind it to VTEP_name (vtep)

admin@SONIC01:~$sudo config vxlan map add vtep 10 1000                      → mapping VNI 1000 to VLAN 10

admin@SONIC01:~$sudo config vxlan map add vtep 20 2000                      → mapping VNI 2000 to VLAN 20

admin@SONIC01:~$sudo config save -y
```

SONiC02

```
admin@SONIC01:~$sudo config vxlan add vtep 2.2.2.2                          → configuring VTEP_name (vtep) and its IP address

admin@SONIC01:~$sudo config vxlan evpn_nvo add nvo vtep                     → create nvo_name (nvo) and bind it to VTEP_name (vtep)

admin@SONIC01:~$sudo config vxlan map add vtep 10 1000                      → mapping VNI 1000 to VLAN 10

admin@SONIC01:~$sudo config vxlan map add vtep 20 2000                      → mapping VNI 2000 to VLAN 20

admin@SONIC01:~$sudo config save -y
```

Step 6. Establish BGP environment for EVPN.

SONiC01(VTYSH)

```
admin@7726:~$ vtysh                                                                 → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal

sonic(config)#router bgp 65100                                                      → assign BGP AS number

sonic(config-router)#neighbor 10.0.0.5 interface remote-as 65100            → assign Ethernet48 IP address to connect via IBGP with peering on same AS

sonic(config-router)# address-family ipv4 unicast                                   → Enter address-family ipv4

sonic(config-router-af)# network 1.1.1.1/32                                         → Announce 1.1.1.1 network

sonic(config-router-af)# exit

sonic(config-router)#address-family l2vpn evpn                                      → enter EVPN setting

sonic(config-router-af)#neighbor 10.0.0.5 activate                                  → activate EVPN for neighbor 10.0.0.5

sonic(config-router-af)#advertise-all-vni                                           → advertise all VNI routing

sonic(config-router-af)#exit
```

SONiC02(VTYSH)

```
admin@7726:~$ vtysh                                                                    → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal

sonic(config)#router bgp 65100                                                         → assign BGP AS number

sonic(config-router)#neighbor 10.0.0.4 interface remote-as 65100          → assign Ethernet48 IP address to connect via
IBGP with peering on same AS

sonic(config-router)# address-family ipv4 unicast                                  → Enter address-family ipv4

sonic(config-router-af)# network 2.2.2.2/32                                          → Announce 2.2.2.2 network

sonic(config-router-af)# exit

sonic(config-router)#address-family l2vpn evpn                                     → enter EVPN setting

sonic(config-router-af)#neighbor 10.0.0.4 activate                                  → activate EVPN for neighbor 10.0.0.4

sonic(config-router-af)#advertise-all-vni                                              → advertise all VNI routing

sonic(config-router-af)#exit
```

Step 7. Check EVPN-VNI status.

SONiC01(VTYSH)

```
sonic# show evpn vni
VNI         Type          VxLAN IF         # MACs         # ARPs          # Remote VTEPs         Tenant VRF
1000         L2            vtep-
10               2                  5                    1                              default
2000         L2            vtep-
20               2                  5                    1                              default



sonic# show evpn vni detail
VNI: 1000
Type: L2
Tenant VRF: default
VxLAN interface: vtep-10
VxLAN ifIndex: 317
Local VTEP IP: 1.1.1.1
Mcast group: 0.0.0.0
Remote VTEPs for this VNI:
2.2.2.2 flood: HER
Number of MACs (local and remote) known for this VNI: 2
Number of ARPs (IPv4 and IPv6, local and remote) known for this VNI: 5
Advertise-gw-macip: No

VNI: 2000
Type: L2
Tenant VRF: default
VxLAN interface: vtep-20
VxLAN ifIndex: 318
Local VTEP IP: 1.1.1.1
Mcast group: 0.0.0.0
Remote VTEPs for this VNI:
2.2.2.2 flood: HER
Number of MACs (local and remote) known for this VNI: 2
Number of ARPs (IPv4 and IPv6, local and remote) known for this VNI: 5
Advertise-gw-macip: No
```

SONiC02(VTYSH)

```
sonic# show evpn vni
VNI         Type          VxLAN IF         # MACs         # ARPs          # Remote VTEPs         Tenant VRF
1000         L2            vtep-
10               2                  5                    1                              default
2000         L2            vtep-
20               2                  5                    1                              default



sonic# show evpn vni detail
VNI: 1000
Type: L2
Tenant VRF: default
VxLAN interface: vtep-10
VxLAN ifIndex: 71
Local VTEP IP: 2.2.2.2
Mcast group: 0.0.0.0
Remote VTEPs for this VNI:
1.1.1.1 flood: HER
Number of MACs (local and remote) known for this VNI: 2
Number of ARPs (IPv4 and IPv6, local and remote) known for this VNI: 5
Advertise-gw-macip: No

VNI: 2000
Type: L2
Tenant VRF: default
VxLAN interface: vtep-20
VxLAN ifIndex: 72
Local VTEP IP: 2.2.2.2
Mcast group: 0.0.0.0
Remote VTEPs for this VNI:
1.1.1.1 flood: HER
Number of MACs (local and remote) known for this VNI: 2
Number of ARPs (IPv4 and IPv6, local and remote) known for this VNI: 5
Advertise-gw-macip: No
```

Step 8. Check BGP EVPN status.

SONiC01(VTYSH)

```
sonic# show bgp l2vpn evpn
BGP table version is 12, local router ID is 192.168.2.253
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path
*> [2]:[0]:[48]:[8c:ea:1b:30:da:47]
1.1.1.1 32768 i
ET:8 RT:65100:1000
*> [2]:[0]:[48]:[8c:ea:1b:30:da:47]:[32]:[192.168.1.1]
1.1.1.1 32768 i
ET:8 RT:65100:1000
*> [2]:[0]:[48]:[8c:ea:1b:30:da:47]:[128]:[fe80::5037:63f4:3aba:bbe9]
1.1.1.1 32768 i
ET:8 RT:65100:1000
*> [3]:[0]:[32]:[1.1.1.1]
1.1.1.1 32768 i
ET:8 RT:65100:1000
Route Distinguisher: ip 192.168.2.253:3

*> [3]:[0]:[32]:[1.1.1.1]
1.1.1.1 32768 i
ET:8 RT:65100:2000
Route Distinguisher: ip 192.168.2.254:2

*>i[3]:[0]:[32]:[2.2.2.2]
2.2.2.2 100 0 i
RT:65100:1000 ET:8
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:4b]
2.2.2.2 100 0 i
RT:65100:2000 ET:8
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:4b]:[32]:[192.168.2.1]
2.2.2.2 100 0 i
RT:65100:2000 ET:8
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:4b]:[128]:[fe80::2a29:a798:37d3:76c2]
2.2.2.2 100 0 i
RT:65100:2000 ET:8
*>i[3]:[0]:[32]:[2.2.2.2]
2.2.2.2 100 0 i
RT:65100:2000 ET:8

Displayed 10 out of 10 total prefixes
```

SONiC02(VTYSH)

```
sonic# show bgp l2vpn evpn
BGP table version is 12, local router ID is 192.168.2.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:47]
1.1.1.1 100 0 i
RT:65100:1000 ET:8
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:47]:[32]:[192.168.1.1]
1.1.1.1 100 0 i
RT:65100:1000 ET:8
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:47]:[128]:[fe80::5037:63f4:3aba:bbe9]
1.1.1.1 100 0 i
RT:65100:1000 ET:8
*>i[3]:[0]:[32]:[1.1.1.1]
1.1.1.1 100 0 i
RT:65100:1000 ET:8
Route Distinguisher: ip 192.168.2.253:3

*>i[3]:[0]:[32]:[1.1.1.1]
1.1.1.1 100 0 i
RT:65100:2000 ET:8
Route Distinguisher: ip 192.168.2.254:2

*> [3]:[0]:[32]:[2.2.2.2]
2.2.2.2 32768 i
ET:8 RT:65100:1000
*> [2]:[0]:[48]:[8c:ea:1b:30:da:4b]
2.2.2.2 32768 i
ET:8 RT:65100:2000
*> [2]:[0]:[48]:[8c:ea:1b:30:da:4b]:[32]:[192.168.2.1]
2.2.2.2 32768 i
ET:8 RT:65100:2000
*> [2]:[0]:[48]:[8c:ea:1b:30:da:4b]:[128]:[fe80::2a29:a798:37d3:76c2]
2.2.2.2 32768 i
ET:8 RT:65100:2000
*> [3]:[0]:[32]:[2.2.2.2]
2.2.2.2 32768 i
ET:8 RT:65100:2000

Displayed 10 out of 10 total prefixes
```

Step 9. Check VNI MAC learning.

SONiC01(VTYSH)

```
sonic# show evpn mac vni all

VNI 1000 #MACs (local and remote) 2

MAC Type Intf/Remote VTEP VLAN Seq #'s
04:f8:f8:6a:f6:91 local Vlan10 10 0/0
8c:ea:1b:30:da:47 local Ethernet52 10 0/0

VNI 2000 #MACs (local and remote) 2

MAC Type Intf/Remote VTEP VLAN Seq #'s
04:f8:f8:6a:f6:91 local Vlan20 20 0/0
8c:ea:1b:30:da:4b remote 2.2.2.2 0/0
```

SONiC02(VTYSH)

```
sonic# show evpn mac vni all

VNI 1000 #MACs (local and remote) 2

MAC Type Intf/Remote VTEP VLAN Seq #'s
04:f8:f8:6b:06:91 local Vlan10 10 0/0
8c:ea:1b:30:da:47 remote 1.1.1.1 0/0

VNI 2000 #MACs (local and remote) 2

MAC Type Intf/Remote VTEP VLAN Seq #'s
04:f8:f8:6b:06:91 local Vlan20 20 0/0
8c:ea:1b:30:da:4b local Ethernet52 20 0/0
```

Step 10. Check ARP learning table. Hosts ARP shall be learnt by both switches.

SONiC01(VTYSH) arp learning

```
sonic# show evpn arp-cache vni all

VNI 1000 #ARP (IPv4 and IPv6, local and remote) 5

IP Type State MAC Remote VTEP Seq #'s
fe80::6f8:f8ff:fe6b:691 local inactive 04:f8:f8:6b:06:91 0/0
fe80::5037:63f4:3aba:bbe9 local active 8c:ea:1b:30:da:47 0/0
192.168.1.254 local active 04:f8:f8:6a:f6:91 0/0
fe80::6f8:f8ff:fe6a:f691 local active 04:f8:f8:6a:f6:91 0/0
192.168.1.1 local active 8c:ea:1b:30:da:47 0/0

VNI 2000 #ARP (IPv4 and IPv6, local and remote) 5

IP Type State MAC Remote VTEP Seq #'s
192.168.2.1 remote active 8c:ea:1b:30:da:4b 2.2.2.2 0/0
fe80::6f8:f8ff:fe6b:691 local inactive 04:f8:f8:6b:06:91 0/0
fe80::6f8:f8ff:fe6a:f691 local active 04:f8:f8:6a:f6:91 0/0
192.168.2.253 local active 04:f8:f8:6a:f6:91 0/0
fe80::2a29:a798:37d3:76c2 remote active 8c:ea:1b:30:da:4b 2.2.2.2 0/0
```

SONiC02(VTYSH) arp learning

```
sonic# show evpn arp-cache vni all

VNI 1000 #ARP (IPv4 and IPv6, local and remote) 5

IP Type State MAC Remote VTEP Seq #'s
fe80::6f8:f8ff:fe6b:691 local active 04:f8:f8:6b:06:91 0/0
fe80::5037:63f4:3aba:bbe9 remote active 8c:ea:1b:30:da:47 1.1.1.1 0/0
fe80::6f8:f8ff:fe6a:f691 local inactive 04:f8:f8:6a:f6:91 0/0
192.168.1.253 local active 04:f8:f8:6b:06:91 0/0
192.168.1.1 remote active 8c:ea:1b:30:da:47 1.1.1.1 0/0

VNI 2000 #ARP (IPv4 and IPv6, local and remote) 5

IP Type State MAC Remote VTEP Seq #'s
192.168.2.1 local active 8c:ea:1b:30:da:4b 0/0
fe80::6f8:f8ff:fe6b:691 local active 04:f8:f8:6b:06:91 0/0
192.168.2.254 local active 04:f8:f8:6b:06:91 0/0
fe80::6f8:f8ff:fe6a:f691 local inactive 04:f8:f8:6a:f6:91 0/0
fe80::2a29:a798:37d3:76c2 local active 8c:ea:1b:30:da:4b 0/0
```

EVPN ARP/ND Suppression

Reference model:

- Switch model name: AS7326-56X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178

Known issue:

1. Creating VLAN-VNI mapping table (Procedure Step 5) **will cause the container crash**.
   Temporary workaround is by **adding IP address on VLAN** (Procedure step 2).
2. The VxLAN's operation status **will always in "oper_down"** even the VxLAN work normally.
3. Doing config reload after adding suppression attribute in config_db.json file **will cause VLAN operation down**. Please do workaround procedure on step 7 below.

Example : EVPN ARP/ND suppression

Purpose:

- In order to suppress ARP request flooding on uplink port, switch will help to reply the ARP reply if the local EVPN database has the ARP entries.

Topology:

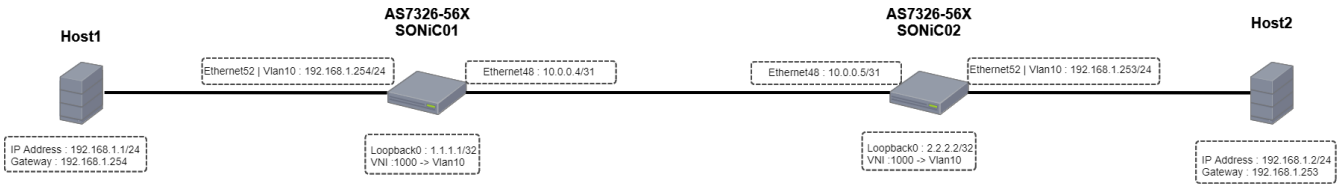

Procedure:

Step 1. Setup VLAN 10. Add Ethernet52 of both switches to VLAN 10. Please refer to VLAN & Inter-VLAN Routing article.

Step 2. Configure IP address to VLAN 10. In this example as shown on topology diagram above, configure the IP address to VLAN 10 of both switches. There's known issue as listed above which requires IP address on VLAN. Please refer to VLAN & Inter-VLAN Routing article.

Step 3. Configure IP address to both Ethernet48 of both switches.

SONiC01

```
admin@SONIC01:~$sudo config interface ip add Ethernet48 10.0.0.4/31
```

SONiC02

```
admin@SONIC02:~$sudo config interface ip add Ethernet48 10.0.0.5/31
```

Step 4: Configure IP address to Loopback0 of both switches.

SONiC01

```
admin@SONIC01:~$ sudo config interface ip remove Loopback0 10.1.0.1/32
admin@SONIC01:~$ sudo config interface ip add Loopback0 1.1.1.1/32
```

SONiC02

```
admin@SONIC02:~$ sudo config interface ip remove Loopback0 10.1.0.1/32
admin@SONIC02:~$ sudo config interface ip add Loopback0 2.2.2.2/32
```

Step 5. Create VxLAN

SONiC01

```
admin@SONIC01:~$sudo config vxlan add vtep 1.1.1.1                          → configuring VTEP_name (vtep) and its IP address

admin@SONIC01:~$sudo config vxlan evpn_nvo add nvo vtep                     → create nvo_name (nvo) and bind it to VTEP_name (vtep)

admin@SONIC01:~$sudo config vxlan map add vtep 10 1000                      → mapping VNI 1000 to VLAN 10

admin@SONIC01:~$sudo config save -y
```

SONiC02

```
admin@SONIC01:~$sudo config vxlan add vtep 2.2.2.2                          → configuring VTEP_name (vtep) and its IP address

admin@SONIC01:~$sudo config vxlan evpn_nvo add nvo vtep                     → create nvo_name (nvo) and bind it to VTEP_name (vtep)

admin@SONIC01:~$sudo config vxlan map add vtep 10 1000                      → mapping VNI 1000 to VLAN 10

admin@SONIC01:~$sudo config save -y
```

Step 6. Establish BGP environment for EVPN.

SONiC01(VTYSH)

```
admin@7726:~$ vtysh                                                         → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal

sonic(config)#router bgp 65100                                              → assign BGP AS number

sonic(config-router)#neighbor 10.0.0.5 interface remote-as internal         → assign Ethernet48 IP address to connect via IBGP with peering on same AS

sonic(config-router)# address-family ipv4 unicast                           → Enter address-family ipv4

sonic(config-router-af)# network 1.1.1.1/32                                 → Announce 1.1.1.1 network

sonic(config-router-af)# exit

sonic(config-router)#address-family l2vpn evpn                              → enter EVPN setting

sonic(config-router-af)#neighbor 10.0.0.5 activate                          → activate EVPN for neighbor 10.0.0.5

sonic(config-router-af)#advertise-all-vni                                   → advertise all VNI routing

sonic(config-router-af)#exit
```

SONiC02(VTYSH)

```
admin@7726:~$ vtysh                                                          → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal

sonic(config)#router bgp 65100                                               → assign BGP AS number

sonic(config-router)#neighbor 10.0.0.4 interface remote-as internal          → assign Ethernet48 IP address to connect via
IBGP with peering on same AS

sonic(config-router)# address-family ipv4 unicast                            → Enter address-family ipv4

sonic(config-router-af)# network 2.2.2.2/32                                  → Announce 2.2.2.2 network

sonic(config-router-af)# exit

sonic(config-router)#address-family l2vpn evpn                               → enter EVPN setting

sonic(config-router-af)#neighbor 10.0.0.4 activate                           → activate EVPN for neighbor 10.0.0.4

sonic(config-router-af)#advertise-all-vni                                    → advertise all VNI routing

sonic(config-router-af)#exit
```

Step 7. Enable ARP suppression on VLAN10. Please configure the same setting below to second switch.

ARP Suppression Setting

```
admin@SONIC01:~$ sudo config neigh_suppress enable 10

admin@SONIC01:~$ sudo config save -y

admin@SONIC01:~$ sudo config reload -y
```

Note : Please do disable then enable ARP suppression after doing config reload (Please refer to workaround below). It's known issue as listed above that doing config reload after enabling ARP suppression may let VLAN operation down.

Workaround

ARP Suppression Workaround

```
admin@SONIC01:~$ sudo config neigh_suppress disable 10

admin@SONIC01:~$ sudo config neigh_suppress enable 10
```

EVPN L2 VxLAN

Reference model:

- Switch model name: AS7326-56X, AS7726-32X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178

Known issue:

1. Creating VLAN-VNI mapping table (Procedure Step 6) will cause the container crash.
   Temporary workaround is by adding IP address on VLAN (Procedure step 2).
2. The VxLAN's operation status will always in "oper_down" even the VxLAN work normally.

Example : VxLAN L2 EVPN

Topology:

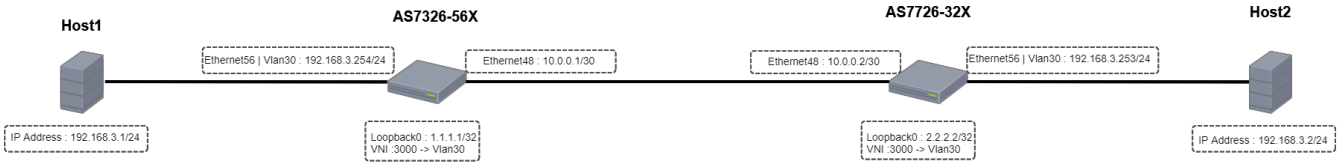

Procedure :

Step1: Create VLAN 30 and add Ethernet56 of both Swicthes to VLAN 30. Please refer to VLAN & Inter-VLAN Routing article.

Step 2: Configure IP address to the VLAN 30. Please refer to VLAN & Inter-VLAN Routing article. In this example, configure 192.168.3.254 /24 and 192.168.3.253/24 to VLAN 30 at both sides.

Step 3: Configure IP address to Loopback0 of both switches.

AS7326

```
admin@7726:~$ sudo config interface ip remove Loopback0 10.1.0.1/32
admin@7726:~$ sudo config interface ip add Loopback0 2.2.2.2/32
```

AS7726

```
admin@7726:~$ sudo config interface ip remove Loopback0 10.1.0.1/32
admin@7726:~$ sudo config interface ip add Loopback0 1.1.1.1/32
```

Step 4: Establish BGP Session between Ethernet48 and announce the network.

AS7326

```
admin@7326:~$ vtysh                                                    → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal
sonic(config)# router bgp 65100                                        → assign BGP AS number
sonic(config-router)# neighbor 10.0.0.2 remote-as 65100   → assign Ethernet48 IP address to connect via IBGP
with peering on same AS
sonic(config-router)# address-family ipv4                              → Enter address-family ipv4
sonic(config-router-af)# network 1.1.1.1/32                            → Announce 1.1.1.1 network
```

AS7726

```
admin@7726:~$ vtysh                                                          → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal
sonic(config)# router bgp 65100                                              → assign BGP AS number
sonic(config-router)# neighbor 10.0.0.1 remote-as 65100                  → assign Ethernet48 IP address to connect via IBGP
with peering on same AS
sonic(config-router)# address-family ipv4                                    → Enter address-family ipv4
sonic(config-router-af)# network 2.2.2.2/32                                  → Announce 2.2.2.2 network
```

Step 5: Announce L2VPN EVPN routes.

```
admin@7326:~$ vtysh                                                          → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal
sonic(config)# router bgp 65100                                              → assign BGP AS number
sonic(config-router)# address-family l2vpn evpn                          → enter EVPN setting
sonic(config-router-af)# neighbor 10.0.0.2 activate                      → activate EVPN for neighbor 10.0.0.2
sonic(config-router-af)# advertise-all-vni                                   → advertise all VNI routing
```

```
admin@7726:~$ vtysh                                                          → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal
sonic(config)# router bgp 65100                                              → assign BGP AS number
sonic(config-router)# address-family l2vpn evpn                          → enter EVPN setting
sonic(config-router-af)# neighbor 10.0.0.1 activate                      → activate EVPN for neighbor 10.0.0.1
sonic(config-router-af)# advertise-all-vni                                   → advertise all VNI routing
```

Note :

1. VNI (VxLAN Network Identifier) : virtual extension of VLAN over IP network.

Step 6: Create VxLAN.

```
admin@7326:~$ sudo config vxlan add vtep 1.1.1.1                             → configuring VTEP_name (vtep) and its IP address
(1.1.1.1)
admin@7326:~$ sudo config vxlan evpn_nvo add nvo vtep                    → create nvo_name (nvo) and bind it to VTEP_name (vtep)
admin@7326:~$ sudo config vxlan map add vtep 30 3000                     → mapping VNI 3000 to VLAN 30.
```

```
admin@7726:~$ sudo config vxlan add vtep 2.2.2.2                             → configuring VTEP_name (vtep) and its IP address
(2.2.2.2)
admin@7726:~$ sudo config vxlan evpn_nvo add nvo vtep                    → create nvo_name (nvo) and bind it to VTEP_name (vtep)
admin@7726:~$ sudo config vxlan map add vtep 30 3000                     → mapping VNI 3000 to VLAN 30.
```

Note :

1. VTEP (VXLAN Tunnel End Point) : an entity that originates and/or terminates VXLAN tunnels which is specified by a source IP address.
   Only one VTEP is allowed on one device. Please use loopback IP address for VTEP's IP address.
2. NVO (Network Virtualization Overlay)
   Only one NVO is allowed on one device.
3. VNI (VxLAN Network Identifier) : virtual extension of VLAN over IP network.

Step 7: Check VxLAN status. Vxlan tunnel will be connected automatically.

```
admin@7726:~$ show vxlan interface
VTEP Information:

       VTEP Name : vtep, SIP : 1.1.1.1
       Source interface : Loopback0

admin@7726:~$ show vxlan tunnel
+---------+---------+-------------------+--------------+
| SIP     | DIP     | Creation Source   | OperStatus   |
+=========+=========+===================+==============+
| 1.1.1.1 | 2.2.2.2 | EVPN              | oper_down    |
+---------+---------+-------------------+--------------+
Total count : 1

admin@7726:~$ show vxlan vlanvnimap
+--------+-------+
| VLAN   | VNI   |
+========+=======+
| Vlan30 | 3000  |
+--------+-------+
Total count : 1
```

Symmetric EVPN IRB

Reference model:

- Switch model name: AS7326-56X
- Edgecore SONiC version: SONiC.Edgecore-SONiC_20200827_110345_ec201911_2020-aug_enhanced_178

Known issue:

1. The VxLAN's operation status **will always in "oper_down"** even the VxLAN work normally.

Example : Symmetric EVPN IRB

Topology:

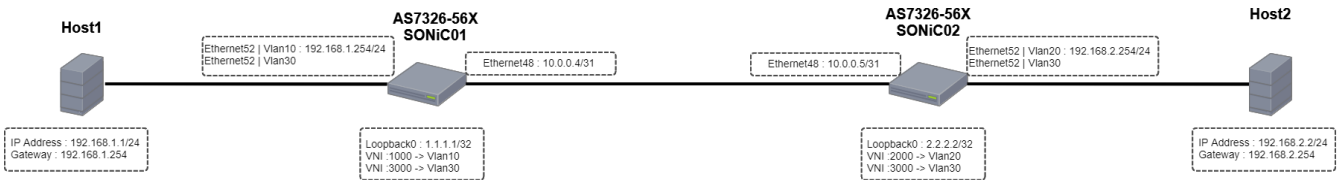

Note :

- BGP, EVPN and VxLAN related articles are available in here, Routing (BGP) , EVPN & VxLAN.

Procedure:

Step 1. Setup VLAN 10 & VLAN 30 on SONiC01 and VLAN 20 & VLAN 30 on SONiC02. Add Ethernet52 of both switches to those created VLAN. Please refer to VLAN & Inter-VLAN Routing article.

Step 2. Configure IP address to both Ethernet48 of both switches.

SONiC01

```
admin@SONIC01:~$sudo config interface ip add Ethernet48 10.0.0.4/31
```

SONiC02

```
admin@SONIC02:~$sudo config interface ip add Ethernet48 10.0.0.5/31
```

Step 3: Configure IP address to Loopback0 of both switches.

SONiC01

```
admin@SONIC01:~$ sudo config interface ip remove Loopback0 10.1.0.1/32
admin@SONIC01:~$ sudo config interface ip add Loopback0 1.1.1.1/32
```

SONiC02

```
admin@SONIC02:~$ sudo config interface ip remove Loopback0 10.1.0.1/32
admin@SONIC02:~$ sudo config interface ip add Loopback0 2.2.2.2/32
```

Step 4. Create VxLAN

SONiC01

```
admin@SONiC01:~$sudo config vxlan add vtep 1.1.1.1                          → configuring VTEP_name (vtep) and its IP address
admin@SONiC01:~$sudo config vxlan evpn_nvo add nvo vtep                     → create nvo_name (nvo) and bind it to VTEP_name (vtep)
admin@SONiC01:~$sudo config vxlan map add vtep 10 1000                      → mapping VNI 1000 to VLAN 10
admin@SONiC01:~$sudo config vxlan map add vtep 30 3000                      → mapping VNI 3000 to VLAN 30
admin@SONiC01:~$sudo config save -y
```

SONiC02

```
admin@SONiC02:~$sudo config vxlan add vtep 2.2.2.2                          → configuring VTEP_name (vtep) and its IP address
admin@SONiC02:~$sudo config vxlan evpn_nvo add nvo vtep                     → create nvo_name (nvo) and bind it to VTEP_name (vtep)
admin@SONiC02:~$sudo config vxlan map add vtep 20 2000                      → mapping VNI 2000 to VLAN 20
admin@SONiC02:~$sudo config vxlan map add vtep 30 3000                      → mapping VNI 3000 to VLAN 30
admin@SONiC02:~$sudo config save -y
```

Step 5. Configure VRF Setting

SONiC01

```
admin@SONiC01:~$ sudo config vrf add Vrf01                                  → create VRF
admin@SONiC01:~$ sudo config interface vrf bind Vlan30 Vrf01                → bind Vlan30 to Vrf01
admin@SONiC01:~$ sudo config interface vrf bind Vlan10 Vrf01                → bind Vlan10 to Vrf01
admin@SONiC01:~$ sudo config vrf add_vrf_vni_map Vrf01 3000                 → map Vrf01 VNI 3000
admin@SONiC01:~$ sudo config interface ip add Vlan10 192.168.1.254/24       → Configure IP address on Vlan10
```

SONiC02

```
admin@SONiC02:~$ sudo config vrf add Vrf01                                  → create VRF
admin@SONiC02:~$ sudo config interface vrf bind Vlan30 Vrf01                → bind Vlan30 to Vrf01
admin@SONiC02:~$ sudo config interface vrf bind Vlan20 Vrf01                → bind Vlan20 to Vrf01
admin@SONiC02:~$ sudo config vrf add_vrf_vni_map Vrf01 3000                 → map Vrf01 VNI 3000
admin@SONiC02:~$ sudo config interface ip add Vlan20 192.168.2.254/24       → Configure IP address on Vlan20
```

Step 6. Save configuration

Both Switches

```
admin@SONiC01:~$ sudo config save -y
```

Step 7. Establish BGP environment for EVPN.

SONiC01(VTYSH)

```
admin@7726:~$ vtysh                                                    → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal

sonic(config)#router bgp 65100                                         → assign BGP AS number

sonic(config-router)#neighbor 10.0.0.5 interface remote-as 65100       → assign Ethernet48 IP address to connect via
IBGP with peering on same AS

sonic(config-router)# address-family ipv4 unicast                      → Enter address-family ipv4

sonic(config-router-af)# network 1.1.1.1/32                            → Announce 1.1.1.1 network

sonic(config-router-af)# exit

sonic(config-router)#address-family l2vpn evpn                         → enter EVPN setting

sonic(config-router-af)#neighbor 10.0.0.5 activate                     → activate EVPN for neighbor 10.0.0.5

sonic(config-router-af)#advertise-all-vni                              → advertise all VNI routing

sonic(config-router-af)# end

sonic# configure terminal

sonic(config)# vrf Vrf01                                               → Enter Vrf01 Setting

sonic(config-vrf)# vni 3000                                            → Announce map VNI 3000
on Vrf01

sonic(config-vrf)# end

sonic# write
```

SONiC02(VTYSH)

```
admin@7726:~$ vtysh                                                    → enter vtysh shell

Hello, this is FRRouting (version 7.2.1-sonic).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

sonic# configure terminal

sonic(config)#router bgp 65100                                         → assign BGP AS number

sonic(config-router)#neighbor 10.0.0.4 interface remote-as 65100       → assign Ethernet48 IP address to connect via
IBGP with peering on same AS

sonic(config-router)# address-family ipv4 unicast                      → Enter address-family ipv4

sonic(config-router-af)# network 2.2.2.2/32                            → Announce 2.2.2.2 network

sonic(config-router-af)# exit

sonic(config-router)#address-family l2vpn evpn                         → enter EVPN setting

sonic(config-router-af)#neighbor 10.0.0.4 activate                     → activate EVPN for neighbor 10.0.0.4

sonic(config-router-af)#advertise-all-vni                              → advertise all VNI routing

sonic(config-router-af)# end

sonic# configure terminal

sonic(config)# vrf Vrf01                                               → Enter Vrf01 Setting

sonic(config-vrf)# vni 3000                                            → Announce map VNI 3000
on Vrf01

sonic(config-vrf)# end

sonic# write
```

Step 8. Check EVPN-VNI status.

SONiC01(VTYSH)

```
sonic# show evpn vni detail
VNI: 1000
Type: L2
Tenant VRF: Vrf01
VxLAN interface: vtep-10
VxLAN ifIndex: 71
Local VTEP IP: 1.1.1.1
Mcast group: 0.0.0.0
No remote VTEPs known for this VNI
Number of MACs (local and remote) known for this VNI: 2
Number of ARPs (IPv4 and IPv6, local and remote) known for this VNI: 3
Advertise-gw-macip: No

VNI: 3000
Type: L3
Tenant VRF: Vrf01
Local Vtep Ip: 1.1.1.1
Vxlan-Intf: vtep-30
SVI-If: Vlan30
State: Up
VNI Filter: none
Router MAC: 04:f8:f8:6a:f6:91
L2 VNIs: 1000
```

SONiC02(VTYSH)

```
sonic# show evpn vni detail
VNI: 2000
Type: L2
Tenant VRF: Vrf01
VxLAN interface: vtep-20
VxLAN ifIndex: 71
Local VTEP IP: 2.2.2.2
Mcast group: 0.0.0.0
No remote VTEPs known for this VNI
Number of MACs (local and remote) known for this VNI: 2
Number of ARPs (IPv4 and IPv6, local and remote) known for this VNI: 3
Advertise-gw-macip: No

VNI: 3000
Type: L3
Tenant VRF: Vrf01
Local Vtep Ip: 2.2.2.2
Vxlan-Intf: vtep-30
SVI-If: Vlan30
State: Up
VNI Filter: none
Router MAC: 04:f8:f8:6b:06:91
L2 VNIs: 2000
```

Step 9. Check BGP EVPN status.

SONiC01(VTYSH)

```
sonic# show bgp summary

IPv4 Unicast Summary:
BGP router identifier 188.188.98.39, local AS number 65100 vrf-id 0
BGP table version 6
RIB entries 7, using 1288 bytes of memory
Peers 1, using 20 KiB of memory

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
Ethernet48 4 65100 12 15 0 0 0 00:01:29 3

Total number of neighbors 1

L2VPN EVPN Summary:
BGP router identifier 188.188.98.39, local AS number 65100 vrf-id 0
BGP table version 0
RIB entries 7, using 1288 bytes of memory
Peers 1, using 20 KiB of memory

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
Ethernet48 4 65100 12 15 0 0 0 00:01:29 4

Total number of neighbors 1
```

SONiC02(VTYSH)

```
sonic# show bgp summary

IPv4 Unicast Summary:
BGP router identifier 188.188.98.40, local AS number 65100 vrf-id 0
BGP table version 4
RIB entries 7, using 1288 bytes of memory
Peers 1, using 20 KiB of memory

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
Ethernet48 4 65100 12 12 0 0 0 00:01:29 3

Total number of neighbors 1

L2VPN EVPN Summary:
BGP router identifier 188.188.98.40, local AS number 65100 vrf-id 0
BGP table version 0
RIB entries 7, using 1288 bytes of memory
Peers 1, using 20 KiB of memory

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
Ethernet48 4 65100 12 12 0 0 0 00:01:29 4

Total number of neighbors 1
```

Step 10. Check routing and EVPN learning.

SONiC01(VTYSH)

```
sonic# show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP,
O - OSPF, I - IS-IS, B - BGP, E - EIGRP, N - NHRP,
T - Table, v - VNC, V - VNC-Direct, A - Babel, D - SHARP,
F - PBR, f - OpenFabric,
> - selected route, * - FIB route, q - queued route, r - rejected route

K>* 0.0.0.0/0 [0/202] via 188.188.1.1, eth0, 00:02:49
C>* 1.1.1.1/32 is directly connected, Loopback20, 00:02:34
B>* 2.2.2.2/32 [200/0] via 10.0.0.5, Ethernet48, 00:02:28
C>* 10.0.0.4/31 is directly connected, Ethernet48, 00:02:30
C>* 188.188.0.0/16 is directly connected, eth0, 00:02:49




sonic# show bgp l2vpn evpn
BGP table version is 4, local router ID is 188.188.98.39
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path
*> [2]:[0]:[48]:[8c:ea:1b:30:da:47]
1.1.1.1 32768 i
ET:8 RT:65100:1000 RT:65100:3000 Rmac:04:f8:f8:6a:f6:91
*> [2]:[0]:[48]:[8c:ea:1b:30:da:47]:[32]:[192.168.1.1]
1.1.1.1 32768 i
ET:8 RT:65100:1000 RT:65100:3000 Rmac:04:f8:f8:6a:f6:91
*> [2]:[0]:[48]:[8c:ea:1b:30:da:47]:[128]:[fe80::5037:63f4:3aba:bbe9]
1.1.1.1 32768 i
ET:8 RT:65100:1000
*> [3]:[0]:[32]:[1.1.1.1]
1.1.1.1 32768 i
ET:8 RT:65100:1000
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:4b]
2.2.2.2 100 0 i
RT:65100:2000 RT:65100:3000 ET:8 Rmac:04:f8:f8:6b:06:91
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:4b]:[32]:[192.168.2.2]
2.2.2.2 100 0 i
RT:65100:2000 RT:65100:3000 ET:8 Rmac:04:f8:f8:6b:06:91
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:4b]:[128]:[fe80::2a29:a798:37d3:76c2]
2.2.2.2 100 0 i
RT:65100:2000 ET:8
*>i[3]:[0]:[32]:[2.2.2.2]
2.2.2.2 100 0 i
RT:65100:2000 ET:8
Route Distinguisher: ip 192.168.1.254:2

*> [5]:[0]:[24]:[192.168.1.0]
1.1.1.1 0 32768 ?
ET:8 RT:65100:3000 Rmac:04:f8:f8:6a:f6:91
Route Distinguisher: ip 192.168.2.254:2

*>i[5]:[0]:[24]:[192.168.2.0]
2.2.2.2 0 100 0 ?
RT:65100:3000 ET:8 Rmac:04:f8:f8:6b:06:91

Displayed 10 out of 10 total prefixes
```

SONiC02(VTYSH)

```
sonic# show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP,
O - OSPF, I - IS-IS, B - BGP, E - EIGRP, N - NHRP,
T - Table, v - VNC, V - VNC-Direct, A - Babel, D - SHARP,
F - PBR, f - OpenFabric,
> - selected route, * - FIB route, q - queued route, r - rejected route

K>* 0.0.0.0/0 [0/202] via 188.188.1.1, eth0, 00:02:49
B>* 1.1.1.1/32 [200/0] via 10.0.0.4, Ethernet48, 00:02:28
C>* 2.2.2.2/32 is directly connected, Loopback20, 00:02:34
C>* 10.0.0.4/31 is directly connected, Ethernet48, 00:02:29
C>* 188.188.0.0/16 is directly connected, eth0, 00:02:49




sonic# show bgp l2vpn evpn
BGP table version is 4, local router ID is 188.188.98.40
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:47]
1.1.1.1 100 0 i
RT:65100:1000 RT:65100:3000 ET:8 Rmac:04:f8:f8:6a:f6:91
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:47]:[32]:[192.168.1.1]
1.1.1.1 100 0 i
RT:65100:1000 RT:65100:3000 ET:8 Rmac:04:f8:f8:6a:f6:91
*>i[2]:[0]:[48]:[8c:ea:1b:30:da:47]:[128]:[fe80::5037:63f4:3aba:bbe9]
1.1.1.1 100 0 i
RT:65100:1000 ET:8
*>i[3]:[0]:[32]:[1.1.1.1]
1.1.1.1 100 0 i
RT:65100:1000 ET:8
*> [2]:[0]:[48]:[8c:ea:1b:30:da:4b]
2.2.2.2 32768 i
ET:8 RT:65100:2000 RT:65100:3000 Rmac:04:f8:f8:6b:06:91
*> [2]:[0]:[48]:[8c:ea:1b:30:da:4b]:[32]:[192.168.2.2]
2.2.2.2 32768 i
ET:8 RT:65100:2000 RT:65100:3000 Rmac:04:f8:f8:6b:06:91
*> [2]:[0]:[48]:[8c:ea:1b:30:da:4b]:[128]:[fe80::2a29:a798:37d3:76c2]
2.2.2.2 32768 i
ET:8 RT:65100:2000
*> [3]:[0]:[32]:[2.2.2.2]
2.2.2.2 32768 i
ET:8 RT:65100:2000
Route Distinguisher: ip 192.168.1.254:2

*>i[5]:[0]:[24]:[192.168.1.0]
1.1.1.1 0 100 0 ?
RT:65100:3000 ET:8 Rmac:04:f8:f8:6a:f6:91
Route Distinguisher: ip 192.168.2.254:2

*> [5]:[0]:[24]:[192.168.2.0]
2.2.2.2 0 32768 ?
ET:8 RT:65100:3000 Rmac:04:f8:f8:6b:06:91

Displayed 10 out of 10 total prefixes
```