



# Clocking Mode Verification for Gen5 Test System



PCIe<sup>®</sup> Gen5 Update, SBExpress-RM5

Verification for SSD Drives

## Table of Contents

Table of Contents .....	1
1. PCIe Clocking Modes in SANBlaze PCIe Gen5 System.....	3
2. Different PCIe Clocking Modes Verification .....	4
2.1 Link Error Counters and PCIe AER Flags.....	5
2.1.1 Link Error Counters in PEX89144 .....	5
2.1.2 PCIe AER Flags .....	5
2.2 Reference Clock Monitoring .....	6
2.3 Skip Ordered-Sets in Bus Trace Verification .....	7
3. Common Clocking Mode Without SSC .....	7
3.1 Configure Clocking Mode as Common Clock Without SSC.....	8
3.2 Check Link Error Counters.....	11
3.3 Clock Scope Shot from PCIe Switch Side .....	12
3.4 Clock Scope Shot from Drive Side .....	13
3.5 Skip Ordered-Sets from Bus Trace .....	14
4. Common Clocking Mode With SSC .....	15
4.1 Configure Clocking Mode as Common Clock With SSC .....	16
4.2 Check Link Error Counters.....	18
4.3 Clock Scope Shot from PCIe Switch Side .....	19
4.4 Clock Scope Shot from Drive Side .....	20
4.5 Skip Ordered-Sets from Bus Trace .....	21
5. Separate Reference Clock Without SSC (Static SRNS).....	22
5.1 Configure Clocking Mode as Separate Clock Without SSC .....	23
5.2 Check Link Error Counters.....	25
5.3 Clock Scope Shot from PCIe Switch Side .....	26
5.4 Clock Scope Shot from Drive Side .....	27
5.5 Skip Ordered-Sets from Bus Trace .....	28
6. Separate Reference Clock Independent SSC (Static SRIS) .....	29
6.1 Configure Clocking Mode as Separate Clock With SSC .....	29
6.2 Check Link Error Counters.....	32
6.3 Clock Scope Shot from PCIe Switch Side .....	33

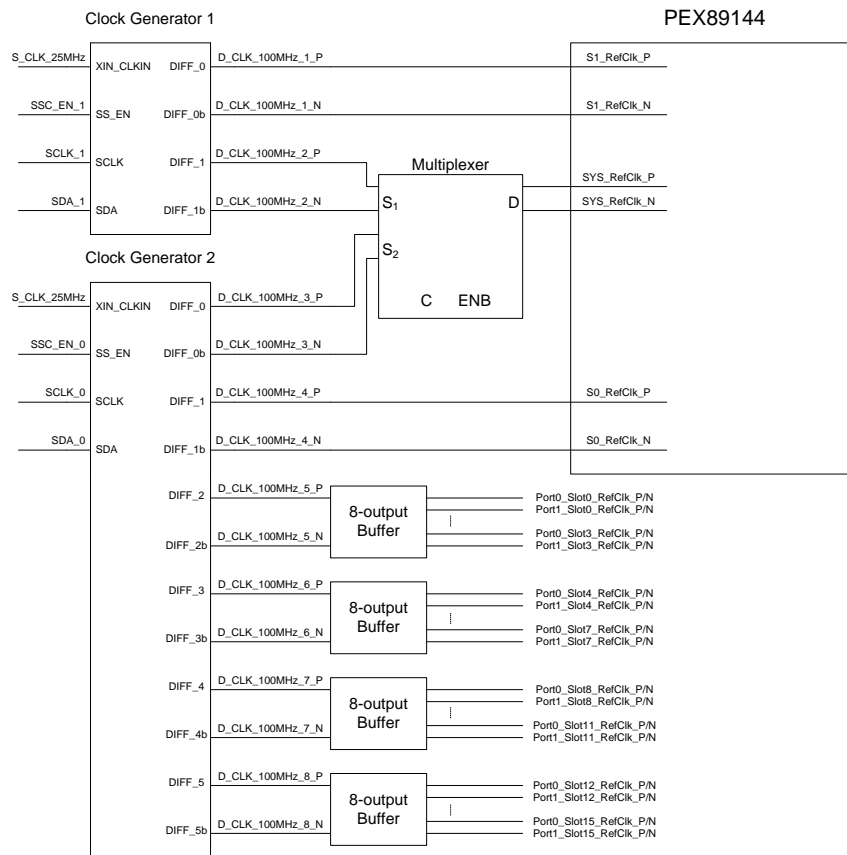
6.4 Clock Scope Shot from Drive Side .....	34
6.5 Skip Ordered-Sets from Bus Trace .....	35
7. RefClk Without SSC for PEX Only (No RefClk to Target Drive, Dynamic SRNS) .....	36
7.1 Configure Clocking Mode as Dynamic SRNS .....	37
7.2 Check Link Error Counters.....	39
7.3 Clock Scope Shot from PCIe Switch Side .....	40
7.4 Clock Scope Shot from Drive Side .....	41
7.5 Skip Ordered-Sets from Bus Trace .....	42
8. RefClk with SSC for PEX Only (No RefClk to Target Drive, Dynamic SRIS).....	43
8.1 Configure Clocking Mode as Dynamic SRIS.....	44
8.2 Check Link Error Counters.....	46
8.3 Clock Scope Shot from PCIe Switch Side .....	47
8.4 Clock Scope Shot from Drive Side .....	48
8.5 Skip Ordered-Sets from Bus Trace .....	49
9. Common Clock Without SSC vs. Common Clock With SSC.....	50
9.1 Drive Side .....	50
9.2 Switch Side .....	52
10. Static SRNS vs. Static SRIS.....	54
10.1 Drive Side .....	54
10.2 Switch Side .....	56
11. Dynamic SRNS vs. Dynamic SRIS .....	58
11.1 Drive Side .....	58
11.2 Switch Side .....	60
12. Summary .....	62

# 1. PCIe Clocking Modes in SANBlaze PCIe Gen5 System

The SANBlaze SBExpress PCIe Gen5 Test System is using the Broadcom® PEX89144 PCI Express (PCIe) 5.0 switch, which supports the following PCIe clocking modes:

- Separate reference clock independent SSC (SRIS)
  - Dynamic SRIS: no clock to the drive
  - Static SRIS: provide clock to the drive
- Separate reference clock without SSC (SRNS)
  - Dynamic SRNS: no clock to the drive
  - Static SRNS: provide clock to the drive
- Spread spectrum clocking (SSC) isolation
- Common clocking mode (without SSC or with SSC)

In order to meet customer’s requirement of PCIe clocking modes verification, we implemented the PCIe clocking in our PCIe Gen5 system (board with PN= 600-950003) as follows:



The clocking mode image, RefClk source and SSC can be selected with CLI commands running in VLUN as shown in the following table 1.

Table 1: CLI Commands for Clocking Mode Selection

Clock Mode	CLI Command
Common Clock Without SSC	iRiser -d -3 set SW1EESEL0 iRiser -d -3 set SW2EESEL1 iRiser -d -3 clear SW6SRIS iRiser -d -3 set SW7SSCL
Common Clock With SSC	iRiser -d -3 clear SW1EESEL0 iRiser -d -3 set SW2EESEL1 iRiser -d -3 clear SW6SRIS iRiser -d -3 clear SW7SSCL
Separate Clock Without SSC (SRNS)	iRiser -d -3 clear SW1EESEL0 iRiser -d -3 set SW2EESEL1 iRiser -d -3 set SW6SRIS iRiser -d -3 set SW7SSCL
Separate Clock With SSC (SRIS)	iRiser -d -3 set SW1EESEL0 iRiser -d -3 clear SW2EESEL1 iRiser -d -3 set SW6SRIS iRiser -d -3 clear SW7SSCL

The clocking mode image, RefClk source and SSC can also be controlled by the back panel switch 1, 2, 6 and 7 of the system as follows, but not recommended. We prefer users to use the CLI commands above.



## 2. Different PCIe Clocking Modes Verification

The SANBlaze PCIe Gen5 system allows us to verify all different clocking modes below for customers' drives:

- Common clocking mode without SSC
- Common clocking mode with SSC
- Separate reference clock without SSC (Static SRNS)
- Separate reference clock independent SSC (Static SRIS)
- RefClk without SSC for PEX only (no RefClk to the target drive, Dynamic SRNS)
- RefClk with SSC for PEX only (no RefClk with SSC to the target drive, Dynamic SRIS)

## 2.1 Link Error Counters and PCIe AER Flags

Link error counters and PCIe AER flags will show if the drive is working well or not in the specified clocking mode. SANBlaze provided one CLI command to dump the Receiver Error Counter, TLP Error Counter, DLLP Error Counter and Recovery Counter in PEX89144. We also can monitor the AER (Advanced Error Reporting) flags like UESa (Uncorrectable Error Status) and CESa (Correctable Error Status) in PCIe registers.

### 2.1.1 Link Error Counters in PEX89144

The Receiver Error Counter is register 0x0BF4 in PEX89144. Please notice it is a station register. When we read and write/clear this register we need to pick up the correct port, also set bit 31 when write/clear. This register returns the receiver error count for the port specified in the Port Select field of this register. It uses 8 bits for receiver error counts and the count saturates at 255.

The TLP error count register is 0x0FAC in PEX89144. It is using full 32-bit and counts the quantity of TLPs received with bad LCRC, or quantity of TLPs with a Sequence Number Mismatch error. The Counter saturates at FFFFFFFFh and does not roll over to 00000000h.

The DLLP error count register is 0x0FB0 in PEX89144. It is using full 32-bit and counts the quantity of DLLPs received with bad LCRC, or quantity of DLLPs with a Sequence Number Mismatch error. The Counter saturates at FFFFFFFFh and does not roll over to 00000000h.

The Recover Diagnostic Register has address as 0x0BC4 in PEX89144 and access type RW, RO and VAR. It is a station level register and using 16-bit to count the quantity of recovery events. It returns the number of times the selected port entered the Recovery state after first reaching the L0 state or after the count was last cleared.

### 2.1.2 PCIe AER Flags

The PCIe AER Uncorrectable Error Status (UESa) register indicates error detection status of individual errors. Errors that are defined as non-Function-specific are logged in the PF. Only Function-specific errors are logged in the VFs. It is a 32-bit flag register and all bits are defined as follows:

*Table 9-33 Uncorrectable Error Status Register Changes*

Bit Location	PF and VF Register Differences From Base	PF Attributes	VF Attributes
4	<u>Data Link Protocol Error Status</u>	Base	0b
5	<u>Surprise Down Error Status</u>	Base	0b
13	<u>Flow Control Protocol Error Status</u>	Base	0b
17	<u>Receiver Overflow Status</u>	Base	0b
18	<u>Malformed TLP Status</u>	Base	0b
19	<u>ECRC Error Status</u>	Base	0b

The PCIe AER Correctable Error Status (CESta) register indicates error detection status of individual correctable errors. Errors that are defined as non-Function-specific are logged in the PF. Only Function-specific errors are logged in the VFs. It is also a 32-bit flag register and all bits are defined as follows:

*Table 9-36 Correctable Error Status Register Changes*


Bit Location	PF and VF Register Differences From Base	PF Attributes	VF Attributes
0	<u>Receiver Error Status</u>	Base	0b
6	<u>Bad TLP Status</u>	Base	0b
7	<u>Bad DLLP Status</u>	Base	0b
8	<u>REPLAY_NUM Rollover Status</u>	Base	0b
12	<u>Replay Timer Timeout Status</u>	Base	0b
15	<u>Header Log Overflow Status</u> If the VF implements Header Log sharing (see <a href="#">Section 9.4.2.1</a> ), this bit is hardwired to 0b.	Base	Base / 0b

The UESTa and CESta from both parent (root complex downstream port) and child/target (SSD) can be dumped with command “setpci” based on their BDF address. The link errors (Receiver Error, Bad TLP and Bad DLLP) show up in PEX89144 (root complex downstream port) should match the parent AER CESta flag instead the target AER CESta flag. During measurement we need to clear all PEX89144 link error counts at the beginning of the test then let them accumulate if happened during the test. But AER UESTa and CESta registers needs to clear at the beginning of each sample loop (for example sample every 1 minute per loop) because they are error flags and have to clear so they can flag again when got error.


## 2.2 Reference Clock Monitoring

Monitoring the reference clock in both PCIe switch side and drive side is the most straightforward way to verify the clocking mode is working as expected or not.

We wired out the reference clock to the PCIe switch with test points as follows for monitoring:



We also wired out the reference clock to the drive slot with test points as follows for monitoring:



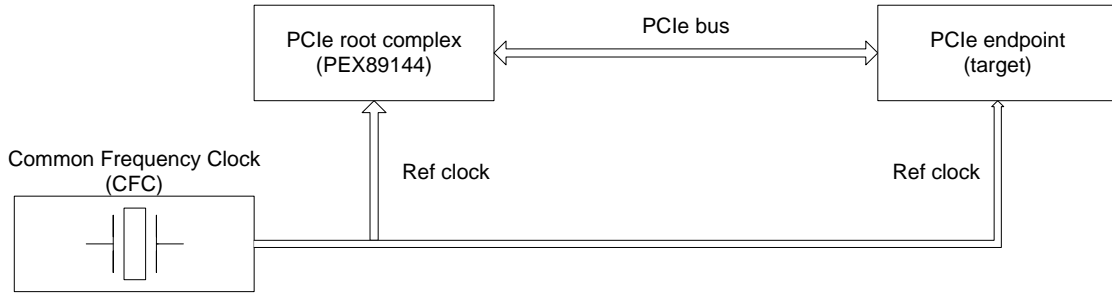
### 2.3 Skip Ordered-Sets in Bus Trace Verification

When drive didn't detect any external Reference Clock and it will use its own internal Reference Clock with SSC on, which will send extra skip ordered-sets. Monitoring the skip ordered-sets in bus trace can also be used for specified clocking mode verification.

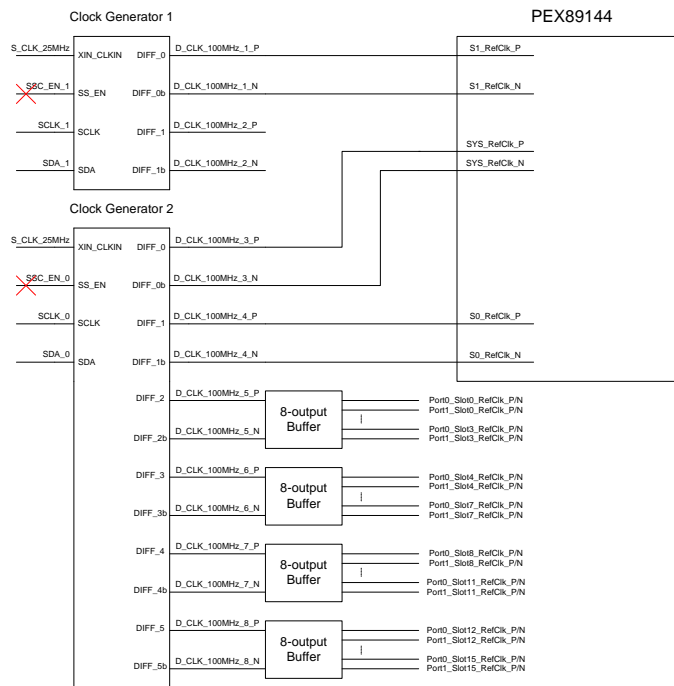
## 3. Common Clocking Mode Without SSC

The common clocking mode without SSC can be described as following chart:





This configuration is the default PCIe clocking mode in the SANBlaze Gen5 system as follows:



The SSC\_EN\_0 and SSC\_EN\_1 for Clock Generator 1 and 2 are disabled by the Back Panel Switches (SW6 is up and SW7 is down, which is the default setting as well). The S0\_RefClk\_P/N and S1\_RefClk\_P/N are ignored in this case because they are for SSC isolation only. The EEPROM image on the SANBlaze SBExpress Gen5 system specified all ports SerDes are clocked with SYS\_RefClk\_P/N. Both PEX89144 (with the MUX selection) and targets got common RefClk without SSC from Clock Generator 2.

With this default clocking mode, I sample the link error counter registers every 1 minute and I didn't notice any receiver errors, TLP errors, DLLP errors and Recovery events happened.

### 3.1 Configure Clocking Mode as Common Clock Without SSC

The default clocking mode of SANBlaze Gen5 system is Common Clock Without SSC. We can configure the system to use this common clock without SSC as follows:

```
[root@vln-100-158-IPMI-151 Linux]# iriser -d -3 set SW1ESEL0
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
WARNING: Systemtype for parent slot ID ffffffff not known. Is this an iRiser5P in a non-SANBlaze system?
INFO: psnToPCIe unable to determine systemtype. Use USB interface for slot=34, assume iRiser4
INFO: id=0 for Motherboard assume RM5, ID=0x2005
```



0	0	Down	128	128	2048	Gen4	/	Gen5	x4	/	x4	Common	88	Up
0	4	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
0	6	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
0	8	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
0	10	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
0	12	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
0	14	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
1	16	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
1	18	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
1	20	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
1	22	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
1	24	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
1	26	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
1	28	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
1	30	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	88	Down
2	32	Down	128	128	2048	-	/	Gen5	-	/	x16	Common	9A	Down
3	48	Host	128	128	2048	Gen4	/	Gen5	x16	/	x16	Common	85	Up
4	64	Down	128	128	2048	-	/	Gen5	-	/	x16	Common	9D	Down
5	80	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	9D	Down
5	82	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	9D	Down
5	84	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	9D	Down
5	86	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	9D	Down
5	88	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	9D	Down
5	90	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	9D	Down
5	92	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	9D	Down
5	94	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	9D	Down
6	96	Down	128	128	2048	-	/	Gen5	-	/	x8	Common	A8	Down
6	104	Down	128	128	2048	Gen2	/	Gen5	x1	/	x8	Common	A8	Up
7	112	Down	128	128	2048	-	/	Gen5	-	/	x16	Common	A8	Down
8	128	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	AD	Down
8	130	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	AD	Down
8	132	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	AD	Down
8	134	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	AD	Down
8	136	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	AD	Down
8	138	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	AD	Down
8	140	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	AD	Down
8	142	Down	128	128	2048	-	/	Gen5	-	/	x2	Common	AD	Down

Host Ports:

BusPrim: Primary Bus number in hex, BusSec: Secondary Bus number in hex  
 BusSub: Subordinate Bus number in hex, Assigned DS Ports: In decimal

Port	BusPrim	BusSec	BusSub	Assigned DS Ports
48	85	86	B5	0, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 32, 64, 80, 82, 84, 86, 88, 90, 92, 94, 96, 104, 112, 128, 130, 132, 134, 136, 138, 140, 142,

Downstream Ports:

Port: Port number in decimal, DS buses: Downstream bus range  
 PhysSlot: Chassis Physical slot in decimal, HostPort: Associated Host Port in decimal  
 VID: Vendor ID, DID: Device ID  
 SVD: Sub-vendor ID, SID: Sub-device ID  
 Attached Device Details: All values are in hex

Port	DS buses	PhysSlot	HostPort	Attached Device Details				ClassCode	SubClassCode
				VID	DID	SVD	SID		
0	89->89	0	48	-	-	-	-	-	-
4	8B->8B	2	48	-	-	-	-	-	-
6	8C->8C	3	48	-	-	-	-	-	-
8	8D->8D	4	48	-	-	-	-	-	-
10	8E->8E	5	48	-	-	-	-	-	-
12	8F->8F	6	48	-	-	-	-	-	-
14	90->90	7	48	-	-	-	-	-	-
16	91->91	8	48	-	-	-	-	-	-
18	92->92	9	48	-	-	-	-	-	-
20	93->93	10	48	-	-	-	-	-	-
22	94->94	11	48	-	-	-	-	-	-
24	95->95	12	48	-	-	-	-	-	-
26	96->96	13	48	-	-	-	-	-	-
28	97->97	14	48	-	-	-	-	-	-
30	98->98	15	48	-	-	-	-	-	-
32	9B->9B	16	48	-	-	-	-	-	-
64	9E->9E	32	48	-	-	-	-	-	-
80	9F->9F	40	48	-	-	-	-	-	-
82	A0->A0	41	48	-	-	-	-	-	-
84	A1->A1	42	48	-	-	-	-	-	-
86	A2->A2	43	48	-	-	-	-	-	-
88	A3->A3	44	48	-	-	-	-	-	-
90	A4->A4	45	48	-	-	-	-	-	-
92	A5->A5	46	48	-	-	-	-	-	-
94	A6->A6	47	48	-	-	-	-	-	-
96	A9->A9	48	48	-	-	-	-	-	-
104	AA->AA	52	48	-	-	-	-	-	-
112	AB->AB	56	48	-	-	-	-	-	-
128	AE->AE	64	48	-	-	-	-	-	-
130	AF->AF	65	48	-	-	-	-	-	-
132	B0->B0	66	48	-	-	-	-	-	-
134	B1->B1	67	48	-	-	-	-	-	-
136	B2->B2	68	48	-	-	-	-	-	-

```

138 B3->B3 69 48 - - - - -
140 B4->B4 70 48 - - - - -
142 B5->B5 71 48 - - - - -
PEX89144 A0> quit

```

Check the back panel switch 1, 2, 6 and 7 of the system again as follows to ensure they are correct:

```

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked

```

```

GPIO|31          24|23          16|15          8|7          0
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
B P D D T T D S S S S S S S S D D D D D D D D D D D D D D
i w i i P P P r r W W W W W W W W u u u u u u u u u u u u u
t r s s s 3 2 1 v 8 7 6 5 4 3 2 1 a a a a a a a a a a a a a a a
3 O U W | | | R D S S S F C U E E 1 1 1 1 1 1 1 1 1 1 1 1 1
1 n s t | | | s i S r r h s E E 1 1 1 1 1 1 1 0 0 0 0 0 0 0
N R b c | | | t s C I c a l S S 5 4 3 2 1 0 9 8 7 6 5 4 3 2 1 0
/ s L h | | | L R L S P i l e l l 1 1 1 1 1 1 1 1 1 1 1 1 1
C t D g | | | | | | | | | | | | | | | | | | | | | | | | | | |
o o o o o o o o o o o o o o o o o o o o o o o o o o o o o
1 1 1 1 0 1 1 1 1 1 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
Dir:[0x214]0xffffffff
Val:[0x210]0xefdffff

```

### 3.2 Check Link Error Counters

We can clear the PCIe link error counters using the command line as follows:

```

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5 -r # clear all PCIe link error counters
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4 (RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1h detect.Active

```

After clearing the PCIe link error counters now we can dump them as follows when some I/O running in the background and dump every other 60 seconds.

```

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4 (RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1h detect.Active

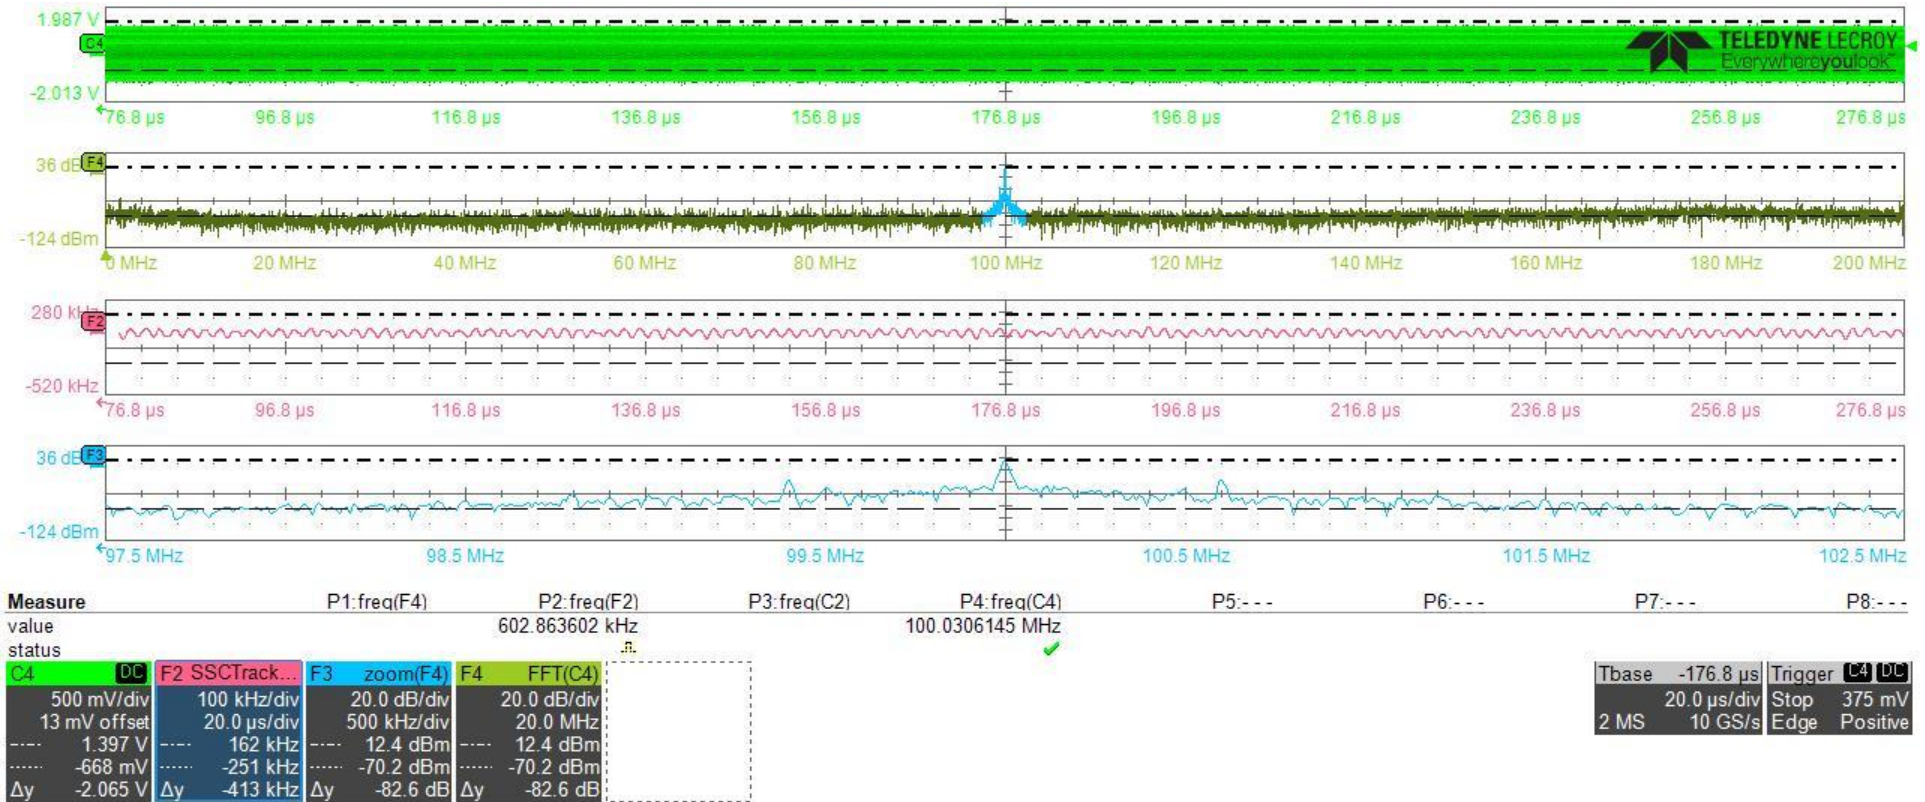
[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4 (RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4 (RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4 (RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1h detect.Active

```

### 3.3 Clock Scope Shot from PCIe Switch Side



- The channel “C4” is the reference clock obtained from PCIe switch side, and the measurement shows it is around 100 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 100 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 100 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts and displays the instantaneous frequency variation of the clock as a function of time.
- From the scope shot above we can tell the clock in PCIe switch side has no SSC turned on.

### 3.4 Clock Scope Shot from Drive Side



- The channel “C4” is the reference clock obtained from PCIe switch side, and the measurement shows it is around 100 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 100 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 100 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts and displays the instantaneous frequency variation of the clock as a function of time.
- From the scope shot above we can tell the clock on the drive side does not have SSC turned on.


### 3.5 Skip Ordered-Sets from Bus Trace

Trace View / PCIe Revision 4.0																
Packet	Direction	Length	Symbol	SKIP Symbols	END	P	FR	SR	CRC	MP	UM	MT	RN	Payload	Time Delta	Time Stamp
0	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	649.000 ns	-0000.000000649000 s
1	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	251.870 ns	0000.000000000000 s
2	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	*	*	*	*						260.870 ns	0000.000000260000 s
3	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	251.870 ns	0000.000000529000 s
4	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	*	*	*	*						260.870 ns	0000.000000789000 s
5	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	249.870 ns	0000.000001058000 s
6	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	*	*	*	*						261.870 ns	0000.000001316000 s
7	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	251.870 ns	0000.000001586000 s
8	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	*	*	*	*						257.870 ns	0000.000001846000 s
9	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	252.000 ns	0000.000002112000 s
10	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	*	*	*	*						8.000 ns	0000.000002364000 s
11	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	1	*	*	*						260.870 ns	0000.000002372000 s
12	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	251.870 ns	0000.000002641000 s
13	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	1	*	*	*						260.870 ns	0000.000002901000 s
14	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	251.870 ns	0000.000003170000 s
15	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	*	*	*	*						259.870 ns	0000.000003430000 s
16	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	249.870 ns	0000.000003698000 s
17	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	*	*	*	*						261.870 ns	0000.000003956000 s
18	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	249.870 ns	0000.000004226000 s
19	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	*	*	*	*						260.870 ns	0000.000004484000 s
20	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	251.870 ns	0000.000004753000 s
21	R→	16.0 x4	SKIP	AAAAAAAAA ...	E1	*	*	*	*						260.870 ns	0000.000005013000 s
22	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	*	*	*	25	0	0	7	0	156	98.000 ns	0000.000005282000 s
23	R→	16.0 x4	SKIP	AAAAAAAAA ...	78	0	0	0	25	0	0	7	0	156	162.000 ns	0000.000005380000 s

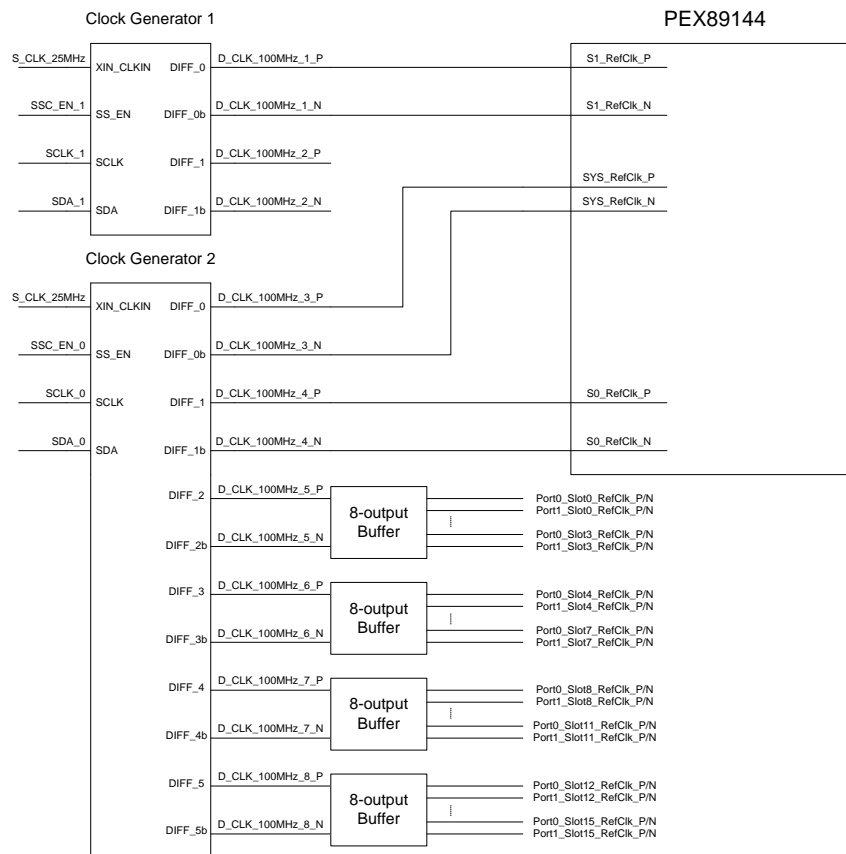
You can see the drive sent out a skip ordered-set around every 260 ns, which is normal for PCIe Gen4 drives when using external refclk and it won't send extra skip ordered-sets since SSC is not turned on.

## 4. Common Clocking Mode With SSC

The common clocking mode with SSC can be described in following chart:



This configuration can be easily achieved with the SANBlaze Gen5 system as follows:



The SSC\_EN\_0 and SSC\_EN\_1 for both Clock Generators are enabled by the Back Panel Switch (both SW6 and SW7 are up). The S0\_RefClk\_P/N and S1\_RefClk\_P/N are ignored in this case because they are for SSC isolation only. The EEPROM image on the SANBlaze Gen5 system specified all ports SerDes are clocked with SYS\_RefClk\_P/N. Both PEX89144 (with the MUX selection) and targets got common RefClk with SSC from Clock Generator 2.

With this clocking mode I sample the link error counter registers every 1 minute and I didn't notice any receiver errors, TLP errors, DLLP errors and Recovery events happened.



## 4.1 Configure Clocking Mode as Common Clock With SSC

We can configure the SANBlaze Gen5 system to use this common clock with SSC as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 clear SW1EESEL0
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
INFO: Clearing bit 16 at addr 0x210
00000210: efdfffff
00000210: efdeffff
00000210: efdeffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW2EESEL1
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: efdeffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 clear SW6SRIS
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
INFO: Clearing bit 21 at addr 0x210
00000210: efdeffff
00000210: efdeffff
00000210: efdeffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 clear SW7SSCL
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
INFO: Clearing bit 22 at addr 0x210
00000210: efdeffff
00000210: ef9effff
00000210: ef9effff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked
GPIO|31|24|23|16|15|8|7|0|
-----|-----|-----|-----|-----|-----|-----|-----|
B P D | D T | T D | S S | S S | S S | S S | D D | D D | D D | D D | D D | D D | D D | D D |
i w i | i P | P P | r W | W W | W W | W W | u u | u u | u u | u u | u u | u u | u u | u u |
t r s | s 3 | 2 1 | v 8 | 7 6 | 5 4 | 3 2 | l a | a a | a a | a a | a a | a a | a a | a a |
3 0 U | W | | | R | D S | S F | C U | E | l | l | l | l | l | l | l | l | l | l | l | l | l | l | l | l |
n i n | s t | | | | | s | i | S | R | r | h | s | E | E | l | l | l | l | l | l | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
N R | b c | | | | | t | s | C | I | c | a | l | S | S | 5 | 4 | 3 | 2 | 1 | 0 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
/ | s | L | h | | | | | L | R | L | S | P | i | L | e | L | L | L | L | L | L | L | L | L | L | L | L | L | L |
C t | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
D | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
g | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
-----|-----|-----|-----|-----|-----|-----|-----|
o o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o |
-----|-----|-----|-----|-----|-----|-----|-----|
1 1 1 | 1 1 | 1 1 | 1 1 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 |
Dir: [0x214]0xffffffff
-----|-----|-----|-----|-----|-----|-----|-----|
1 1 1 | 1 1 | 1 1 | 1 1 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 | 0 0 |
Val: [0x210]0xef9effff
```

After the configuration above we need to shut down the system then power it up as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# shutdown now
```

When the full system is powered off, wait for 10 seconds then power it up. After boot up you can check the clocking mode as follows and ensure it is “common and SSC/SRIS”:

```
[root@vlnun-100-158-IPMI-151 vlnun]# mount 192.168.1.80:/home/sanblaze.local /mnt
[root@vlnun-100-158-IPMI-151 vlnun]# cd /mnt/shynes/Broadcom/tools/G4XTOOLS/Linux/
[root@vlnun-100-158-IPMI-151 Linux]# ./g4Xdiagnostics.x86_64 -sdb /dev/ttyACM3
```

-----  
g4Xdiagnostics v0.0.0.8 - Broadcom Inc. (c) 2021 (Bld-37.46.11.64.6.0)  
-----

```
PEX89144 A0> port
Port: Port number in decimal Stn: Station number in decimal
```

Type: Port type Down-Downstream, Host-Upstream, Fab-Fabric Mode, Mgmt-Management  
MRR: Max Read Request Size in Bytes MPSS: Max Payload Size Supported in Bytes  
MPS: Max Payload Size in Bytes Link: Link Width by Negotiated/Maximum  
Speed: Link Speed by Negotiated/Maximum ClkMode: Clock Mode  
PBus: Primary Bus number in hex, Stat: Link Status

Stn	Port	Type	MRR	MPS	MPSS	LinkSpeed	LinkWidth	ClkMode	PBus	Stat
0	0	Down	512	256	256	Gen5 / Gen5	x16 / x2	Common	88	Up
0	2	Fab	512	256	256	Gen5 / Gen5	x16 / x2	SSC/SRIS	88	Up
0	4	Down	512	256	256	Gen5 / Gen5	x16 / x2	Common	88	Up
0	6	Down	512	256	256	Gen5 / Gen5	x16 / x2	SSC/SRIS	88	Up
0	8	Down	128	256	2048	- / Gen5	- / x2	Common	88	Down
0	10	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	88	Down
0	12	Down	128	256	2048	- / Gen5	- / x2	Common	88	Down
0	14	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	88	Down
1	16	Down	128	256	2048	- / Gen5	- / x2	Common	88	Down
1	18	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	88	Down
1	20	Down	128	256	2048	- / Gen5	- / x2	Common	88	Down
1	22	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	88	Down
1	24	Down	128	256	2048	- / Gen5	- / x2	Common	88	Down
1	26	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	88	Down
1	28	Down	128	256	2048	- / Gen5	- / x2	Common	88	Down
1	30	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	88	Down
2	32	Down	128	256	2048	- / Gen5	- / x16	Common	9A	Down
3	48	Host	128	256	2048	Gen4 / Gen5	x16 / x16	Common	85	Up
4	64	Down	128	256	2048	- / Gen5	- / x16	Common	9D	Down
5	80	Down	128	256	2048	- / Gen5	- / x2	Common	9D	Down
5	82	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	9D	Down
5	84	Down	128	256	2048	- / Gen5	- / x2	SSC	9D	Down
5	86	Down	128	256	2048	- / Gen5	- / x2	Common	9D	Down
5	88	Down	128	256	2048	- / Gen5	- / x2	SRIS	9D	Down
5	90	Down	128	256	2048	- / Gen5	- / x2	Common	9D	Down
5	92	Down	128	256	2048	- / Gen5	- / x2	SSC	9D	Down
5	94	Down	128	256	2048	- / Gen5	- / x2	SSC	9D	Down
6	96	Down	128	256	2048	- / Gen5	- / x8	Common	A8	Down
6	104	Down	128	256	2048	Gen2 / Gen5	x1 / x8	Common	A8	Up
7	112	Down	128	256	2048	- / Gen5	- / x16	Common	A8	Down
8	128	Down	128	256	2048	- / Gen5	- / x2	Common	AD	Down
8	130	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	AD	Down
8	132	Down	128	256	2048	- / Gen5	- / x2	Common	AD	Down
8	134	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	AD	Down
8	136	Down	128	256	2048	- / Gen5	- / x2	Common	AD	Down
8	138	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	AD	Down
8	140	Down	128	256	2048	- / Gen5	- / x2	Common	AD	Down
8	142	Down	128	256	2048	- / Gen5	- / x2	SSC/SRIS	AD	Down

Downstream to Upstream configured groups can not be displayed in this mode.

Downstream Ports:

Port: Port number in decimal, DSBuses: Downstream bus range  
PhysSlot: Chassis Phsyiscal slot in decimal, HostPort : Associated Host Port in decimal  
VID: Vendor ID, DID : Device ID  
SVD: Sub-vendor ID, SID : Sub-device ID  
Attached Device Details: All values are in hex

Port	DSBuses	PhysSlot	HostPort	Attached Device Details					
				VID	DID	SVD	SID	ClassCode	SubClassCode
0	89->89	0	48	-	-	-	-	-	-
2	8A->8A	1	48	-	-	-	-	-	-
4	8B->8B	2	48	-	-	-	-	-	-
6	8C->8C	3	48	-	-	-	-	-	-
8	8D->8D	4	48	-	-	-	-	-	-
10	8E->8E	5	48	-	-	-	-	-	-
12	8F->8F	6	48	-	-	-	-	-	-
14	90->90	7	48	-	-	-	-	-	-
16	91->91	8	48	-	-	-	-	-	-
18	92->92	9	48	-	-	-	-	-	-
20	93->93	10	48	-	-	-	-	-	-
22	94->94	11	48	-	-	-	-	-	-
24	95->95	12	48	-	-	-	-	-	-
26	96->96	13	48	-	-	-	-	-	-
28	97->97	14	48	-	-	-	-	-	-
30	98->98	15	48	-	-	-	-	-	-
32	9B->9B	16	48	-	-	-	-	-	-
64	9E->9E	32	48	-	-	-	-	-	-
80	9F->9F	40	48	-	-	-	-	-	-
82	A0->A0	41	48	-	-	-	-	-	-
84	A1->A1	42	48	-	-	-	-	-	-
86	A2->A2	43	48	-	-	-	-	-	-
88	A3->A3	44	48	-	-	-	-	-	-
90	A4->A4	45	48	-	-	-	-	-	-
92	A5->A5	46	48	-	-	-	-	-	-
94	A6->A6	47	48	-	-	-	-	-	-
96	A9->A9	48	48	-	-	-	-	-	-
120	0->0	0	48	-	-	-	-	-	-
136	B2->B2	68	48	-	-	-	-	-	-

PEX89144 A0> quit

Check the back panel switch 1, 2, 6 and 7 of the system again as follows to ensure they are correct:

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked
GPIO|31|24|23|16|15|8|7|0|
-----|-----|-----|-----|-----|-----|-----|-----|
B P D T T D S S S S S S D D D D D D D D D D D D
i w i i P P P R W W W W W W U U U U U U U U U U
t r s s 3 2 1 v 8 7 6 5 4 3 2 1 a a a a a a a a a a
3 O U W T R D S S R r h s E E l l l l l l l l l l
l n s t c s i S R r h s E E l l l l l l 1 0 0 0
N R b b c t s C I c a l S S 5 4 3 2 1 0 9 8 7 6 5 4 3 2 1 0
/ s L h L R L S P i l e L L L L L L L L L L L L L L
C t D l l n l l l l l l l l l l l l l l l l l l l l
g s w n l l l l l l l l l l l l l l l l l l l l
o o o o o o o o o o o o o o o o o o o o o o o o o o Dir:[0x214]0xffffffff
1 1 1 1 0 1 1 1 1 1 0 0 1 1 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 Val:[0x210]0xef9effff
```

## 4.2 Check Link Error Counters

We can clear the PCIe link error counters with command line as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5 -r
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active
```

After clearing the PCIe link error counters now we can dump them as follows when some I/O running in the background and dump every other 60 seconds.

```
[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active
```

### 4.3 Clock Scope Shot from PCIe Switch Side



- The channel “C4” is the reference clock obtained from PCIe switch side, and the measurement shows it is around 99.7 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 99.7 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 99.7 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts and displays the instantaneous frequency variation of the clock as a function of time. You can see the SSC modulation frequency is about 32 kHz which matches clock generator specification (Modulation Frequency: Min = 30 kHz, Typ = 31.5 kHz and Max = 33 kHz).
- From the scope shot above we can tell the clock on the PCIe switch side has SSC turned on, and the modulation technique is down-spread.

## 4.4 Clock Scope Shot from Drive Side



- The channel “C4” is the reference clock obtained from drive side, and the measurement shows it is around 99.65 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 99.65 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 99.65 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts and displays the instantaneous frequency variation of the clock as a function of time. You can see the SSC modulation frequency is about 32 kHz which matches clock generator specification (Modulation Frequency: Min = 30 kHz, Typ = 31.5 kHz and Max = 33 kHz).
- From the scope shot above we can tell the clock on the drive side has SSC turned on, and the modulation technique is down-spread.


## 4.5 Skip Ordered-Sets from Bus Trace

Packet	R←	8.0	SKIP	SKIP Symbols	END	P	LFSR Symbols	Idle	Time Stamp
0	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.027 us	0000 . 000 000 000 000 s
1	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.035 us	0000 . 000 006 043 000 s
2	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.035 us	0000 . 000 012 094 000 s
3	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.021 us	0000 . 000 018 145 000 s
4	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.015 us	0000 . 000 024 182 000 s
5	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.023 us	0000 . 000 030 213 000 s
6	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.035 us	0000 . 000 036 252 000 s
7	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.035 us	0000 . 000 042 303 000 s
8	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.025 us	0000 . 000 048 354 000 s
9	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.015 us	0000 . 000 054 395 000 s
10	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.021 us	0000 . 000 060 426 000 s
11	R←	x4	SKIP	AAAAAAAAAA...	E1	1	** ** *	6.033 us	0000 . 000 066 463 000 s
12	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.037 us	0000 . 000 072 512 000 s
13	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.025 us	0000 . 000 078 565 000 s
14	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.017 us	0000 . 000 084 606 000 s
15	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.021 us	0000 . 000 090 639 000 s
16	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.031 us	0000 . 000 096 676 000 s
17	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.037 us	0000 . 000 102 723 000 s
18	R←	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
19	R→	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	3.343 us	0000 . 000 108 776 000 s
19	R→	x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	Idle	Time Stamp
								503.750 ns	0000 . 000 112 119 000 s

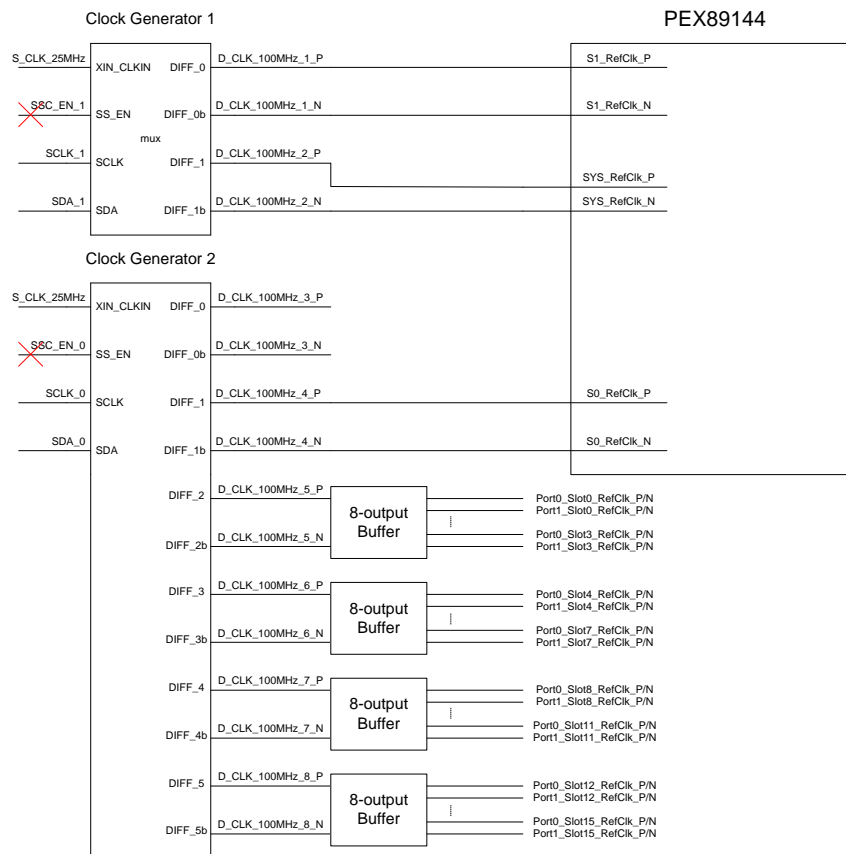
You can see the drive sent out a skip ordered-set around every 6  $\mu$ s, which is normal for PCIe Gen3 drives when the drive is using an external reference clock. Although SSC is turned on, the drive doesn't know the external refclk has SSC turned on or not, so no extra skip ordered-set is sent out.

## 5. Separate Reference Clock Without SSC (Static SRNS)

The separate reference clock without SSC (Static SRNS) can be described in following chart:



This configuration can be easily achieved with the SANBlaze Gen5 system as follows:



The SSC\_EN\_0 and SSC\_EN\_1 for Clock Generator 1 and 2 are disabled by the Back Panel Switches (both SW6 and SW7 are down). The S0\_RefClk\_P/N and S1\_RefClk\_P/N are ignored in this case because they are for SSC isolation only. The EEPROM image on the SANBlaze Gen5 system specified all ports SerDes are clocked with SYS\_RefClk\_P/N. PEX89144 got separate RefClk without SSC from Clock Generator 1 with the MUX selection, and all targets got separate RefClk without SSC from Clock Generator 2.

With this clocking mode I sample the link error counter registers every 1 minute and no receiver errors, TLP errors, DLLP errors and Recovery events happened.

# 5.1 Configure Clocking Mode as Separate Clock Without SSC

We can configure the SANBlaze Gen5 system to use this separate clock without SSC as follows:

```

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 clear SW1EESEL0
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
INFO: Clearing bit 16 at addr 0x210
00000210: efdfefdf
00000210: efdfefdf
00000210: efdfefdf
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW2EESEL1
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: efdfefdf
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW6SRIS
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: effeffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW7SSCL
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: effeffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

```

```

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked
GPIO|31| 24|23| 16|15| 8|7| 0|
|-----|-----|-----|-----|-----|
| B | P | D | D | T | T | S | S | S | S | S | S | D | D | D | D | D | D | D | D | D | D | D | D |
| i | w | i | P | P | P | R | W | W | W | W | W | u | u | u | u | u | u | u | u | u | u | u | u |
| t | r | s | s | 3 | 2 | 1 | v | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | a | a | a | a | a | a | a | a | a | a |
| 3 | O | U | W | | | | R | D | S | S | F | C | U | E | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| 1 | n | s | t | | | | s | i | S | R | r | h | s | E | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 |
| N | R | b | c | | | | t | s | C | I | C | a | l | S | S | 5 | 4 | 3 | 2 | 1 | 0 | 9 | 8 | 7 | 6 | 5 | 4 |
| / | s | L | h | | | | L | R | L | S | P | i | L | e | L | L | L | L | L | L | L | L | L | L | L |
| C | t | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-----|-----|-----|-----|-----|
| o | o | o | o | o | o | o | o | o | o | o | o | o | o | o | o | o | o | o | o | o | o | o | o |
| 1 | 1 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
Dir: [0x214]0xffffffff
Val: [0x210]0xfffffff

```

After the configuration above we need to shut down the system then power it up as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# shutdown now
```

When the full system is powered off, wait for 10 seconds then power it up. After boot up you can check the clocking mode as follows and ensure it is "SRIS" because all separate clock modes will show up as "SRIS", no matter if it has SSC signal enabled or not.

```

[root@vlnun-100-158-IPMI-151 vlnun]# mount 192.168.1.80:/home/sanblaze.local /mnt
[root@vlnun-100-158-IPMI-151 vlnun]# cd /mnt/shyres/Broadcom/tools/G4XTOOLS/Linux/
[root@vlnun-100-158-IPMI-151 Linux]# ./g4Xdiagnostics.x86_64 -sdb /dev/ttyACM3

```

```

-----
g4Xdiagnostics v0.0.0.8 - Broadcom Inc. (c) 2021 (Bld-37.46.11.64.6.0)
-----
PEX89144 A0> port
Port: Port number in decimal Stn: Station number in decimal
Type: Port type Down-Downstream, Host-Upstream, Fab-Fabric Mode, Mgmt-Management
MRR: Max Read Request Size in Bytes MPSS: Max Payload Size Supported in Bytes
MPS: Max Payload Size in Bytes Link: Link Width by Negotiated/Maximum
Speed: Link Speed by Negotiated/Maximum ClkMode: Clock Mode
PBus: Primary Bus number in hex, Stat: Link Status

```



Stn	Port	Type	MRR	MPS	MPSS	LinkSpeed	LinkWidth	ClkMode	PBus	Stat
0	0	Down	128	128	2048	Gen4 / Gen5	x4 / x4	SRIS	88	Up
0	4	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	6	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	8	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	10	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	12	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	14	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	16	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	18	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	20	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	22	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	24	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	26	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	28	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	30	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
2	32	Down	128	128	2048	- / Gen5	- / x16	SRIS	9A	Down
3	48	Host	128	128	2048	Gen4 / Gen5	x16 / x16	SRIS	85	Up
4	64	Down	128	128	2048	- / Gen5	- / x16	SRIS	9D	Down
5	80	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	82	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	84	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	86	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	88	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	90	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	92	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	94	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
6	96	Down	128	128	2048	- / Gen5	- / x8	SRIS	A8	Down
6	104	Down	128	128	2048	Gen2 / Gen5	x1 / x8	Common	A8	Up
7	112	Down	128	128	2048	- / Gen5	- / x16	SRIS	A8	Down
8	128	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	130	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	132	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	134	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	136	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	138	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	140	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	142	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down

Host Ports:

BusPrim: Primary Bus number in hex, BusSec: Secondary Bus number in hex  
 BusSub: Subordinate Bus number in hex, Assigned DS Ports: In decimal

Port BusPrim BusSec BusSub Assigned DS Ports

-----  
 48 85 86 B5 0, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 32, 64, 80, 82, 84, 86, 88, 90, 92, 94, 96, 104, 112, 128, 130, 132, 134, 136, 138, 140, 142,

Downstream Ports:

Port: Port number in decimal, DSBuses: Downstream bus range  
 PhysSlot: Chassis Phsyiscal slot in decimal, HostPort: Associated Host Port in decimal  
 VID: Vendor ID, DID: Device ID  
 SVD: Sub-vendor ID, SID: Sub-device ID  
 Attached Device Details: All values are in hex

Port	DSBuses	PhysSlot	HostPort	Attached Device Details					
				VID	DID	SVD	SID	ClassCode	SubClassCode
0	89->89	0	48	-	-	-	-	-	-
4	8B->8B	2	48	-	-	-	-	-	-
6	8C->8C	3	48	-	-	-	-	-	-
8	8D->8D	4	48	-	-	-	-	-	-
10	8E->8E	5	48	-	-	-	-	-	-
12	8F->8F	6	48	-	-	-	-	-	-
14	90->90	7	48	-	-	-	-	-	-
16	91->91	8	48	-	-	-	-	-	-
18	92->92	9	48	-	-	-	-	-	-
20	93->93	10	48	-	-	-	-	-	-
22	94->94	11	48	-	-	-	-	-	-
24	95->95	12	48	-	-	-	-	-	-
26	96->96	13	48	-	-	-	-	-	-
28	97->97	14	48	-	-	-	-	-	-
30	98->98	15	48	-	-	-	-	-	-
32	9B->9B	16	48	-	-	-	-	-	-
64	9E->9E	32	48	-	-	-	-	-	-
80	9F->9F	40	48	-	-	-	-	-	-
82	A0->A0	41	48	-	-	-	-	-	-
84	A1->A1	42	48	-	-	-	-	-	-
86	A2->A2	43	48	-	-	-	-	-	-
88	A3->A3	44	48	-	-	-	-	-	-
90	A4->A4	45	48	-	-	-	-	-	-
92	A5->A5	46	48	-	-	-	-	-	-
94	A6->A6	47	48	-	-	-	-	-	-
96	A9->A9	48	48	-	-	-	-	-	-
104	AA->AA	52	48	-	-	-	-	-	-
112	AB->AB	56	48	-	-	-	-	-	-
128	AE->AE	64	48	-	-	-	-	-	-
130	AF->AF	65	48	-	-	-	-	-	-
132	B0->B0	66	48	-	-	-	-	-	-
134	B1->B1	67	48	-	-	-	-	-	-

```

136 B2->B2 68 48 - - - - -
138 B3->B3 69 48 - - - - -
140 B4->B4 70 48 - - - - -
142 B5->B5 71 48 - - - - -
PEX89144 A0> quit

```

Check the back panel switch 1, 2, 6 and 7 of the system again as follows to ensure they are correct:

```

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked
GPIO|31      24|23      16|15      8|7      0
+---+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| B  | P  | D  | T  | T  | D  | S  | S  | S  | S  | S  | S  | D  | D  | D  | D  | D  | D  | D  | D  | D  | D  | |
| i  | w  | i  | P  | P  | P  | R  | W  | W  | W  | W  | W  | u  | u  | u  | u  | u  | u  | u  | u  | u  | u  |
| t  | r  | s  | 3  | 2  | 1  | v  | 8  | 7  | 6  | 5  | 4  | 3  | 2  | 1  | a  | a  | a  | a  | a  | a  | a  |
| 3  | O  | U  | W  |   |   |   | R  | D  | S  | S  | F  | C  | U  | E  |   |   |   |   |   |   |   |   |
| i  | n  | s  | t  |   |   |   | s  | i  | S  | R  | r  | h  | s  | E  |   |   |   |   |   |   |   |   |
| N  | R  | b  | c  |   |   |   | t  | s  | I  | C  | I  | c  | a  | l  | S  |   |   |   |   |   |   |   |
| /  | s  | L  | h  | d  |   |   | L  | R  | L  | S  |   | P  | i  | L  | e  |   |   |   |   |   |   |   |
| C  | t  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
|   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
+---+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  | o  |
|   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
+---+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1  | 1  | 1  | 0  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  |
|   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
+---+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Dir: [0x214]0xffffffff
Val: [0x210]0xeffeffff

```

## 5.2 Check Link Error Counters

We can clear the PCIe link error counters using the command line as follows:

```

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5 -r
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIEDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

```

After clearing the PCIe link error counters now we can dump them as follows when some I/O running in the background and dump every other 60 seconds.

```

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIEDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIEDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIEDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIEDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active


```

### 5.3 Clock Scope Shot from PCIe Switch Side



- The channel “C4” is the reference clock obtained from PCIe switch side, and the measurement shows it is around 100 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 100 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 100 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts and displays the instantaneous frequency variation of the clock as a function of time.
- From the scope shot above we can tell the clock on the switch side does not have SSC turned on.

## 5.4 Clock Scope Shot from Drive Side



- The channel “C4” is the reference clock obtained from PCIe switch side, and the measurement shows it is around 100 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 100 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 100 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts and displays the instantaneous frequency variation of the clock as a function of time. The 624.02 kHz signal is just jitter picked up and may be from power supply.
- From the scope shot above we can tell the clock on the drive side does not have SSC turned on.


## 5.5 Skip Ordered-Sets from Bus Trace

Packet	R←	8.0 x4	SKIP	SKIP Symbols	END	P	LFSR Symbols	Idle	Time Stamp
0	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 000 000 000 s
1	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 006 029 000 s
2	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 012 058 000 s
3	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 018 087 000 s
4	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 024 116 000 s
5	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.011 us	0000 . 000 030 145 000 s
6	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	1	** ** *	6.013 us	0000 . 000 036 172 000 s
7	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	1	** ** *	6.013 us	0000 . 000 042 201 000 s
8	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 048 230 000 s
9	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 054 259 000 s
10	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 060 288 000 s
11	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 066 317 000 s
12	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.011 us	0000 . 000 072 346 000 s
13	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 078 373 000 s
14	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 084 402 000 s
15	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	6.013 us	0000 . 000 090 431 000 s
16	R←	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	367.000 ns	0000 . 000 096 460 000 s
17	R→	8.0 x4	SKIP	AAAAAAAAAA...	E1	*	** ** *	503.750 ns	0000 . 000 096 827 000 s

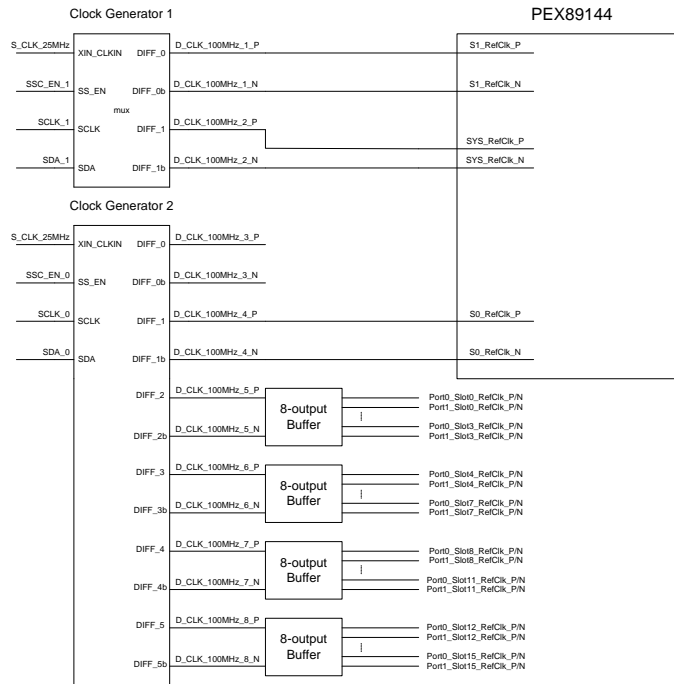
You can see the drive sent out a skip ordered-set around every 6  $\mu$ s, which is normal for PCIe Gen3 drives when the drive is using an external reference clock. The drive doesn't know if the external refclk has SSC turned on or not, so no extra skip ordered-set is sent out.

## 6. Separate Reference Clock Independent SSC (Static SRIS)

The separate reference clock independent SSC (Static SRIS) can be described in following chart:



This configuration can be easily achieved with the SANBlaze Gen5 system as follows:



The SSC\_EN\_0 and SSC\_EN\_1 for Clock Generator 1 and 2 are enabled by the Back Panel Switches (SW6 is down and SW7 is up). The S0\_RefClk\_P/N and S1\_RefClk\_P/N are ignored in this case because they are for SSC isolation only. The EEPROM image on the SANBlaze Gen5 system specified all ports SerDes are clocked with SYS\_RefClk\_P/N. PEX89144 got separate RefClk with SSC from Clock Generator 1 with the MUX selection, and all targets got separate RefClk with SSC from Clock Generator 2.

With this clocking mode I sample the link error counter registers every 1 minute and I didn't notice any receiver errors, TLP errors, DLLP errors and Recovery events happened.

### 6.1 Configure Clocking Mode as Separate Clock With SSC

We can configure the SANBlaze Gen5 system to use this separate clock with SSC as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW1ESEL0
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: ef9fffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000
```

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 clear SW2EESEL1
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
INFO: Clearing bit 17 at addr 0x210
00000210: ef9fffff
00000210: ef9dffff
00000210: ef9dffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000
```

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW6SRIS
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: ffffffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: ffffffff
```

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 clear SW7SSCL
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
INFO: Clearing bit 22 at addr 0x210
00000210: efbdfbff
00000210: efbdfbff
00000210: efbdfbff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000
```

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked
```

```
GPIO[31 24|23 16|15 8|7 0]
B P D D T T D S S S S S S S S D D D D D D D D D D D D D D
i w i i P P P r W W W W W W W W W W U U U U U U U U U U U U
t r s s s 3 2 1 v 8 7 6 5 4 3 2 1 a a a a a a a a a a a a a a a
3 O U W | | | | | | | | | | | | | | | | | | | | | | | | | |
n s t | | | | | | | | | | | | | | | | | | | | | | | | | |
N R b c | | | | | | | | | | | | | | | | | | | | | | | | | |
/ s L h d | | | | | | | | | | | | | | | | | | | | | | | | | |
C t | | | | | | | | | | | | | | | | | | | | | | | | | |
o o o o o o o o o o o o o o o o o o o o o o o o o o o o Dir: [0x214]0xffffffff
1 1 1 0 1 1 1 1 1 1 0 1 1 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 Val: [0x210]0xefbdfbff
```

After the configuration above we need to shut down the system then power it up as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# shutdown now
```

When the full system is powered off, wait for 10 seconds then power it up. After boot up you can check the clocking mode as follows and ensure it is "SRIS" because all separate clock modes will show up as "SRIS", no matter it has SSC enabled or not.

```
[root@vlnun-100-158-IPMI-151 vlnun]# mount 192.168.1.80:/home/sanblaze.local /mnt
[root@vlnun-100-158-IPMI-151 vlnun]# cd /mnt/shyunes/Broadcom/tools/G4XTOOLS/Linux/
[root@vlnun-100-158-IPMI-151 Linux]# ./g4Xdiagnostics.x86_64 -sdb /dev/ttyACM3
```

```
-----
g4Xdiagnostics v0.0.0.8 - Broadcom Inc. (c) 2021 (Bld-37.46.11.64.6.0)
-----
```

```
PEX89144 A0> port
Port: Port number in decimal Stn: Station number in decimal
Type: Port type Down-Downstream, Host-Upstream, Fab-Fabric Mode, Mgmt-Management
MRR: Max Read Request Size in Bytes MPSS: Max Payload Size Supported in Bytes
MPS: Max Payload Size in Bytes Link: Link Width by Negotiated/Maximum
Speed: Link Speed by Negotiated/Maximum ClkMode: Clock Mode
PBus: Primary Bus number in hex, Stat: Link Status
```

Stn	Port	Type	MRR	MPS	MPSS	LinkSpeed	LinkWidth	ClkMode	PBus	Stat
0	0	Down	512	256	256	Gen5 / Gen5	x16 / x2	SRIS	88	Up
0	2	Fab	512	256	256	Gen5 / Gen5	x16 / x2	SRIS	88	Up
0	4	Down	512	256	256	Gen5 / Gen5	x16 / x2	SRIS	88	Up
0	6	Fab	512	256	256	Gen5 / Gen5	x16 / x2	SRIS	88	Up
0	8	Fab	512	256	256	Gen5 / Gen5	x16 / x16	SRIS	88	Up
1	24	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	26	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	28	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	30	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
2	32	Down	128	128	2048	- / Gen5	- / x16	SRIS	9A	Down
3	48	Host	128	128	2048	Gen4 / Gen5	x16 / x16	Common	85	Up
4	64	Down	128	128	2048	- / Gen5	- / x16	SRIS	9D	Down
5	80	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	82	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down

```

5 84 Down 128 128 2048 - / Gen5 - / x2 SRIS 9D Down
5 86 Down 128 128 2048 - / Gen5 - / x2 SRIS 9D Down
5 88 Down 128 128 2048 - / Gen5 - / x2 SRIS 9D Down
5 90 Down 128 128 2048 - / Gen5 - / x2 SRIS 9D Down
5 92 Down 128 128 2048 - / Gen5 - / x2 SRIS 9D Down
5 94 Down 128 128 2048 - / Gen5 - / x2 SRIS 9D Down
6 96 Down 128 128 2048 - / Gen5 - / x8 SRIS A8 Down
6 104 Down 128 128 2048 Gen2 / Gen5 x1 / x8 SSC A8 Up
7 112 Down 128 128 2048 - / Gen5 - / x16 SRIS A8 Down
8 128 Down 128 128 2048 - / Gen5 - / x2 SRIS AD Down
8 130 Down 128 128 2048 - / Gen5 - / x2 SRIS AD Down
8 132 Down 128 128 2048 - / Gen5 - / x2 SRIS AD Down
8 134 Down 128 128 2048 - / Gen5 - / x2 SRIS AD Down
8 136 Down 128 128 2048 - / Gen5 - / x2 SRIS AD Down
8 138 Down 128 128 2048 - / Gen5 - / x2 SRIS AD Down
8 140 Down 128 128 2048 - / Gen5 - / x2 SRIS AD Down
8 142 Down 128 128 2048 - / Gen5 - / x2 SRIS AD Down

```

Host Ports:

```

BusPrim: Primary Bus number in hex,      BusSec: Secondary Bus number in hex
BusSub:  Subordinate Bus number in hex,   Assigned DS Ports: In decimal

```

```

Port  BusPrim  BusSec  BusSub  Assigned DS Ports
-----
48    85        86      B5      0, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 32, 64, 80, 82, 84, 86, 88, 90, 92, 94,
96, 104, 112, 128, 130, 132, 134, 136, 138, 140, 142,

```

Downstream Ports:

```

Port: Port number in decimal,      DSBus: Downstream bus range
PhysSlot: Chassis Physical slot in decimal,  HostPort : Associated Host Port in decimal
VID: Vendor ID,                    DID : Device ID
SVD: Sub-vendor ID,                SID : Sub-device ID
Attached Device Details: All values are in hex

```


Port	DSBuses	PhysSlot	HostPort	Attached Device Details					
				VID	DID	SVD	SID	ClassCode	SubClassCode
0	89->89	0	48	-	-	-	-	-	-
4	8B->8B	2	48	-	-	-	-	-	-
6	8C->8C	3	48	-	-	-	-	-	-
8	8D->8D	4	48	-	-	-	-	-	-
10	8E->8E	5	48	-	-	-	-	-	-
12	8F->8F	6	48	-	-	-	-	-	-
14	90->90	7	48	-	-	-	-	-	-
16	91->91	8	48	-	-	-	-	-	-
18	92->92	9	48	-	-	-	-	-	-
20	93->93	10	48	-	-	-	-	-	-
22	94->94	11	48	-	-	-	-	-	-
24	95->95	12	48	-	-	-	-	-	-
26	96->96	13	48	-	-	-	-	-	-
28	97->97	14	48	-	-	-	-	-	-
30	98->98	15	48	-	-	-	-	-	-
32	9B->9B	16	48	-	-	-	-	-	-
64	9E->9E	32	48	-	-	-	-	-	-
80	9F->9F	40	48	-	-	-	-	-	-
82	A0->A0	41	48	-	-	-	-	-	-
84	A1->A1	42	48	-	-	-	-	-	-
86	A2->A2	43	48	-	-	-	-	-	-
88	A3->A3	44	48	-	-	-	-	-	-
90	A4->A4	45	48	-	-	-	-	-	-
92	A5->A5	46	48	-	-	-	-	-	-
94	A6->A6	47	48	-	-	-	-	-	-
96	A9->A9	48	48	-	-	-	-	-	-
104	AA->AA	52	48	-	-	-	-	-	-
112	AB->AB	56	48	-	-	-	-	-	-
128	AE->AE	64	48	-	-	-	-	-	-
130	AF->AF	65	48	-	-	-	-	-	-
132	B0->B0	66	48	-	-	-	-	-	-
134	B1->B1	67	48	-	-	-	-	-	-
136	B2->B2	68	48	-	-	-	-	-	-
138	B3->B3	69	48	-	-	-	-	-	-
140	B4->B4	70	48	-	-	-	-	-	-
142	B5->B5	71	48	-	-	-	-	-	-

Check the back panel switch 1, 2, 6 and 7 of the system again as follows to ensure they are correct:

```

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked

```







## 6.3 Clock Scope Shot from PCIe Switch Side



- The channel “C4” is the reference clock obtained from the PCIe switch side, and the measurement shows it is around 100 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 100 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 100 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts/displays the instantaneous frequency variation of the clock as a function of time.
- From the scope shot above we can tell the clock in the switch side has no SSC turned on, which is not expected. Turns out when in both static SRIS mode and dynamic SRIS mode, the PEX89144 chip will not link so we can only provide refclk without SSC, and the PEX89144 flash image will use the internal PLL to turn on SSC for the PCIe switch. There’s no way for us to check the clock SSC which is generated inside the chip.

## 6.4 Clock Scope Shot from Drive Side



- The channel “C4” is the reference clock obtained from drive side, and the measurement shows it is around 99.67 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 99.67 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 99.67 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts and displays the instantaneous frequency variation of the clock as a function of time. You can see the SSC modulation frequency is about 31.9 kHz which matches clock generator specification (Modulation Frequency: Min = 30 kHz, Typ = 31.5 kHz and Max = 33 kHz).
- From the scope shot above we can tell the clock in the drive side has SSC turned on, and the modulation technique is down-spread.

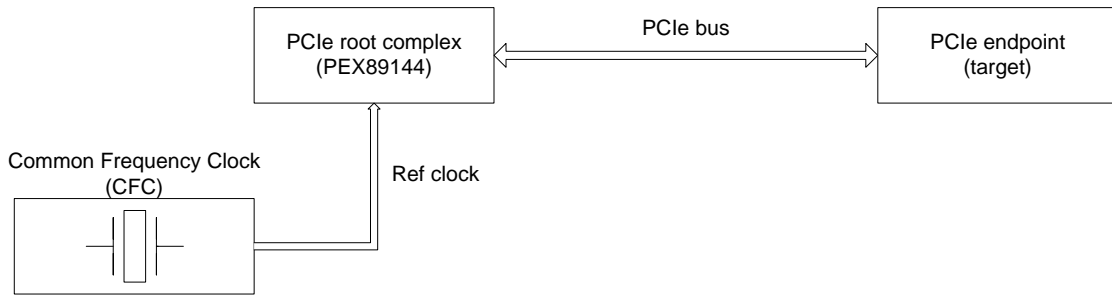
## 6.5 Skip Ordered-Sets from Bus Trace

Trace View / PCIe Revision 4.0									
Packet	R←	8.0	SKIP	SKIP Symbols	END	P	LFSR Symbols	Idle	Time Stamp
39246	x4			AAAAAAAAAA...	E1	*	** ** *	6.027 us	0000 . 000 000 000 000 s
39247	x4			AAAAAAAAAA...	E1	*	** ** *	6.015 us	0000 . 000 006 043 000 s
39248	x4			AAAAAAAAAA...	E1	*	** ** *	6.021 us	0000 . 000 012 074 000 s
39249	x4			AAAAAAAAAA...	E1	*	** ** *	6.029 us	0000 . 000 018 111 000 s
39250	x4			AAAAAAAAAA...	E1	*	** ** *	6.037 us	0000 . 000 024 156 000 s
39251	x4			AAAAAAAAAA...	E1	*	** ** *	6.029 us	0000 . 000 030 209 000 s
39252	x4			AAAAAAAAAA...	E1	*	** ** *	6.019 us	0000 . 000 036 254 000 s
39253	x4			AAAAAAAAAA...	E1	*	** ** *	6.017 us	0000 . 000 042 289 000 s
39254	x4			AAAAAAAAAA...	E1	*	** ** *	6.029 us	0000 . 000 048 322 000 s
39255	x4			AAAAAAAAAA...	E1	1	** ** *	6.037 us	0000 . 000 054 367 000 s
39256	x4			AAAAAAAAAA...	E1	*	** ** *	6.031 us	0000 . 000 060 420 000 s
39257	x4			AAAAAAAAAA...	E1	*	** ** *	6.019 us	0000 . 000 066 467 000 s
39258	x4			AAAAAAAAAA...	E1	*	** ** *	6.017 us	0000 . 000 072 502 000 s
39259	x4			AAAAAAAAAA...	E1	*	** ** *	6.025 us	0000 . 000 078 535 000 s
39261	x4			AAAAAAAAAA...	E1	*	** ** *	6.037 us	0000 . 000 084 576 000 s
39262	x4			AAAAAAAAAA...	E1	*	** ** *	6.033 us	0000 . 000 090 629 000 s
39268	x4			AAAAAAAAAA...	E1	*	** ** *	6.021 us	0000 . 000 096 678 000 s

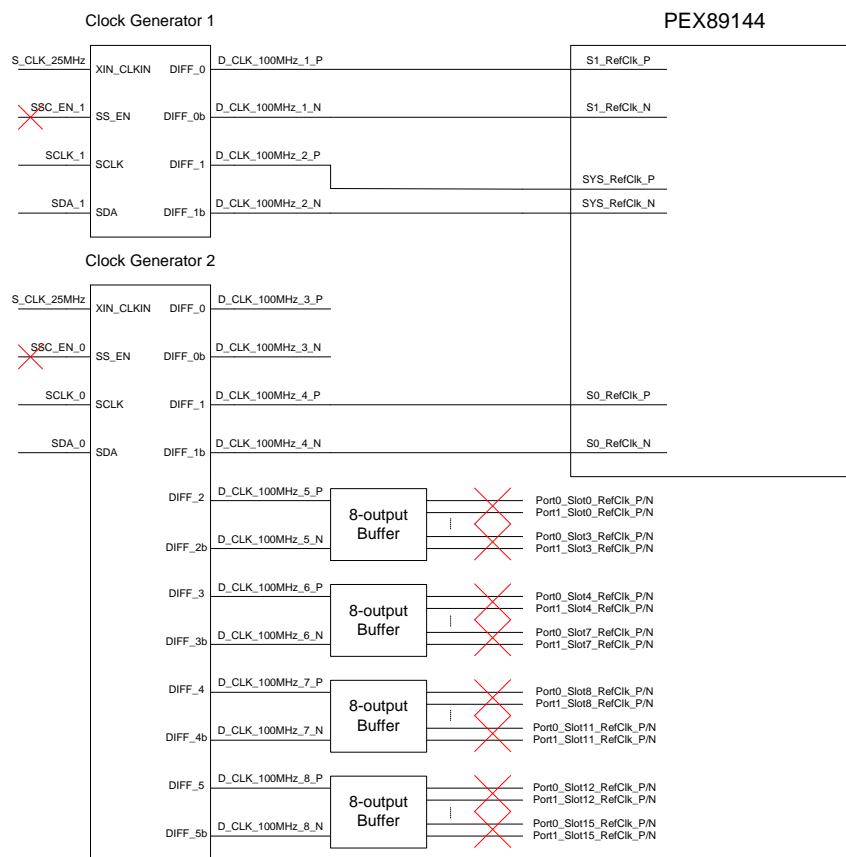
You can see the drive sent out a skip ordered-set around every 6  $\mu$ s, which is normal for PCIe Gen3 drives when drive is using an external reference clock. Although SSC is turned on, the drive doesn't know the external refclk has SSC turned on or not so no extra skip ordered-set is sent out.

## 7. RefClk Without SSC for PEX Only (No RefClk to Target Drive, Dynamic SRNS)

The RefClk without SSC for PEX only (no RefClk to the target drive) can be described in following chart:



This PCIe clocking mode can be easily achieved in the SANBlaze Gen5 system as follows:



The SSC\_EN\_0 and SSC\_EN\_1 for Clock Generator 1 and 2 are disabled by the Back Panel Switches (both SW6 and SW7 are down). The S0\_RefClk\_P/N and S1\_RefClk\_P/N are ignored in this case because they are for SSC isolation only. The EEPROM image on the SANBlaze Gen5 system specified all ports SerDes are clocked with SYS\_RefClk\_P/N. The PEX89144 got separate RefClk without SSC from Clock Generator 1 with the MUX selection. The RefClk for all targets can be disabled through **sb\_i2c2** commands. In general you don't have to disable the RefClk for all targets, just disable the RefClk for the slot you are testing. For example, the following commands can be used to disable/enable RefClk for slot 7:

```
sb_i2c2 -n 1 -d 7 -f EN_REFCLK_P0 -w 0      # disable RefClk for slot 7
sb_i2c2 -n 1 -d 7 -f EN_REFCLK_P0 -w 1      # enable RefClk for slot 7
```

You also can tape the RefClk on the PCIe riser so no RefClk for targets. After disabling the refclk you may need to reset the drive to bring it up on the bus.

With this clocking mode I sample the link error counter registers every 1 minute and notice there were no receiver errors, TLP errors, DLLP errors or Recovery events.

## 7.1 Configure Clocking Mode as Dynamic SRNS

We can configure the SANBlaze Gen5 system to use this dynamic SRNS as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 clear SW1EESEL0
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
INFO: Clearing bit 16 at addr 0x210
00000210: efdfffff
00000210: efdfffff
00000210: efdfffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW2EESEL1
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: efdfffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW6SRIS
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: efdfffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW7SSCL
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: efdfffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked
GPIO|31      24|23      16|15      8|7      0
+-----+-----+-----+-----+
| B P D | D T | T T | S S | S S | S S | S S | S S | D D | D D | D D | D D | D D | D D | D D | D D | | | | | | | | | | | | | | |
| i w i | i P | P P | r W | W W | W W | W W | W W | u u | u u | u u | u u | u u | u u | u u | u u |
| t r s | s s | 3 2 | 1 v | 8 7 | 6 5 | 4 3 | 2 1 | a a | a a | a a | a a | a a | a a | a a | a a |
| 3 0 U | W | | | | R | D | S | S | F | C | U | E | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| 1 n s | t | | | | s | i | S | R | r | h | s | E | E | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| N R b | c | | | | t | s | C | I | c | a | l | S | S | 5 | 4 | 3 | 2 | 1 | 0 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| / s L | h | | | | L | R | L | S | P | i | L | e | L | L | L | L | L | L | L | L | L | L | L | L | L | L | L | L |
| C t | d | | | | g | s | t | w | n | l | l | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
+-----+-----+-----+-----+
| o o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o | o o |
| 1 1 1 | 0 1 | 1 1 | 1 1 | 1 1 | 1 1 | 1 1 | 1 1 | 0 1 | 1 1 | 1 1 | 1 1 | 1 1 | 1 1 | 1 1 | 1 1 |
Dir: [0x214]0xffffffff
Val: [0x210]0xefffffff
```

After the configuration above, we need to shut down the system then power it up as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# shutdown now
```

When the full system is powered off, wait for 10 seconds then power it up. After boot up you can check the clocking mode as follows and ensure it is "SRIS" because all separate clock modes will show up as "SRIS", no matter if it has SSC enabled or not.

```
[root@vlnun-100-158-IPMI-151 vlnun]# mount 192.168.1.80:/home/sanblaze.local /mnt
[root@vlnun-100-158-IPMI-151 vlnun]# cd /mnt/shyenes/Broadcom/tools/G4XTOOLS/Linux/
[root@vlnun-100-158-IPMI-151 Linux]# ./g4Xdiagnostics.x86_64 -sdb /dev/ttyACM3
```

PEX89144 A0> port

Port: Port number in decimal                     Stn: Station number in decimal  
Type: Port type Down-Downstream, Host-Upstream, Fab-Fabric Mode, Mgmt-Management  
MRR: Max Read Request Size in Bytes           MPSS: Max Payload Size Supported in Bytes  
MPS: Max Payload Size in Bytes               Link: Link Width by Negotiated/Maximum  
Speed: Link Speed by Negotiated/Maximum     ClkMode: Clock Mode  
PBus: Primary Bus number in hex,             Stat: Link Status

Stn	Port	Type	MRR	MPS	MPSS	LinkSpeed	LinkWidth	ClkMode	PBus	Stat
0	0	Down	128	128	2048	Gen4 / Gen5	x4 / x4	SRIS	88	Up
0	4	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	6	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	8	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	10	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	12	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
0	14	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	16	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	18	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	20	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	22	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	24	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	26	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	28	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	30	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
2	32	Down	128	128	2048	- / Gen5	- / x16	SRIS	9A	Down
3	48	Host	128	128	2048	Gen4 / Gen5	x16 / x16	SRIS	85	Up
4	64	Down	128	128	2048	- / Gen5	- / x16	SRIS	9D	Down
5	80	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	82	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	84	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	86	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	88	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	90	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	92	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	94	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
6	96	Down	128	128	2048	- / Gen5	- / x8	SRIS	A8	Down
6	104	Down	128	128	2048	Gen2 / Gen5	x1 / x8	Common	A8	Up
7	112	Down	128	128	2048	- / Gen5	- / x16	SRIS	A8	Down
8	128	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	130	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	132	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	134	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	136	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	138	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	140	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	142	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down

Host Ports:

BusPrim: Primary Bus number in hex,             BusSec: Secondary Bus number in hex  
BusSub: Subordinate Bus number in hex,         Assigned DS Ports: In decimal

Port	BusPrim	BusSec	BusSub	Assigned DS Ports
48	85	86	B5	0, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 32, 64, 80, 82, 84, 86, 88, 90, 92, 94, 96, 104, 112, 128, 130, 132, 134, 136, 138, 140, 142,

Downstream Ports:

Port: Port number in decimal,                     DSBuses: Downstream bus range  
PhysSlot: Chassis Physical slot in decimal,     HostPort: Associated Host Port in decimal  
VID: Vendor ID,                                     DID: Device ID  
SVD: Sub-vendor ID,                                SID: Sub-device ID  
Attached Device Details: All values are in hex

Port	DSBuses	PhysSlot	HostPort	Attached Device Details				
				VID	DID	SVD	SID	ClassCode
0	89->89	0	48	-	-	-	-	-
4	8B->8B	2	48	-	-	-	-	-
6	8C->8C	3	48	-	-	-	-	-
8	8D->8D	4	48	-	-	-	-	-
10	8E->8E	5	48	-	-	-	-	-
12	8F->8F	6	48	-	-	-	-	-
14	90->90	7	48	-	-	-	-	-
16	91->91	8	48	-	-	-	-	-
18	92->92	9	48	-	-	-	-	-
20	93->93	10	48	-	-	-	-	-
22	94->94	11	48	-	-	-	-	-
24	95->95	12	48	-	-	-	-	-
26	96->96	13	48	-	-	-	-	-
28	97->97	14	48	-	-	-	-	-
30	98->98	15	48	-	-	-	-	-
32	9B->9B	16	48	-	-	-	-	-
64	9E->9E	32	48	-	-	-	-	-
80	9F->9F	40	48	-	-	-	-	-
82	A0->A0	41	48	-	-	-	-	-
84	A1->A1	42	48	-	-	-	-	-
86	A2->A2	43	48	-	-	-	-	-
88	A3->A3	44	48	-	-	-	-	-
90	A4->A4	45	48	-	-	-	-	-
92	A5->A5	46	48	-	-	-	-	-
94	A6->A6	47	48	-	-	-	-	-
96	A9->A9	48	48	-	-	-	-	-
104	AA->AA	52	48	-	-	-	-	-

```
112 AB->AB 56 48 - - - - -
128 AE->AE 64 48 - - - - -
130 AF->AF 65 48 - - - - -
132 B0->B0 66 48 - - - - -
134 B1->B1 67 48 - - - - -
136 B2->B2 68 48 - - - - -
138 B3->B3 69 48 - - - - -
140 B4->B4 70 48 - - - - -
142 B5->B5 71 48 - - - - -
PEX89144 A0> quit
```

Check the back panel switch 1, 2, 6 and 7 of the system again as follows to ensure they are correct:

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked
GPIO 31 24 23 16 15 8 7 0
 B P D T T D S S S S S S D D D D D D D D D D D D D D
 t r s s 3 2 1 v p r W W W W W W W W U U U U U U U U U U
 3 O U W P P R R D S S F C U E L L L L L L L L L L L L L L
 n s t c i s R R r h s E 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0
 N R b c t s C I c a l S S 5 4 3 2 1 1 0 9 8 7 6 5 4 3 2 1 0
 / s L h D L R L S P i L e L L L L L L L L L L L L L L L
 C t d g t t t t t t t t t t t t t t t t t t t t t t t t
 o o o o o o o o o o o o o o o o o o o o o o o o o o o o Dir: [0x214]0xffffffff
 1 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 Val: [0x210]0xeffeffff
```

Now we can disable the clock to the drive slot and reset the drive:

```
[root@vlnun-100-158-IPMI-151 Linux]# sb i2c2 -d 7 -f EN_REFCLK_P0
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x6b] 0x00[02:02] EN_REFCLK_P0 Def=0x01 Cur=0x01

[root@vlnun-100-158-IPMI-151 Linux]# sb i2c2 -d 7 -f EN_REFCLK_P0 -w 0
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x6b] 0x00[02:02] EN_REFCLK_P0 Def=0x01 Cur=0x00

[root@vlnun-100-158-IPMI-151 Linux]# sb i2c2 -d 7 -f EN_REFCLK_P0
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x6b] 0x00[02:02] EN_REFCLK_P0 Def=0x01 Cur=0x00

[root@vlnun-100-158-IPMI-151 Linux]# sb i2c2 -d 7 -f PORT0_PERST_L -w 0
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x60] 0x05[05:04] PORT0_PERST_L Def=0x01 Cur=0x00

[root@vlnun-100-158-IPMI-151 Linux]# sb i2c2 -d 7 -f PORT0_PERST_L -w 1
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x60] 0x05[05:04] PORT0_PERST_L Def=0x01 Cur=0x01
```

7.2 Check Link Error Counters

We can clear the PCIe link error counters with command line as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5 -r
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 0 1h detect.Active
```

After clearing the PCIe link error counters now we can dump them as follows when some I/O running in the background and dump every other 60 seconds.

```
[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active
```



### 7.3 Clock Scope Shot from PCIe Switch Side



- The channel “C4” is the reference clock obtained from the PCIe switch side, and the measurement shows it is around 100 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 100 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 100 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts and displays the instantaneous frequency variation of the clock as a function of time.
- From the scope shot above we can tell the clock in switch side has no SSC turned on.

## 7.4 Clock Scope Shot from Drive Side



- The channel “C4” is the reference clock obtained from drive side, and the measurement shows it is around 100 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 100 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 100 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts/displays the instantaneous frequency variation of the clock as a function of time.
- From the scope shot above we can tell the clock in drive side has no SSC turned on.
- Wait a minute, we already turned off the refclk to this slot, why do we still see the clock from the test point ion the drive side? Turns out it is related to the **sb\_i2c2** command used to turn off the refclk to the drive which will cut the refclk on the riser, not the clock wired out in the test points on the drive side.

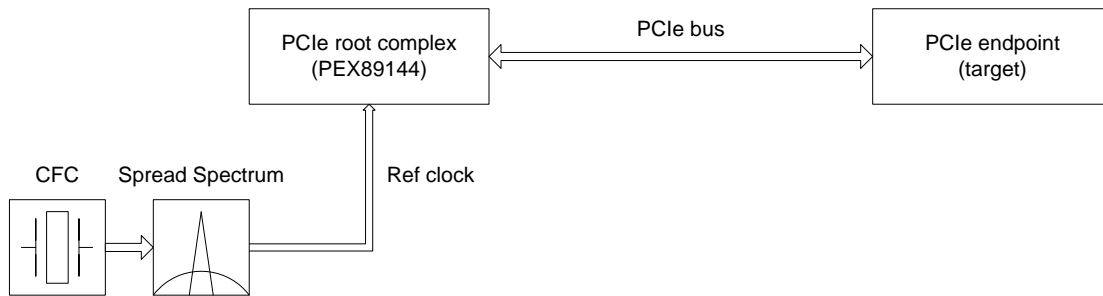
## 7.5 Skip Ordered-Sets from Bus Trace

Trace View / PCIe Revision 4.0									
Packet	R←	8.0	SKIP	SKIP Symbols	END	P	LFSR Symbols	Idle	Time Stamp
0		x4		AAAAAAAAAA...	E1	1	** ** *	554.750 ns	0000 . 000 000 000 000 s
1		x4		AAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 000 571 000 s
2		x4		AAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 001 142 000 s
3		x4		AAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 001 713 000 s
4		x4		AAAAAAAAAA...	E1	*	** ** *	556.750 ns	0000 . 000 002 284 000 s
5		x4		AAAAAAAAAA...	E1	1	** ** *	554.750 ns	0000 . 000 002 857 000 s
6		x4		AAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 003 428 000 s
7		x4		AAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 003 999 000 s
8		x4		AAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 004 570 000 s
9		x4		AAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 005 141 000 s
10		x4		AAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
11		x4		AAAAAAAAAA...	E1	*	** ** *	447.000 ns	0000 . 000 005 712 000 s
12		x4		AAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
13		x4		AAAAAAAAAA...	E1	*	** ** *	124.000 ns	0000 . 000 006 159 000 s
14		x4		AAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
15		x4		AAAAAAAAAA...	E1	*	** ** *	396.000 ns	0000 . 000 006 283 000 s
16		x4		AAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
17		x4		AAAAAAAAAA...	E1	*	** ** *	175.000 ns	0000 . 000 006 679 000 s
18		x4		AAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
19		x4		AAAAAAAAAA...	E1	*	** ** *	345.000 ns	0000 . 000 006 854 000 s
20		x4		AAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
21		x4		AAAAAAAAAA...	E1	*	** ** *	226.000 ns	0000 . 000 007 199 000 s
22		x4		AAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
23		x4		AAAAAAAAAA...	E1	*	** ** *	294.000 ns	0000 . 000 007 425 000 s

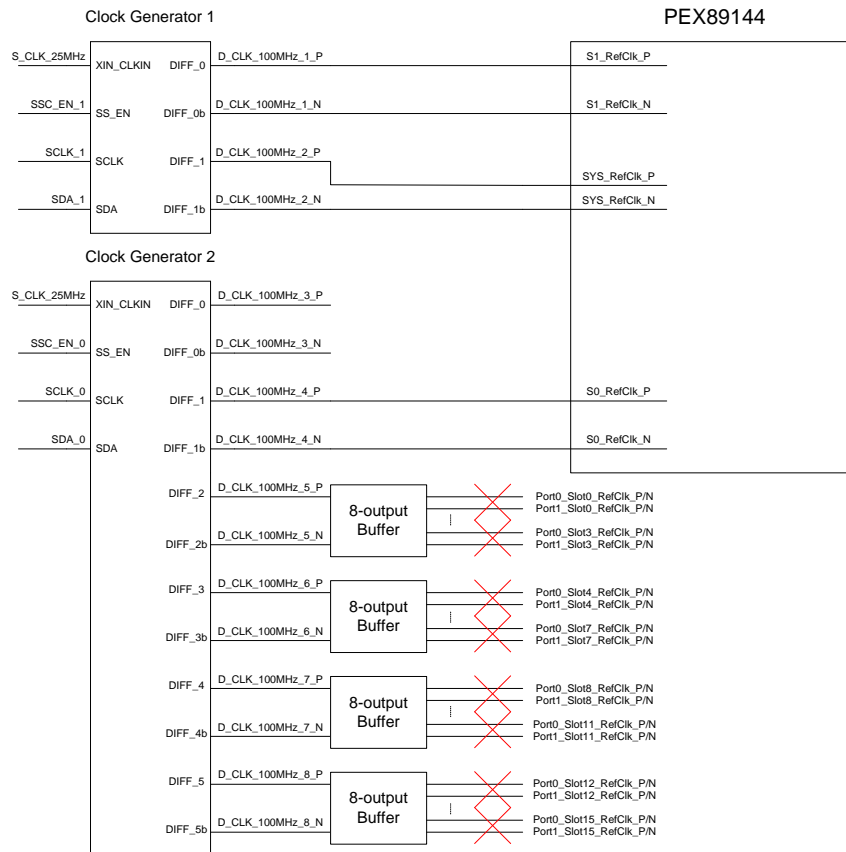
You can see the drive sent out a skip ordered-set around every 571 ns, which is much more frequent than situation with external refclk. With external refclk the drive sent out a skip ordered-set around every 6  $\mu$ s. When the drive didn't detect an external reference clock, it will use the internal refclk with SSC, so the drive will send out extra skip ordered-sets.

## 8. RefClk with SSC for PEX Only (No RefClk to Target Drive, Dynamic SRIS)

The RefClk with SSC for PEX only (no RefClk to the target drive) can be described in following chart:



This PCIe clocking mode can be easily achieved in the SANBlaze Gen5 system as follows:



The SSC\_EN\_0 and SSC\_EN\_1 for both Clock Generators are enabled by the Back Panel Switches (SW6 is down and SW7 is up). The S0\_RefClk\_P/N and S1\_RefClk\_P/N are ignored in this case because they are for SSC isolation only. The EEPROM image on the SANBlaze Gen5 system specified all ports SerDes are clocked with SYS\_RefClk\_P/N. The PEX89144 got separate RefClk with SSC from Clock Generator 1 with the MUX selection. The RefClk for all targets can be disabled through sb\_i2c2 commands below. For example, the following commands can be used to disable/enable RefClk for slot 7:

```
sb_i2c2 -n 1 -d 7 -f EN_REFCLK_PO -w 0      # disable RefClk for slot 7
sb_i2c2 -n 1 -d 7 -f EN_REFCLK_PO -w 1      # enable RefClk for slot 7
```

You also can tape the RefClk on the PCIe riser so no RefClk for targets. After disabling the refclk you may need to reset the drive to bring it up on the bus.

With this clocking mode I sample the link error counter registers every 1 minute and no receiver errors, TLP errors, DLLP errors or Recovery events happened.

## 8.1 Configure Clocking Mode as Dynamic SRIS

We can configure the SANBlaze Gen5 system to use this dynamic SRIS as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW1EESEL0
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: ef9fffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 clear SW2EESEL1
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
INFO: Clearing bit 17 at addr 0x210
00000210: ef9fffff
00000210: ef9dffff
00000210: ef9dffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 set SW6SRIS
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
00000210: ffffffff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: ffffffff

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 clear SW7SSCL
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
INFO: Clearing bit 22 at addr 0x210
00000210: efbdfbff
00000210: efbdfbff
00000210: efbdfbff
INFO: Storing current GPIO in non-volatile eeprom
0000020c: 00000000

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked
GPIO | 31 | 24 | 23 | 16 | 15 | 8 | 7 | 0 |
-----|-----|-----|-----|-----|-----|-----|-----|-----|
B P D D T T D S S S S S S D D D D D D D D D D D D D D D D
i w i i P P P r W W W W W W W W u u u u u u u u u u u u u u
t r s s 3 2 1 v 8 7 6 5 4 3 2 1 a a a a a a a a a a a a a a a a
3 O U W | | | R D S S S F C U E E 1 1 1 1 1 1 1 1 1 1 1 1 1 1
1 n s t | | | s i S R r h s E E 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0
N R b c | | | t s C I c a l S S 5 4 3 2 1 0 9 8 7 6 5 4 3 2 1 0
/ s L h | | | L R L S P i n l e e L L L L L L L L L L L L L L
C t D | | | s t w n l l | | | | | | | | | | | | | | | |
g | | | t r l 0 | | | | | | | | | | | | | | | | | | | | |
o o o o o o o o o o o o o o o o o o o o o o o o o o o o Dir: [0x214]0xffffffff
1 1 1 0 1 1 1 1 1 0 1 1 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 Val: [0x210]0xfbdffff
```

After the configuration above, we need to shut down the system then power it up as follows:

```
[root@vlnun-100-158-IPMI-151 Linux]# shutdown now
```

When the full system is powered off, wait for 10 seconds then power it up. After boot up you can check the clocking mode as follows and ensure it is “SRIS” because all separate clock modes will show up as “SRIS”, no matter if it has SSC signal enabled or not.

```
[root@vlnun-100-158-IPMI-151 vlnun]# mount 192.168.1.80:/home/sanblaze.local /mnt
[root@vlnun-100-158-IPMI-151 vlnun]# cd /mnt/shynes/Broadcom/tools/G4XTOOLS/Linux/
[root@vlnun-100-158-IPMI-151 Linux]# ./g4Xdiagnostics.x86_64 -sdb /dev/ttyACM3
```

-----  
g4Xdiagnostics v0.0.0.8 - Broadcom Inc. (c) 2021 (Bld-37.46.11.64.6.0)

```

-----
PEX89144 A0> port
Port: Port number in decimal          Stn: Station number in decimal
Type: Port type Down-Downstream, Host-Upstream, Fab-Fabric Mode, Mgmt-Management
MRR: Max Read Request Size in Bytes  MPSS: Max Payload Size Supported in Bytes
MPS: Max Payload Size in Bytes       Link: Link Width by Negotiated/Maximum
Speed: Link Speed by Negotiated/Maximum  ClkMode: Clock Mode
PBus: Primary Bus number in hex,      Stat: Link Status

```

Stn	Port	Type	MRR	MPS	MPSS	LinkSpeed	LinkWidth	ClkMode	PBus	Stat
0	0	Down	512	256	256	Gen5 / Gen5	x16 / x2	SRIS	88	Up
0	2	Fab	512	256	256	Gen5 / Gen5	x16 / x2	SRIS	88	Up
0	4	Down	512	256	256	Gen5 / Gen5	x16 / x2	SRIS	88	Up
0	6	Fab	512	256	256	Gen5 / Gen5	x16 / x2	SRIS	88	Up
0	8	Fab	512	256	256	Gen5 / Gen5	x16 / x16	SRIS	88	Up
1	24	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	26	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	28	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
1	30	Down	128	128	2048	- / Gen5	- / x2	SRIS	88	Down
2	32	Down	128	128	2048	- / Gen5	- / x16	SRIS	9A	Down
3	48	Host	128	128	2048	Gen4 / Gen5	x16 / x16	Common	85	Up
4	64	Down	128	128	2048	- / Gen5	- / x16	SRIS	9D	Down
5	80	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	82	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	84	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	86	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	88	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	90	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	92	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
5	94	Down	128	128	2048	- / Gen5	- / x2	SRIS	9D	Down
6	96	Down	128	128	2048	- / Gen5	- / x8	SRIS	A8	Down
6	104	Down	128	128	2048	Gen2 / Gen5	x1 / x8	SSC	A8	Up
7	112	Down	128	128	2048	- / Gen5	- / x16	SRIS	A8	Down
8	128	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	130	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	132	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	134	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	136	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	138	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	140	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down
8	142	Down	128	128	2048	- / Gen5	- / x2	SRIS	AD	Down

```

Host Ports:
BusPrim: Primary Bus number in hex,   BusSec: Secondary Bus number in hex
BusSub: Subordinate Bus number in hex, Assigned DS Ports: In decimal

```

Port	BusPrim	BusSec	BusSub	Assigned DS Ports
48	85	86	B5	0, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 32, 64, 80, 82, 84, 86, 88, 90, 92, 94, 96, 104, 112, 128, 130, 132, 134, 136, 138, 140, 142,

```

Downstream Ports:
Port: Port number in decimal,          DSBuses: Downstream bus range
PhysSlot: Chassis Phsyiscal slot in decimal, HostPort : Associated Host Port in decimal
VID: Vendor ID,                        DID : Device ID
SVD: Sub-vendor ID,                    SID : Sub-device ID
Attached Device Details: All values are in hex

```

Port	DSBuses	PhysSlot	HostPort	Attached Device Details					
				VID	DID	SVD	SID	ClassCode	SubClassCode
0	89->89	0	48	-	-	-	-	-	-
4	8B->8B	2	48	-	-	-	-	-	-
6	8C->8C	3	48	-	-	-	-	-	-
8	8D->8D	4	48	-	-	-	-	-	-
10	8E->8E	5	48	-	-	-	-	-	-
12	8F->8F	6	48	-	-	-	-	-	-
14	90->90	7	48	-	-	-	-	-	-
16	91->91	8	48	-	-	-	-	-	-
18	92->92	9	48	-	-	-	-	-	-
20	93->93	10	48	-	-	-	-	-	-
22	94->94	11	48	-	-	-	-	-	-
24	95->95	12	48	-	-	-	-	-	-
26	96->96	13	48	-	-	-	-	-	-
28	97->97	14	48	-	-	-	-	-	-
30	98->98	15	48	-	-	-	-	-	-
32	9B->9B	16	48	-	-	-	-	-	-
64	9E->9E	32	48	-	-	-	-	-	-
80	9F->9F	40	48	-	-	-	-	-	-
82	A0->A0	41	48	-	-	-	-	-	-
84	A1->A1	42	48	-	-	-	-	-	-
86	A2->A2	43	48	-	-	-	-	-	-
88	A3->A3	44	48	-	-	-	-	-	-
90	A4->A4	45	48	-	-	-	-	-	-
92	A5->A5	46	48	-	-	-	-	-	-
94	A6->A6	47	48	-	-	-	-	-	-
96	A9->A9	48	48	-	-	-	-	-	-
104	AA->AA	52	48	-	-	-	-	-	-
112	AB->AB	56	48	-	-	-	-	-	-
128	AE->AE	64	48	-	-	-	-	-	-
130	AF->AF	65	48	-	-	-	-	-	-
132	B0->B0	66	48	-	-	-	-	-	-
134	B1->B1	67	48	-	-	-	-	-	-

```

136 B2->B2 68 48 - - - - - 
138 B3->B3 69 48 - - - - - 
140 B4->B4 70 48 - - - - - 
142 B5->B5 71 48 - - - - - 
PEX89144 A0> quit

```

Check the back panel switch 1, 2, 6 and 7 of the system again as follows to ensure they are correct:

```

[root@vlnun-100-158-IPMI-151 Linux]# iriser -d -3 show gpio
INFO: Using activefile=/etc/iRiser/1/34/active.cfg
INFO: Use activefile=/etc/iRiser/1/34/active.cfg for cfgfile
INFO: Called to read configuration file /etc/iRiser/1/34/active.cfg activate = 0
Action: Sequence 512 address 0x200
Direction: I/O (UC) user writable signal, i/o (LC) signal locked
GPIO|31|24|23|16|15|8|7|0|
-----|-----|-----|-----|-----|-----|-----|-----|
B P D D T T D S S S S S S S D D D D D D D D D D D D D D D D
i w i i P P P R W W W W W W W U U U U U U U U U U U U U U U
t r s s s 3 2 1 v 8 7 6 5 4 3 2 1 a a a a a a a a a a a a a a
3 O U W R D S S F C U E E 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
1 n s t s i S R r h s E E 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
N R b c t s C I C a l S S 5 4 3 2 1 1 0 9 8 7 6 5 4 3 2 1 0
/ s L h d L R L S P i L e e L L L L L L L L L L L L L L L
C t D s t w n l l 1 1 0
-----|-----|-----|-----|-----|-----|-----|-----|
o o o o o o o o o o o o o o o o o o o o o o o o o o Dir:[0x214]0xffffffff
1 1 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 Val:[0x210]0xefbdfbff

```

Now we can disable the clock to the drive slot and reset the drive:

```

[root@vlnun-100-158-IPMI-151 Linux]# sb_i2c2 -d 7 -f EN_REFCLK_P0
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x6b] 0x00[02:02] EN_REFCLK_P0 Def=0x01 Cur=0x01

[root@vlnun-100-158-IPMI-151 Linux]# sb_i2c2 -d 7 -f EN_REFCLK_P0 -w 0
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x6b] 0x00[02:02] EN_REFCLK_P0 Def=0x01 Cur=0x00

[root@vlnun-100-158-IPMI-151 Linux]# sb_i2c2 -d 7 -f EN_REFCLK_P0
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x6b] 0x00[02:02] EN_REFCLK_P0 Def=0x01 Cur=0x00

[root@vlnun-100-158-IPMI-151 Linux]# sb_i2c2 -d 7 -f PORT0_PERST_L -w 0
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x60] 0x05[05:04] PORT0_PERST_L Def=0x01 Cur=0x00

[root@vlnun-100-158-IPMI-151 Linux]# sb_i2c2 -d 7 -f PORT0_PERST_L -w 1
INFO: System 01[07] SBExpress-RM5 SN=950A2050001 Rev=R03 i2c_main=i2c-7 i2c_priv=i2c-5 SDB=/dev/ttyACM3 VLUN=0
INFO: 07[0x60] 0x05[05:04] PORT0_PERST_L Def=0x01 Cur=0x01

```

## 8.2 Check Link Error Counters

We can clear the PCIe link error counters using the command line as follows:

```

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5 -r
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

```

After clearing the PCIe link error counters now we can dump them as follows when some I/O running in the background and dump every other 60 SECONDS.

```

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

[root@vlnun-100-158-IPMI-151 Linux]# sb_sdb -n 1 -d 1 -l -5
Recovery counter is under development and is for internal use only
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR) 0x0BC4(RCVY) 0x0BB0(LTSSM) State.Substate
1 20(x14) 93:00.0 0 0 0 1 0 0 0 0 0 0 0 0 1h detect.Active

```

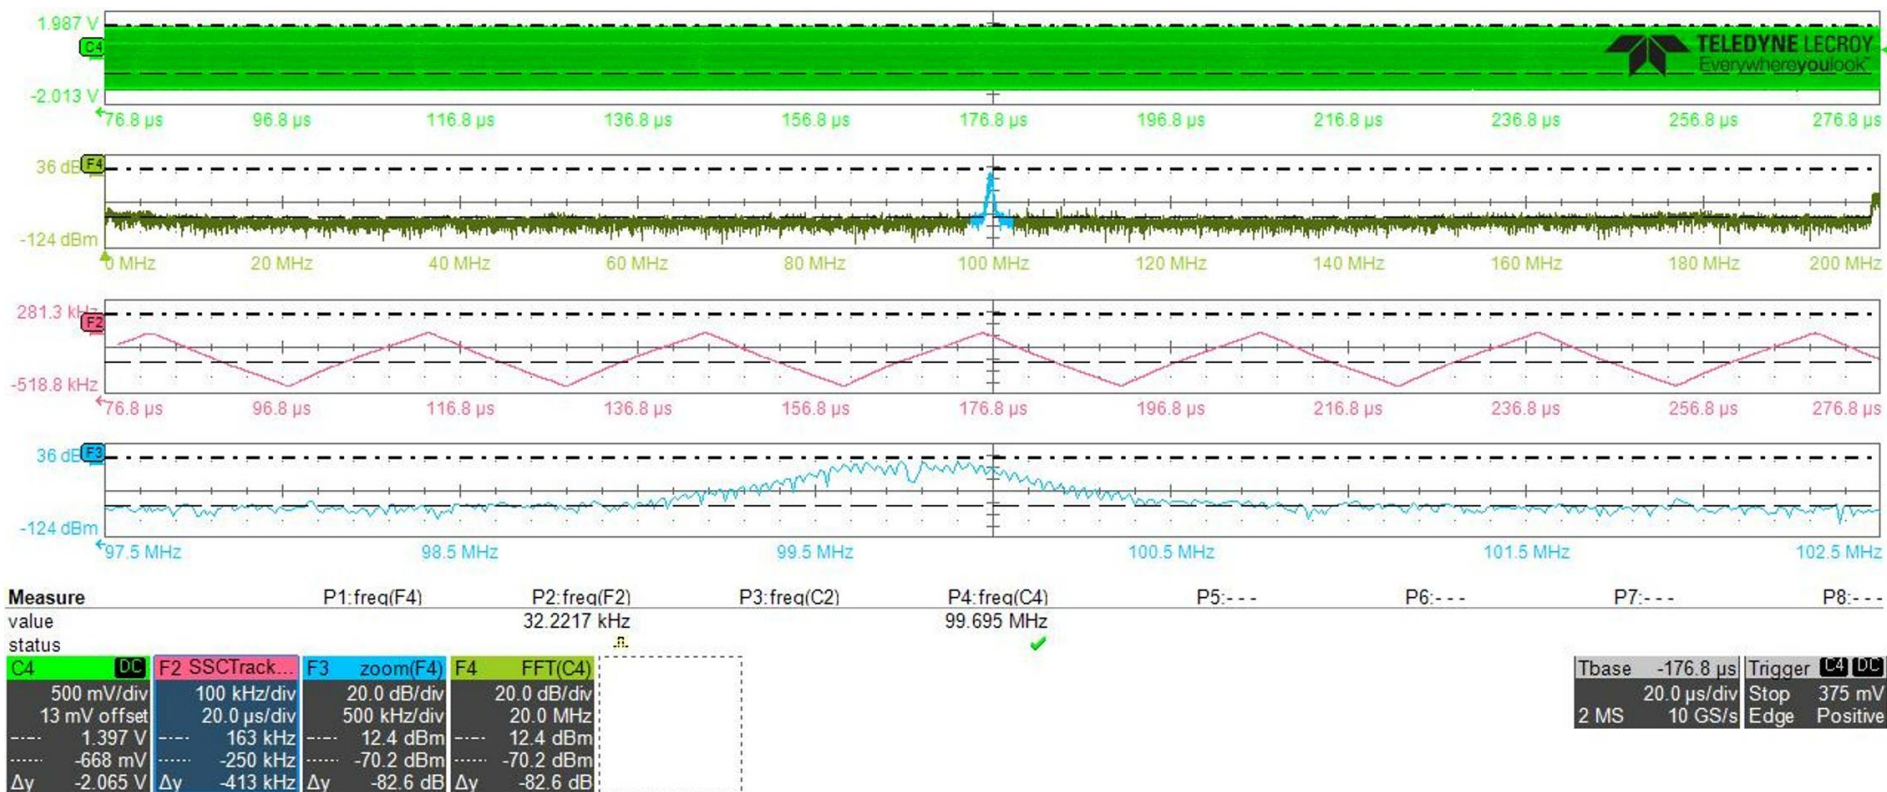
### 8.3 Clock Scope Shot from PCIe Switch Side



- The channel “C4” is the reference clock obtained from PCIe switch side, and the measurement shows it is around 100.6 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 100.6 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 100.6 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts/displays the instantaneous frequency variation of the clock as a function of time.
- From the scope shot above we can tell the clock in switch side has no SSC turned on, which is not expected. Turns out when in both static SRIS mode and dynamic SRIS mode, the PEX89144 chip will not link so we can only provide refclk without SSC, and the PEX89144 flash image will use the internal PLL to turn on SSC for the PCIe switch. There’s no way for us to check the clock SSC which is generated inside the chip.



## 8.4 Clock Scope Shot from Drive Side



- The channel “C4” is the reference clock obtained from drive side, and the measurement shows it is around 99.7 MHz.
- The channel “F4” is the FFT() result of channel “C4” and you can see the spectrum, and the max spectral amplitude is around 99.7 MHz.
- The channel “F3” is the zoom of channel “F4” between 97.5 MHz and 102.5 MHz, and confirmed the max spectral amplitude is around 99.7 MHz.
- The channel “F2” is SSCTrack() result of channel “C4”, which extracts and displays the instantaneous frequency variation of the clock as a function of time. You can see the SSC modulation frequency is about 32.2 kHz which matches clock generator specification (Modulation Frequency: Min = 30 kHz, Typ = 31.5 kHz and Max = 33 kHz).
- From the scope shot above we can tell the clock in drive side has SSC turned on, and the modulation technique is down-spread.
- Wait a minute, we already turned off the refclk to this slot, why do we still see the clock from the test point on the drive side? Turns out it is related to the `sb_i2c2` command used to turn off the refclk to the drive, which will cut the refclk on the riser, not the clock wired out in the test points at drive side.

## 8.5 Skip Ordered-Sets from Bus Trace

Teledyne LeCroy PCIe Protocol Analysis - BETA - [C:\Users\haiyan\_lin1\Documents\PCIe\_traces\T48\_U.2\_dynamic\_SRIS\_2.pex]

File Setup Record Generate Report Search View Tools Window Help

Trace View / PCIe Revision 4.0


Packet	R→	8.0	SKIP	SKIP Symbols	END	P	LFSR Symbols	Idle	Time Stamp
0	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 000 000 000 s
1	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 000 571 000 s
2	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	1	** ** *	554.750 ns	0000 . 000 001 142 000 s
3	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	556.750 ns	0000 . 000 001 713 000 s
4	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 002 286 000 s
5	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 002 857 000 s
6	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	1	** ** *	554.750 ns	0000 . 000 003 428 000 s
7	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 003 999 000 s
8	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 004 570 000 s
9	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	1	** ** *	554.750 ns	0000 . 000 005 141 000 s
10	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 005 712 000 s
11	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	554.750 ns	0000 . 000 006 283 000 s
12	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
13	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	32.000 ns	0000 . 000 006 854 000 s
14	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	1	** ** *	Idle	Time Stamp
15	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	503.750 ns	0000 . 000 006 886 000 s
16	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	1	** ** *	Time Delta	Time Stamp
17	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	19.000 ns	0000 . 000 007 406 000 s
18	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
19	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	501.000 ns	0000 . 000 007 425 000 s
20	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	Time Delta	Time Stamp
21	R→	8.0	SKIP	AAAAAAAAAAAA...	E1	*	** ** *	70.000 ns	0000 . 000 007 926 000 s

You can see the drive sent out a skip ordered-set around every 571 ns, which is much more frequent than the situation with external refclk. With external refclk the drive sent out a skip ordered-set around every 6  $\mu$ s. When the drive didn't detect an external reference clock, it will use internal refclk with SSC, so the drive will send out extra skip ordered-sets.

## 9. Common Clock Without SSC vs. Common Clock With SSC

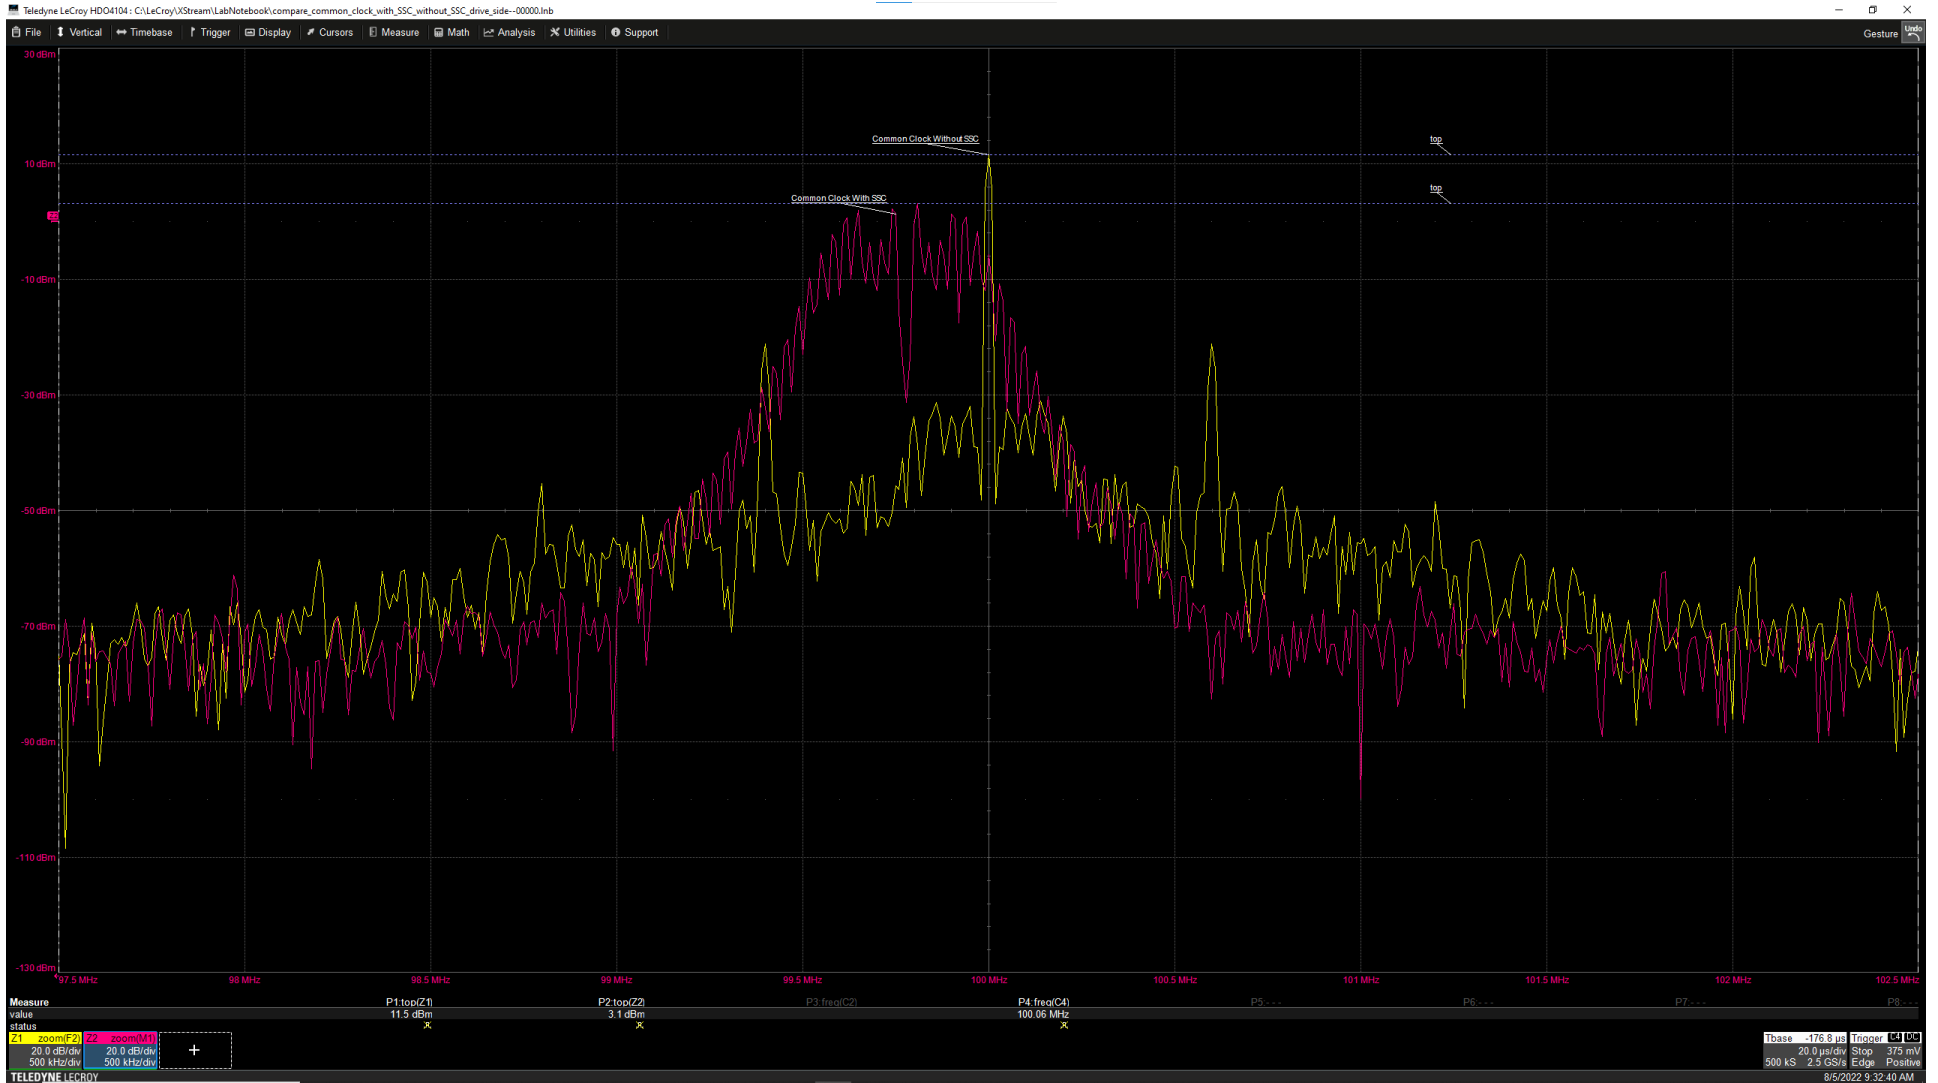
The only difference between common clock without SSC and common clock with SSC is just one has SSC disabled and the other has SSC enabled.

### 9.1 Drive Side

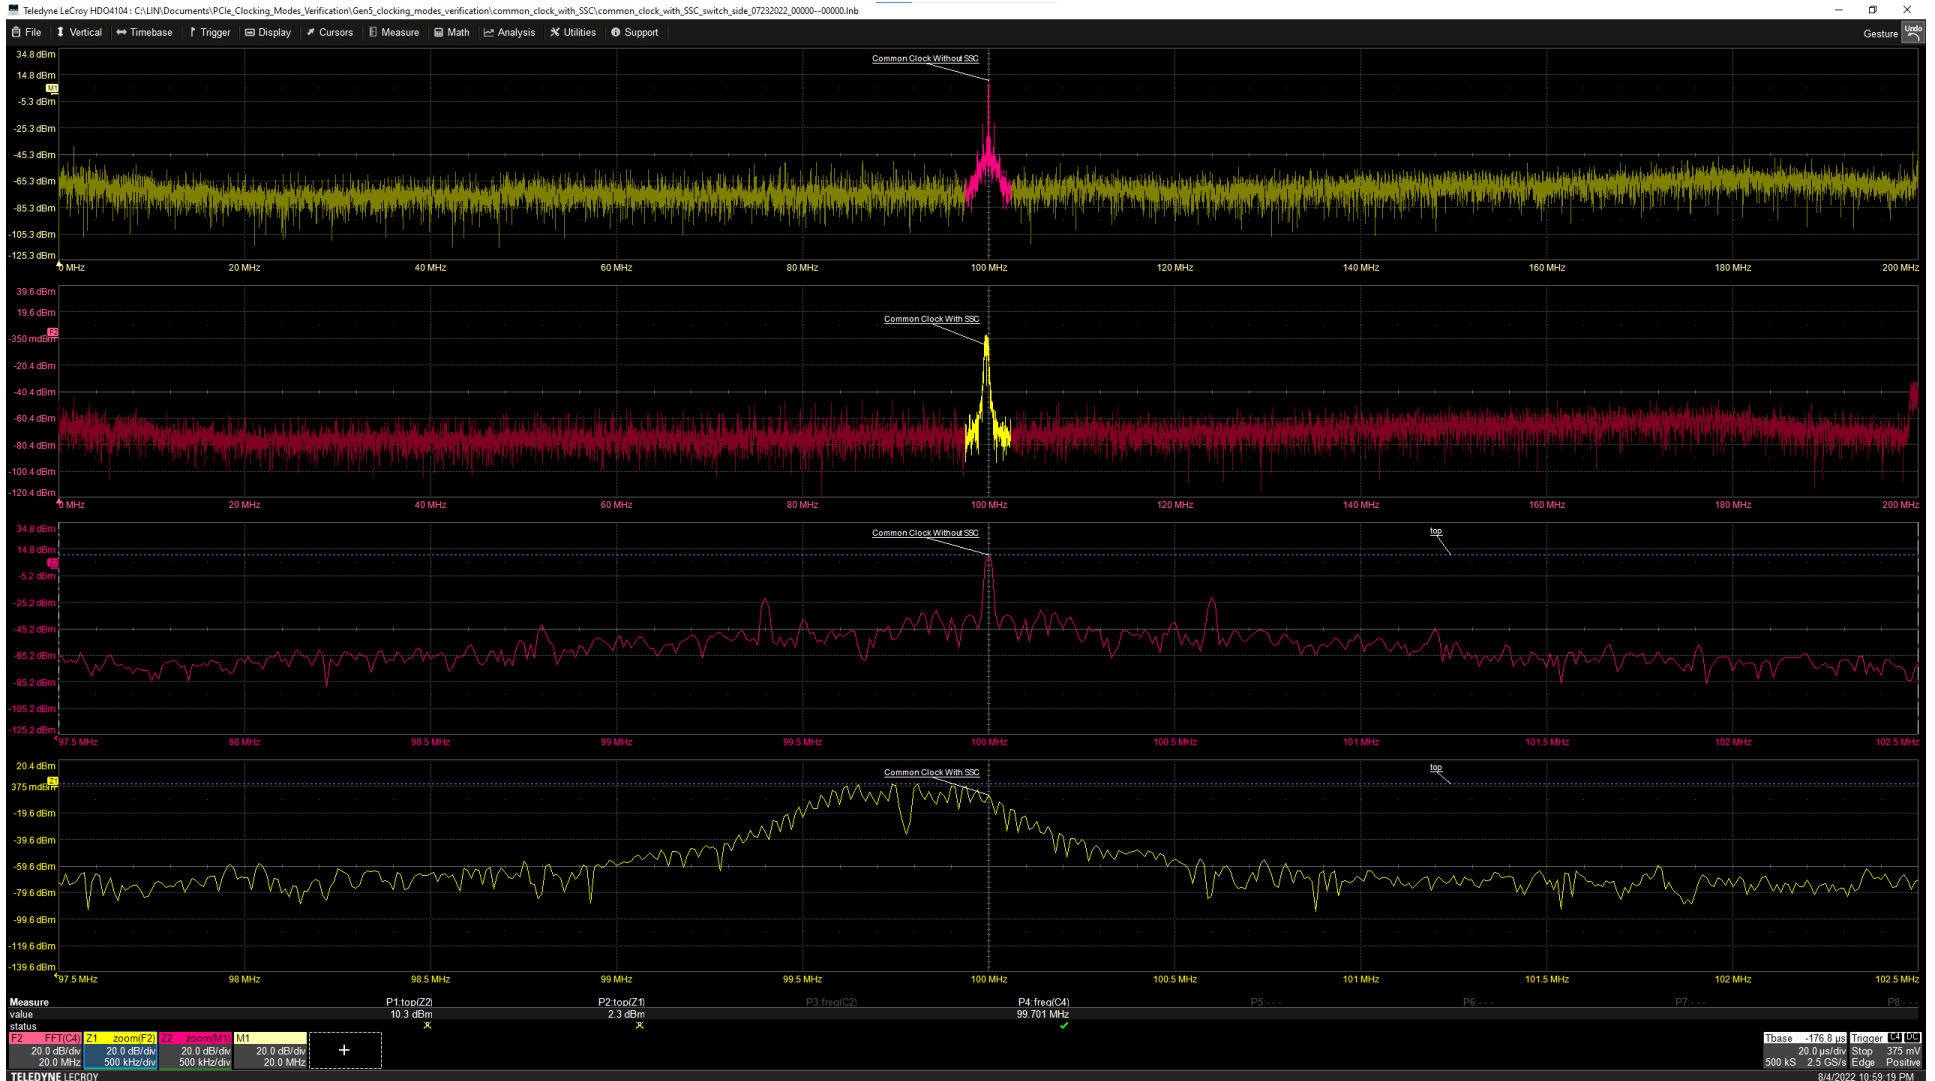


The common clock with SSC is down-spread and shows 8.4 dBm power spectrum amplitude drop from the common clock without SSC for EMI reduction.

Put in the same chart and it will show the difference better as follows:




## 9.2 Switch Side



The common clock with SSC is down-spread and shows 8.0 dBm power spectrum amplitude drop from the common clock without SSC for EMI reduction.

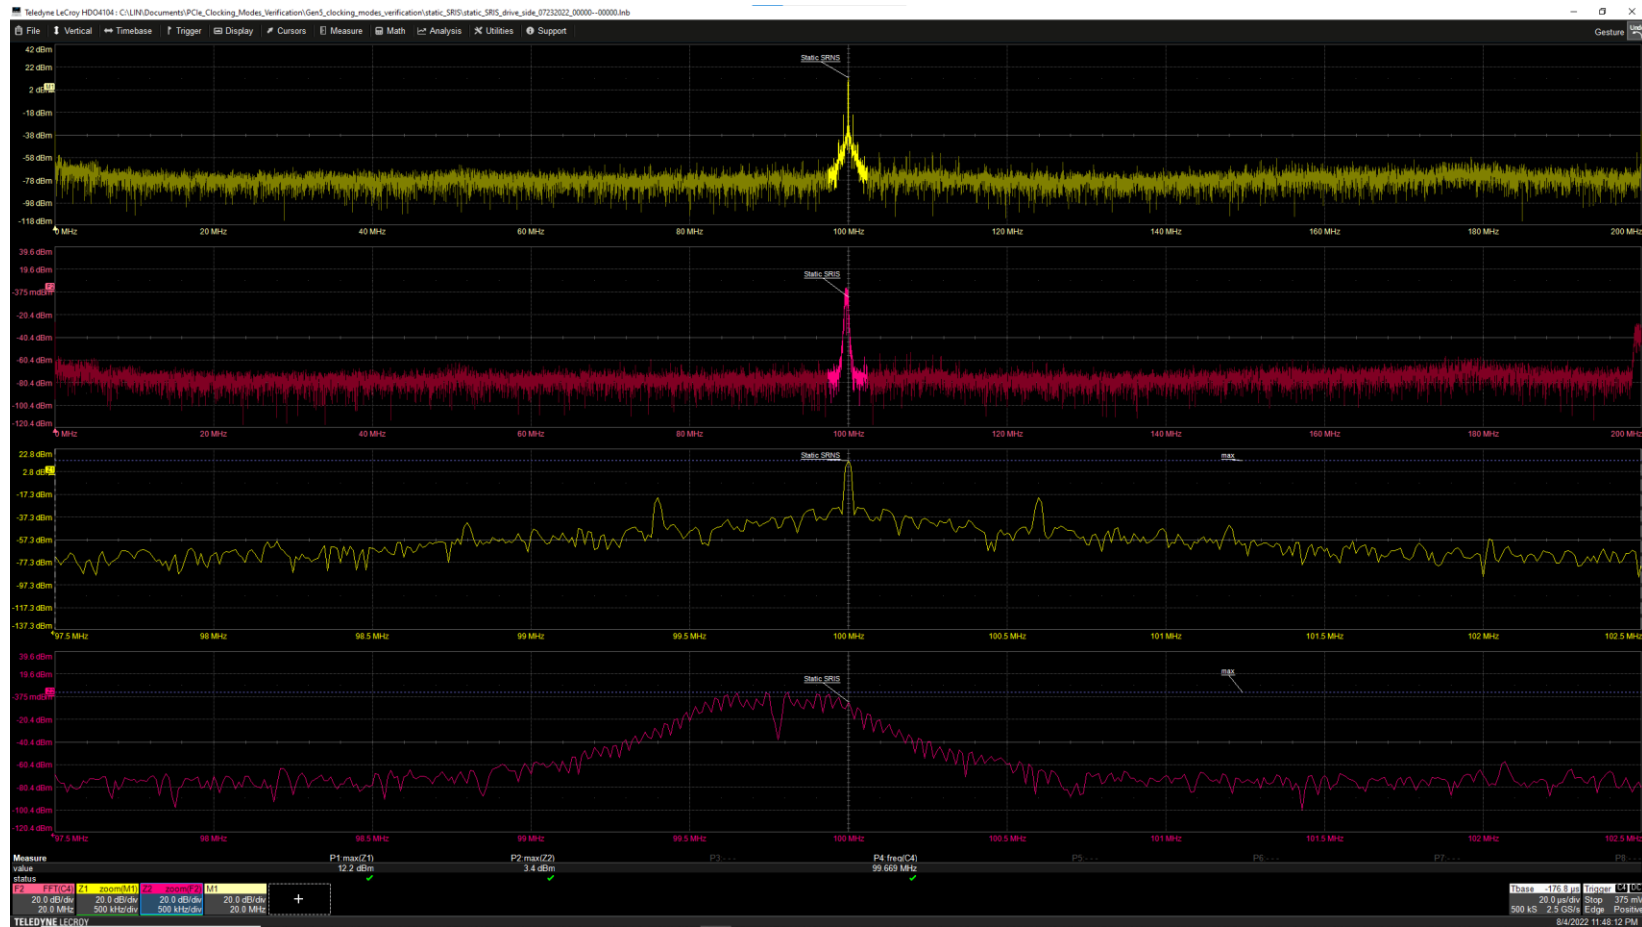
Put in the same chart and it will show the difference better as follows:



## 10. Static SRNS vs. Static SRIS


The only difference between Static SRNS and Static SRIS is just one has SSC disabled and the other has SSC enabled.

### 10.1 Drive Side




The Static SRIS is down-spread and shows 8.8 dBm power spectrum amplitude drop from the Static SRNS for EMI reduction.

Put in the same chart we can better show the difference:






## 10.2 Switch Side



Both Static SRNS and Static SRIS on the switch side are using the clock without SSC turned on, so you can see the power spectrum amplitude is almost the same for both clocks. But inside the switch it is using internal PLL to enable SSC for Static SRIS mode. Please see page 33 for detail.

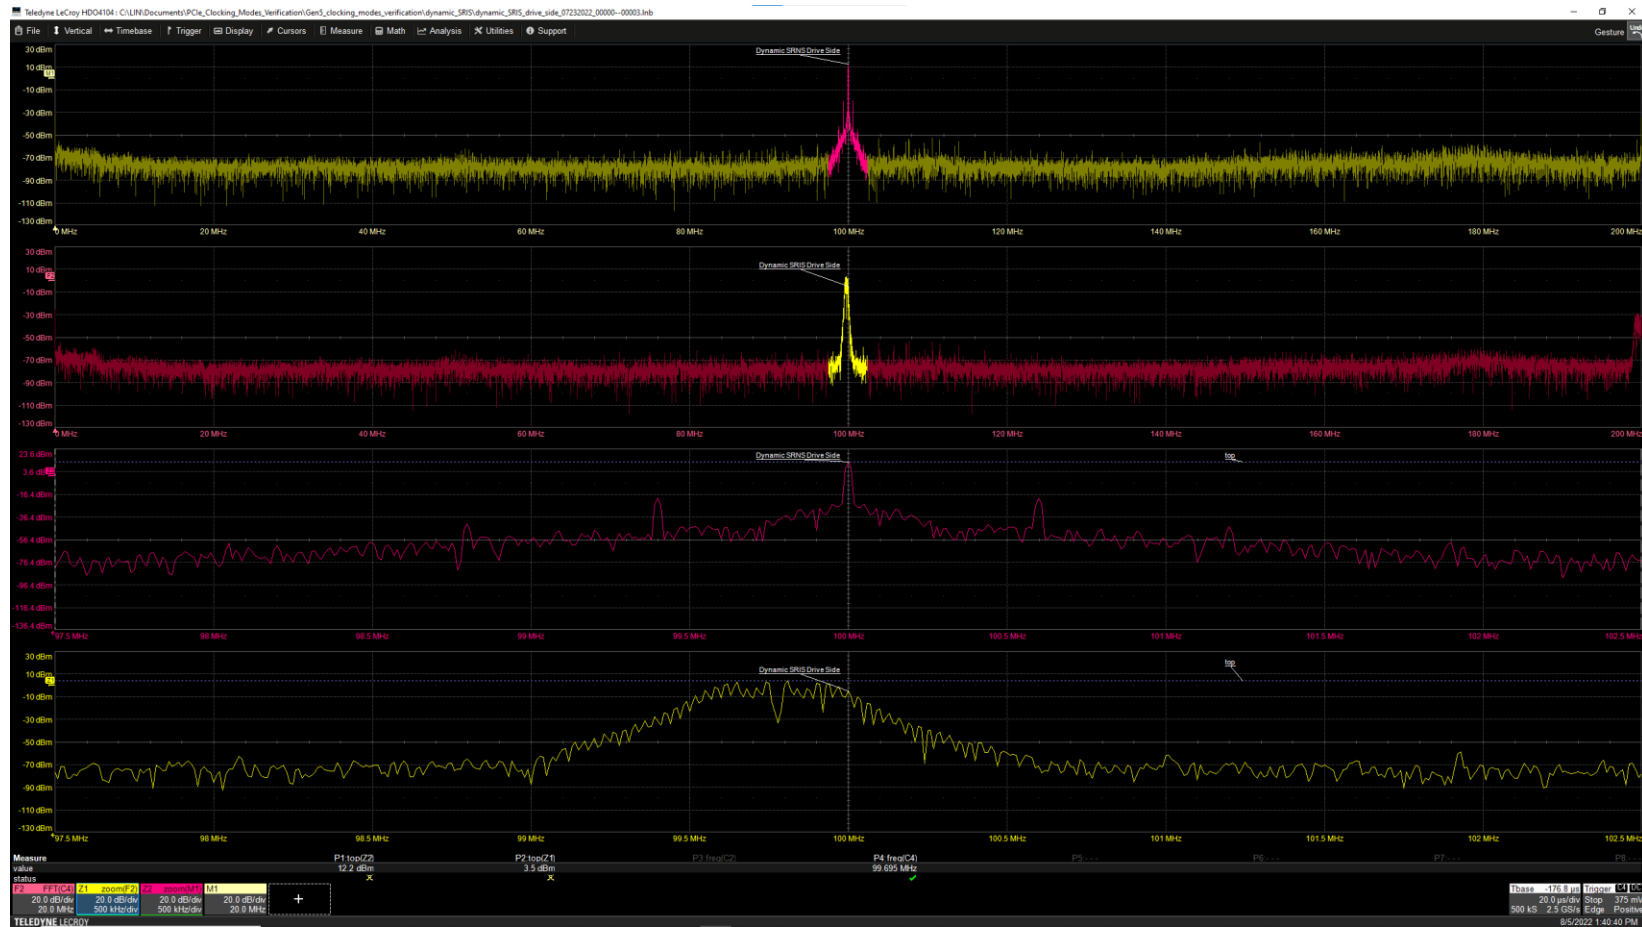
Here are both clocks' FFT in the same chart as follows:




# 11. Dynamic SRNS vs. Dynamic SRIS

The only difference between Dynamic SRNS and Dynamic SRIS is just one has SSC disabled and the other has SSC enabled.

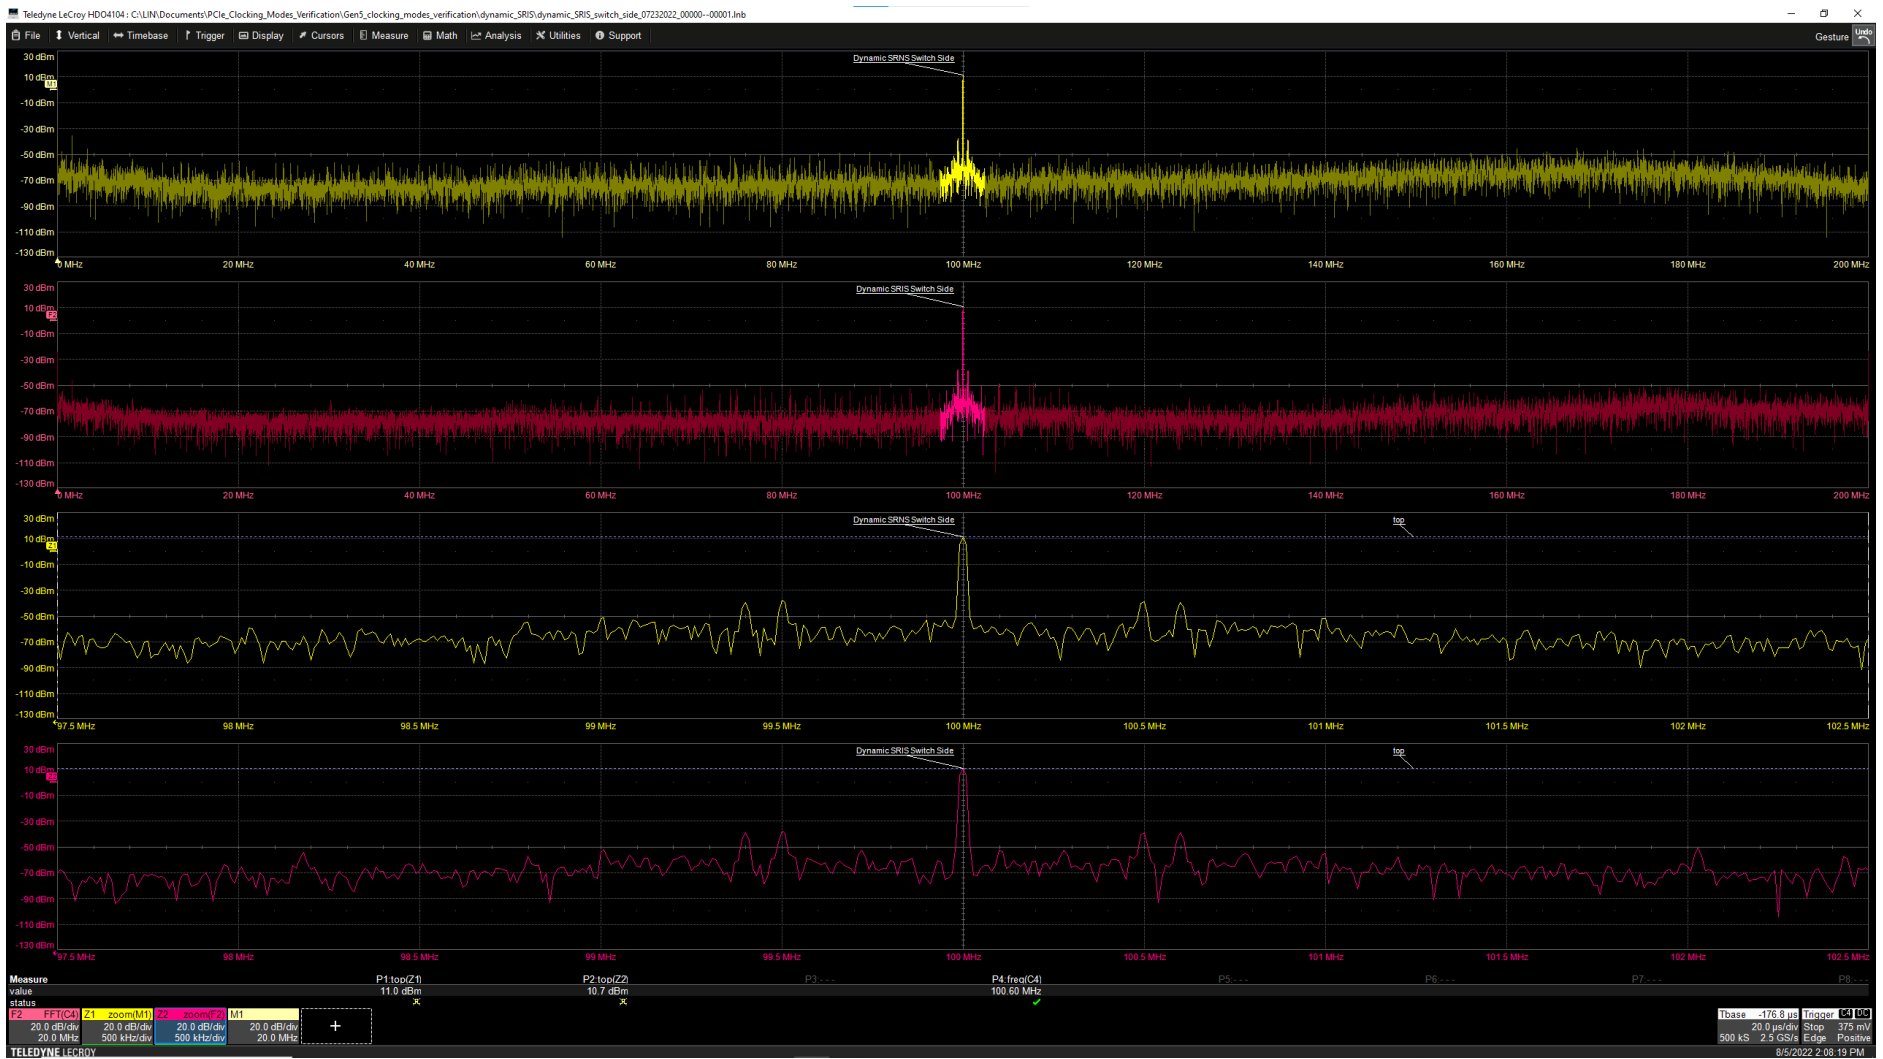
## 11.1 Drive Side



The Dynamic SRIS is down-spread and shows 8.7 dBm power spectrum amplitude drop from the Dynamic SRNS for EMI reduction. The Dynamic SRIS has refclk turned off on the riser, but still shows up at the test points on the drive side. Putting both clocks' FFT in the same chart for better comparison:




## 11.2 Switch Side



Both Dynamic SRNS and Dynamic SRIS on the switch side are using the clock without SSC turned on, so you can see the power spectrum amplitude is almost the same for both clocks. But inside the switch it is using the internal PLL to enable SSC for Dynamic SRIS mode. Please see page 47 for detail.

Put both clocks' FFT in the same chart:



## 12. Summary

Based on the back panel switches readback, PCIe switch ports clock mode read back, clock waveforms from the drive side, clock waveforms from the switch side and Skip Ordered-Sets from bus trace, we confirmed the following 6 clocking modes are working as expected with the SANBlaze SBExpress NVMe Gen5 test system.

- Common clocking mode without SSC
- Common clocking mode with SSC
- Separate reference clock without SSC (Static SRNS)
- Separate reference clock independent SSC (Static SRIS)
- RefClk without SSC for PEX only (no RefClk to the target drive, Dynamic SRNS)
- RefClk with SSC for PEX only (no RefClk with SSC to the target drive, Dynamic SRIS)

Just a couple of items below need to be pointed out:

1. In both static SRIS mode and dynamic SRIS mode, the PEX89144 chip will not link if we turn on SSC. So we can only provide refclk without SSC to the PCIe switch, and enable the PEX89144 flash image to use internal PLL to turn on SSC for the PCIe switch. But there's no way for us to check the clock SSC which is generated inside the chip.
2. Issuing the **sb\_i2c2** command to turn off the refclk to the drive will cut the refclk on the riser, not the clock in the riser slot such as the test points on the drive side in page 7.

For more information, please contact SANBlaze [here](#).