How to install Sonic on Celestica dx010 100g switch

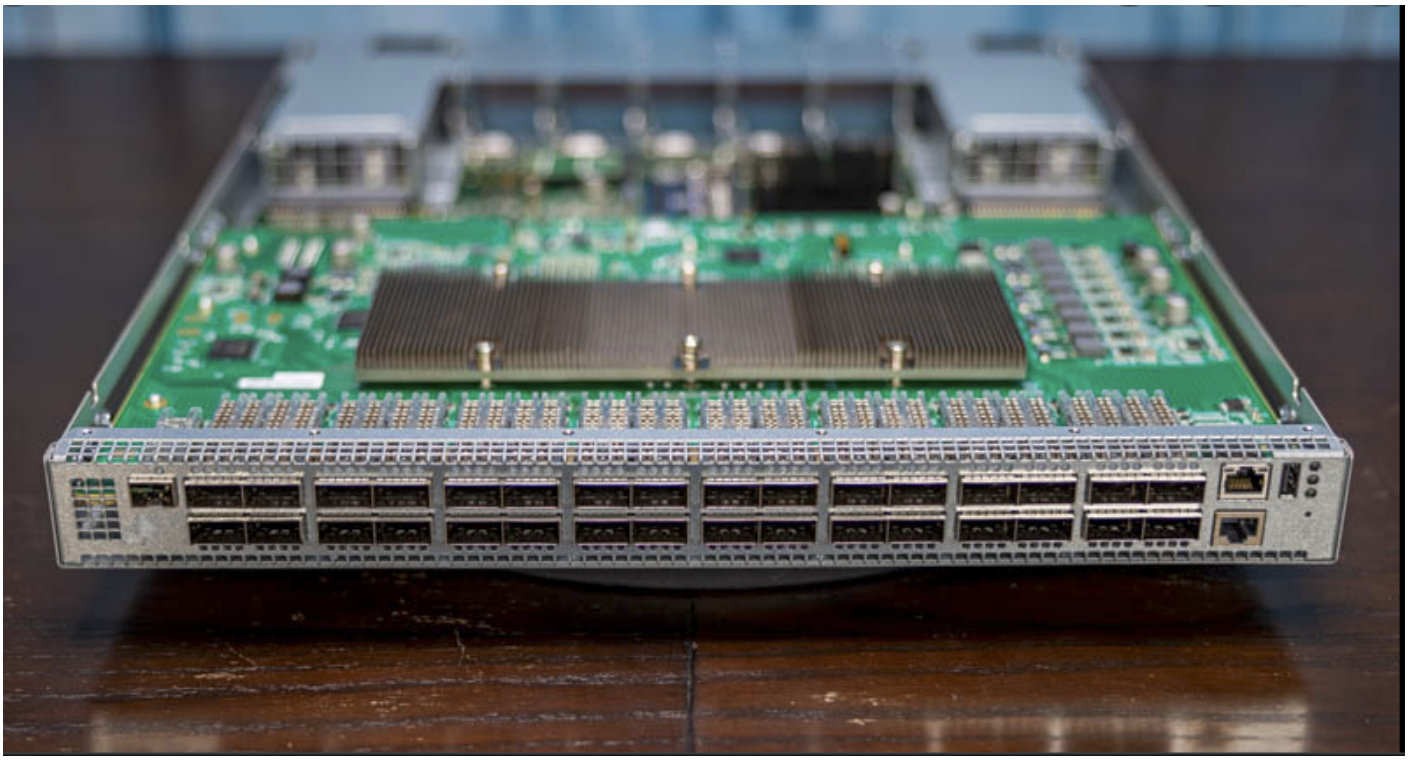

In this document i am going to talk about installation of Sonic Network operating system on celestica dx010 32x100G switches

If you do these steps you'll end up with a l2 switch that can forward 100g per port. Keep in mind that Spanning tree is still an issue on Sonic.

Important : Make sure you check the Sonic HCL, if your model is not there, don't try to do this installation.

Upgrading or re-installing onie

Onie is the Network operating system loader created by open network computing. It is a software to create zero touch provisioning on open compute / Bare metal switches. there are some binaries supplied by open compute platform, and you can compile or cross compile your version of it.

By your version i meant the model of the bare metal switch. before 2016 ONL which is sonic based on, supported powerpc cpu models however with sonic you are stuck with intel and some arm cpu's, no powerpc cpu is supported or can be compiled.

Compiling onie

To compile onie you need to read this documentation and have many headaches. I can't get in to details here because it will use a lot of space but it is a good documentation and pretty straight-forward to do.

Onie build documentation

Basically it will create a folder with many permutations of onie for every device out there.

if your device is a Powerpc platform then you'll need to do a cross-compile

installation or upgrading onie using u-boot

After powercycle press any key to stop u-boot process. take backup of current environment using

```
printenv
```

output should look like below, please copy and paste it somwhere safe.

Click for environment details

```
```=> printenv
active=image2
autoload=no
backup=image1
baudrate=115200
bootactive=fsload 1000000 $active; bootm 1000074
bootargs=root=dev/ram console=ttyS0,115200
bootbackup=fsload 1000000 $backup; bootm 1000074
bootcmd=run check_boot_reason; run nos_bootcmd; run onie_bootcmd
bootdelay=3
bootfile=uImage
bootimage1=fsload 1000000 image1; bootm 1000074
bootimage2=fsload 1000000 image2; bootm 1000074
check_boot_reason=if test -n $onie_boot_reason; then setenv onie_bootargs
boot_reason=$onie_boot_rea
son; run onie_bootcmd; fi;
consoledev=ttyS0
dboot=setenv bootargs root=/dev/ram mtdparts=physmap-
flash.0:126336k(jffs2),4096k(onie),128k(u-boot-
```
```

```
env),512k(u-boot) quiet console=ttyS0,$baudrate $dbootfrom && bootimg $dkaddr $draddr
$ddaddr
ddaddr=0x0D000000
dhcp_user-class=powerpc-quanta_ly2r-r0_uboot
dhcp_vendor-class-identifier=powerpc-quanta_ly2r-r0
diag_bootcmd=run dsdboot || run flash_bootcmd || run dnetboot
dkaddr=0x0A000000
dmemerase=mw $dkaddr 0 100; mw $draddr 0 100; mw $ddaddr 0 100; setenv dbootfrom
dnetboot=run dmemerase; run dnetcheck; if test -n $dbootfrom && test $dbootfrom = netboot;
then run
dboot; fi
dnetcheck=echo Checking diag from tftp server...; if tftp $dkaddr uImage && tftp $draddr
uInitrd2m &
& tftp $ddaddr ly2r.dtb; then setenv dbootfrom netboot; fi
draddr=0x0B000000
dsdboot=run dmemerase; run dsdcheck; if test -n $dbootfrom && test $dbootfrom = sdboot;
then run dbo
ot; fi
dsdcheck=echo Checking diag from Secure Digital Card...; if mmcinfo && ext2load mmc 0:1
$dkaddr uIma
ge && ext2load mmc 0:1 $draddr uInitrd2m && ext2load mmc 0:1 $ddaddr ly2r.dtb; then setenv
dbootfrom
 sdboot; fi
ethact=eTSEC1
ethaddr=a8:1e:84:58:67:10
flash_bootcmd=setenv bootargs root=/dev/ram mtdparts=physmap-
flash.0:58368k(jffs2),3968k(ramdisk),30
72k(kernel),128k(dtb),60800k(empty),4096k(onie),128k(u-boot-env),512k(u-boot)
console=ttyS0,$baudrat
e; bootm EBCE0000 EB900000 EBFE0000
hostname=LY2RP2020
hwconfig=usb1:dr_mode=host,phy_type=ulpi,esdhc
ipaddr=192.168.2.1
loadaddr=1000000
loads_echo=1
mtddevnum=0
mtdids=nor0=nor0
     [0/0]
mtdparts=nor0:65536k(jffs2),60800k(open),4096k(ONIE),128k(u-boot-env),512k(u-boot)
```

```
nfsboot=setenv bootargs root=/dev/nfs rw nfsroot=$serverip:$rootpath
ip=$ipaddr:$serverip:$gatewayip:$netmask:$hostname:$netdev:off
console=$consoledev,$baudrate $othbootargs;tftp $loadaddr $bootfile;tftp $fdtaddr
$fdtfile;bootm $loadaddr - $fdtaddr
nos_bootcmd=run bootactive; run bootbackup; run bootimage1; run bootimage2
onie_args=run onie_initargs onie_platformargs
onie_bootcmd=echo Loading Open Network Install Environment ...; echo Platform:
$onie_platform ; echo Version : $onie_version ; cp.b $onie_start $loadaddr ${onie_sz.b} &&
run onie_args && bootm ${loadaddr}
onie_dropbear_dss_host_key=begin-
base64@600@d#AAAAB3NzaC1kc3MAAACBAOCuKIyK2CzC6aUUNVfJLj6SlJ+anzWWUbuqtyn4#oty/BYzbG7rk+9ex
B/I9SOBMf0n8BHGJOdS+hXVFe8CAGf0fI7AlHF0zTZCN#kZAW70T175b4AZC8zCegABBKv86s0s7OT4qOXLCeYjCM9
gMWHRDAYHDfzktY#vBHapERPhIHzAAAAFQDfAxkC7+W/gxwO8BAg9Se08laRBwAAAIEAyjX1GFoX#oqPyEVn/qQ85x
Xpo3fr3UNT/NmZZGXQFZ+pUnHDOO6OwTZB0XLpPxzQ8WUVJ#Bt9d2tLC5Sw6rCp7odbp0Od9uWOc7ljt7HKtQVuodB
/yjt9KbQc69sOY25bC#efY2c0IcZ9LiNbQbH4COR6zGMLzA2XgHQ7uJ7xmV1fsAAACBAJP+ZxxFAehm#H9Ac5HUHYG
4ithFYK7l5MHRrA/pEGuqwUeezUKO9B8soCP9r4SoT/vLzUxQ5#TMwyANyi7BRjKnkjrRjJ/r6o9t5PS7YuGm/3G6h
OL7UHJoDIDFuznL9NG+f4#GhHF6NElGGnpmATry0+5gTN2KDYlZLoMriGaGBLVAAAAFQCDpjU05iU1Lbci#IvGrJfE
o2HQweg==#====#
onie_dropbear_rsa_host_key=begin-
base64@600@r#AAAAB3NzaC1yc2EAAAADAQABAAAAgQC8AVEAfeuMSVYNPJPULsBYL0KrFOOM#62CmnWNODmX230TZ
YW/DDv/NSOMUNFUR109MpzEyy8WmNqqNl+h0X3dK/KUy#H4iZhIxXXYaRu8Yywfaf/xXNnCcQqdj5JYmNGlaRggcvk
bYhnX2Hus7LrkGP#k0xdkl5erOsUGPmnf5RLNQAAAIAqqn3Qh8MiBi6iOGIwICAEo00U+FufyncX#uNYli7z2JrtAP
g2YDdpdozHtjubTcrKLadIPfhEv5G4WTigx6Mxwcw1oHUon#LoBzMCEjULMpyReO5xQcyV0WTLyPncPgqh4wtT7gr2
ePkEASwMgcIiXu0tV9#Q81k8UOcMZQz5kc3QQAAAEEA8mO+xK2LBWrmT37yh3c4xbg/Naj60CPxzRKR#HF0w6FzL8B
cCAL009z/YHtIJVpoXo321c8DNazZ47hQEoQpv0QAAAEEAxo/Q#1zXuEwsJUqrOE5RZ8KdfsQSEw9I+Pl8pDNBV06M
TbnQPgjMr9beD6NfSo1po#uZmbPVzaakzOtyduVCqCJQ==#====#
onie_initargs=setenv bootargs quiet console=$consoledev,$baudrate
onie_machine=quanta_ly2r
onie_machine_rev=0
onie_platform=powerpc-quanta_ly2r-r0
onie_platformargs=setenv bootargs $bootargs serial_num=${serial#} eth_addr=$ethaddr
$onie_bootargs $onie_debugargs
onie_rescue=setenv onie_boot_reason rescue && boot
onie_start=0xefb60000
onie_sz.b=0x00400000
onie_uninstall=setenv onie_boot_reason uninstall && boot
onie_update=setenv onie_boot_reason update && boot
onie_vendor_id=7244
onie_version=2015.05-v1.2-9d597eb
```

```
partition=nor0,0
platform=quanta_ly2r
ramboot=setenv bootargs root=/dev/ram rw console=$consoledev,$baudrate $othbootargs; tftp
$ramdiskaddr $ramdiskfile;tftp $loadaddr $bootfile;tftp $fdtaddr $fdtfile;bootm $loadaddr
$ramdiskaddr $fdtaddr
rootpath=/opt/nfsroot
serial#=QTFCEA7070054
serverip=192.168.2.11
stderr=serial
stdin=serial
stdout=serial
ver=U-Boot 2010.12 (Oct 14 2016 - 16:28:02) - ONIE 2015.05-v1.2-9d597eb

Environment size: 5059/8188 bytes```
```

We are going to change environment as shown and explanied below,

Set ip address to eth0/management port of bare metal switch

```
setenv ipaddr 4.3.2.2
setenv netmask 255.255.255.252
```

Specify tftp server address

```
setenv serverip 4.3.2.1
```

save changes

```
save
```

Make sure have the same subnet with tftp server After this point all the procedure will be destructive, please keep in mind that you might loose the current boot configuration or configurations within device.

specify space in flash

```
setenv start 0xefb60000
setenv sz.b 0x4a0000
```

upload file to ram

```
tftp onie-quanta_ly2r-r0.bin
```

Disable protection on flash and erase the old onie

```
protect off $start +${sz.b} && erase $start +${sz.b}
```

start flashing

```
cp.b $fileaddr $start ${sz.b} && protect on $start +${sz.b}
```

after this point Onie will be upgraded/installed and procedure will reset the device booting into new onie and it's environment

Installing sonic

Before install

Sonic is a buggy software, there are many missing features, and because of that there are many builds for every vendor the list is here

we are going to use 202111 edition on celestica dx010

installation

Onie will boot up and you will be greeted with onie-discovery process, unless you did special dhcp configurations, we will need to stop onie-discovery

```
onie-discover-stop
```

then the ip addressing

```
ip addr add 4.3.2.2/30 dev eth0
ip link set up dev eth0
```

start installing sonic

```
onie-nos-install tftp://4.3.2.1/sonic-buildimage.broadcom
```

at this point onie installer will do its thing after 5-10 minutes you will be greeted with sonics login screen

After install

Default user name: admin default pass : YourPaSsWoRd By default sonic names every 100g interface with steps of 4, Ethernet0,Ethernet4,Ethernet8 and there are aliases for you to see which port are you working on

```
admin@sonic:~$ show interfaces status
Interface            Lanes    Speed    MTU    FEC    Alias    Vlan    Oper
Admin            Type    Asym PFC
 -----------  --------------  -------  -----  -----  -------  ------  ------  -------
--------------  ----------
Ethernet0              65      10G   9100    N/A   Eth1/1  routed    down     down  QSFP+ or
later         N/A
Ethernet1              66      10G   9100    N/A   Eth1/2  routed    down     down  QSFP+ or
later         N/A
Ethernet2              67      10G   9100    N/A   Eth1/3  routed    down     down  QSFP+ or
later         N/A
Ethernet3              68      10G   9100    N/A   Eth1/4  routed    down     down  QSFP+ or
later         N/A
Ethernet4              69      10G   9100    N/A   Eth2/1  routed    down     down  QSFP+ or
later         N/A
Ethernet5              70      10G   9100    N/A   Eth2/2  routed    down     down  QSFP+ or
later         N/A
Ethernet6              71      10G   9100    N/A   Eth2/3  routed    down     down  QSFP+ or
later         N/A
Ethernet7              72      10G   9100    N/A   Eth2/4  routed    down     down  QSFP+ or
later         N/A
Ethernet8      73,74,75,76     100G   9100    N/A     Eth3  routed    down
up             N/A         N/A
 Ethernet12      77,78,79,80     100G   9100    N/A     Eth4  routed    down
up             N/A         N/A
 Ethernet16      33,34,35,36     100G   9100    N/A     Eth5  routed    down
up             N/A         N/A
 Ethernet20      37,38,39,40     100G   9100    N/A     Eth6  routed    down
up             N/A         N/A
 Ethernet24      41,42,43,44     100G   9100    N/A     Eth7  routed    down
up             N/A         N/A
 Ethernet28      45,46,47,48     100G   9100    N/A     Eth8  routed    down
up             N/A         N/A
 Ethernet32      49,50,51,52     100G   9100    N/A     Eth9  routed    down
up             N/A         N/A
 Ethernet36      53,54,55,56     100G   9100    N/A    Eth10  routed    down
```

```
up             N/A         N/A
 Ethernet40      57,58,59,60     100G   9100    N/A    Eth11  routed    down
up             N/A         N/A
 Ethernet44      61,62,63,64     100G   9100    N/A    Eth12  routed    down
up             N/A         N/A
 Ethernet48      81,82,83,84     100G   9100    N/A    Eth13  routed    down
up             N/A         N/A
 Ethernet52      85,86,87,88     100G   9100    N/A    Eth14  routed    down
up             N/A         N/A
 Ethernet56      89,90,91,92     100G   9100    N/A    Eth15  routed    down
up             N/A         N/A
 Ethernet60      93,94,95,96     100G   9100    N/A    Eth16  routed    down
up             N/A         N/A
 Ethernet64     97,98,99,100     100G   9100    N/A    Eth17  routed    down
up             N/A         N/A
 Ethernet68  101,102,103,104     100G   9100    N/A    Eth18  routed    down
up             N/A         N/A
 Ethernet72  105,106,107,108     100G   9100    N/A    Eth19  routed    down
up             N/A         N/A
 Ethernet76  109,110,111,112     100G   9100    N/A    Eth20  routed    down
up             N/A         N/A
 Ethernet80          1,2,3,4     100G   9100    N/A    Eth21  routed    down
up             N/A         N/A
 Ethernet84          5,6,7,8     100G   9100    N/A    Eth22  routed    down
up             N/A         N/A
 Ethernet88       9,10,11,12     100G   9100    N/A    Eth23  routed    down
up             N/A         N/A
 Ethernet92      13,14,15,16     100G   9100    N/A    Eth24  routed    down
up             N/A         N/A
 Ethernet96              17      25G   9100    N/A  Eth25/1  routed    down     down  QSFP+
or later         N/A
 Ethernet97              18      25G   9100    N/A  Eth25/2  routed    down     down  QSFP+
or later         N/A
 Ethernet98              19      25G   9100    N/A  Eth25/3  routed    down     down  QSFP+
or later         N/A
 Ethernet99              20      25G   9100    N/A  Eth25/4  routed    down     down  QSFP+
or later         N/A
Ethernet100      21,22,23,24     100G   9100    N/A    Eth26  routed    down
up             N/A         N/A
Ethernet104      25,26,27,28     100G   9100    N/A    Eth27  routed    down
```

```
up             N/A         N/A
Ethernet108      29,30,31,32      40G   9100    N/A    Eth28  routed    down       up  QSFP+
or later         N/A
Ethernet112  113,114,115,116     100G   9100    N/A    Eth29  routed    down
up             N/A         N/A
Ethernet116  117,118,119,120      40G   9100    N/A    Eth30  routed    down       up  QSFP+
or later         N/A
Ethernet120  121,122,123,124     100G   9100    N/A    Eth31  routed    down
up             N/A         N/A
Ethernet124  125,126,127,128     100G   9100    N/A    Eth32  routed    down
up             N/A         N/A
```

Sonic has some problems, it's not the news, it's a fact. One of them is this installation media doesn't create a running environment, to do that follow the steps below

```
sudo sonic-cfggen --preset l2 -p -H -k Seastone-DX010 -w
sudo config reload -y
sudo reboot
```

Clear the IP addresses from each interface

```
show runningconfiguration interfaces | grep \| | awk -F'"' '{ print $2 }' | awk -F'|' '{ print
"sudo config interface ip remove "$1" "$2 }' > /var/tmp/remove-l3-ips
bash /var/tmp/remove-l3-ips
rm -f /var/tmp/remove-l3-ips
```

Configure sonic

Layer1

Every port and platform or asic has different abilities or features, for example you can breakout every port on a celestica dx010 how ever an edge-core AS9716-32D will only do breakouts on Ethernet0-20 which is the top row of switch. I was able to work with 3rdparty qsfp modules on sonic however i couldn't find an HCL about qsfps .

Breakouts

You neeed to do everything written on *after install* to be able to do breakouts. Supported breakouts are ;

- 1x100G[100g]
- 1x100G[40g]
- 2x50G[50G]
- 4x25G[25G]
- 4x25G[10G]

to break out 1 port to 4 10g ports

```
sudo config interface breakout Ethernet96 4x25G[10G] -y
```

a normal output should look like below, after this you will need to reboot the switch.

```
admin@sonic:~$ sudo config interface breakout Ethernet96 4x25G[10G] -y

Running Breakout Mode : 1x100G[40G]
Target Breakout Mode : 4x25G[10G]

Ports to be deleted :
 {
    "Ethernet96": "100000"
}
Ports to be added :
 {
    "Ethernet96": "25000",
    "Ethernet97": "25000",
    "Ethernet98": "25000",
    "Ethernet99": "25000"
}

After running Logic to limit the impact

Final list of ports to be deleted :
 {
    "Ethernet96": "100000"
}
Final list of ports to be added :
 {
    "Ethernet96": "25000",
    "Ethernet97": "25000",
    "Ethernet98": "25000",
    "Ethernet99": "25000"
```

```
}
sonic_yang(6):Note: Below table(s) have no YANG models: SNMP, SNMP_COMMUNITY
Below Config can not be verified, It may cause harm to the system
 {}
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`.
```

Dont forget to save :)

```
admin@sonic:~$ sudo config save
Existing files will be overwritten, continue? [y/N]: y
Running command: /usr/local/bin/sonic-cfggen -d --print-data > /etc/sonic/config_db.json
```

You **will reboot** the switch so it can show newly emerged ports.

After reboot you will have Ethernet 96 to 99, set the speeed to 10g on those interfaces

```
config interface speed Ethernet96 10000
```

And start the port,

```
config interface startup Ethernet96
```

Do this for every port you broke and want to work immediately

Layer2

Creating Vlans

here

```
sudo config vlan add 1000 
```

vlans and membership

to see the details

```
show vlan brief
```

tagging, untagged ports

- tagged

```
sudo config vlan member add 1000 Ethernet48
```

- untagged

```
sudo config vlan member add -u 1000 Ethernet48
```

You can use the following commands to edit ports in bulk

```
seq 36 39|xargs -i -t sudo config interface speed Ethernet{} 10000
seq 36 39|xargs -i -t sudo config interface startup Ethernet{}
seq 36 39|xargs -i -t sudo config vlan member add -u 1000 Ethernet{}
```

Layer3

intervlan routing

to add an ip or svi to vlan 1 use the command below

```
admin@sonic:~$ sudo config interface ip add Vlan1000 192.168.1.2/24
```

routed interfaces

as in svi configuration but this time use the port name

```
sudo config interface ip add Ethernet4 192.168.1.2/24
```

Revision #3
Created 2022-11-09 11:52:57 UTC by Mesut Bayrak
Updated 2022-12-05 05:37:00 UTC by Mesut Bayrak