

Testing NVMe L1.1 and L1.2 Power Sub-States with SBExpress-RM5



Verifying CLKREQ# Control for NVMe
Entering and Exiting L1.1 and L1.2 Ultra Low
Power Sub-States

Author:

Vince Asbridge
CEO and Director of
Software Engineering,
SANBlaze

Table of Contents

Introduction	1
Primary PCIe Power States.....	1
Defining L1.1 and L1.2 Sub-States.....	2
Configuring the PCIe LTR Registers for the L1 Substate Test Script.....	3
Controlling CLKREQ#	4
CLKREQ# Signal	4
Modes of CLKREQ# L1.1 L1.2 Testing	5
Scripts for L1.1 L1.2 Sub-state Testing.....	5
CLKREQ# Mode Definition	5
Staging the Mode 1 CLKREQ# Tests.....	6
Selecting "Test Manager" Testing	6
Monitoring the NVMe Sideband Signals	8
SetFeatures NVMe Support.....	8
Mode1 Testing Results	9
Staging the Mode 2 CLKREQ# Tests.....	9
Requirements for Mode2 Testing	9
Execution of the Mode2 CLKREQ# Tests	10
Interpreting and Publishing Test Results	11
Under the Covers for CLKREQ# Testing	11
Initiating Sub-States from the Command Line Interface (CLI).....	11
Use sb_sdb to Check Link State	11
Using SetFeatures to Enable Lowest Power State	12
Turn off the System Watchdog	12
Turn off the SMBus to the Drive	12
Use ASPM to Enable Low Power States	13
Enabling L1.1 or L1.2	13
Additional Capabilities of CLI Commands.....	15
Examples of sb_i2c2 Commands	15
Read and Write CLKREQ_L.....	15
Conclusion.....	16
Determining the Capability of your RM5 System	16
FAQs:	17

Introduction

SANBlaze SBExpress and the *Certified by SANBlaze* test methodology have greatly simplified qualification of PCIe NVMe devices by providing a simple means of creating complex test suites, validating specification compliance, data integrity, power and reset testing, and MI compliance.

With Version 10.5 of the *Certified by SANBlaze* test suite, SANBlaze is introducing a suite of tests designed to verify the operation of NVMe devices in the ultra-low power sub-states L1.1 and L1.2.

L1.x sub-state testing presents a unique, but crucial challenge for producers of NVMe devices destined for use in battery powered devices such as tablets and laptops. These devices will greatly benefit from the new power sub-states which trade a little PCIe latency for the lowest possible power consumption.

The fundamental difference with respect to L1.x sub-states is the use of an out-of-band signal to indicate to the device to "wake back up". The use of an out-of-band signal allows the device to achieve a near zero power consumption state by powering down its transceivers completely and monitoring only the level of the CLKREQ# out-of-band signal.

Primary PCIe Power States

A comprehensive overview of the PCIe power states, and specifically the L1.1 and L1.2 sub-states, can be found in the following PCI-SIG document <https://pcisig.com/making-most-pcie-low-power-features> by Scott Knowlton, Marketing Work Group Co-chair, PCI-SIG which is referenced here.

The primary PCIe power states are:

- L0 – a link which is operating normally.
- L1 – a link state where no data is being transferred so key portions of the PCIe transceiver logic can be turned off.
- L2 – a link state identical to L3 but in which power has not (yet) been removed.
- L3 – when the device is powered off.
- L0s – a link state where data may be being transferred in one direction but not the other, so the two devices on a link can each independently idle their transmitter.

Defining L1.1 and L1.2 Sub-States

As PCIe speeds and lane density have increased, a significant amount of power is consumed simply keeping the PCIe channel alive, but turning the transceiver PHYs off would, by definition, close the communication channel required to turn them back on.

To resolve this dilemma, PCIe NVMe defines the out-of-band signal "CLKREQ#", which in its simplest definition is a "wake up" signal from the host to the end point device to turn the PHY's back on and begin PCIe communication again. Turning the PHYs off can save a significant amount of power and allow the endpoint to consume almost no power while in idle mode, a dramatic improvement for a battery-operated device.

Scott Knowlton's paper describes the benefits of the L1.x sub-states as follows, and I include the text from the paper referenced above for completeness.

"The fundamental idea behind L1 sub-states is to use something other than the high-speed logic inside the PCIe transceivers to wake the devices. The goal is to achieve near zero power consumption with an active state.

That is done by adding additional functionality to an existing PCIe pin (CLKREQ#) to provide a very simple signaling protocol. This allows the PCIe transceivers to turn off their high-speed circuits and rely on the new signaling to wake them up again. In fact, two of these new sub-states were defined: L1.1 and L1.2 providing their own power vs. exit latency trade-off choices. Both L1.1 and L1.2 permit the PCIe transceivers to turn off their PLLs along with their receivers and transmitters, while L1.2 even allows turning off the common mode keeper circuits.

The results are dramatic. Efficient circuit design and modern silicon processes mean that a representative PCI Express 4.0 x4 PHY (4 transceivers plus related digital logic for four lanes) running at the full 16GT/s data rate in L0 consumes somewhere in the range of 400-500mW. Utilizing L1.1, the same PHY's power consumption drops by a factor of around 20x to consume only 20-30mW. Accepting the slightly longer exit latency of L1.2 permits power consumption to fall by another 10x to a mere 2-3mW.

The figure below shows the low power solutions available with the existing L1 state compared to using L1 sub-states. It is expected that the power savings scale linearly for multi-lane links and implementing the L1 sub-states feature reduces power consumption at the increase of the L1 exit latency. Implementing L1 sub-states is key to reducing power consumption for mobile designs using PCIe."




Figure 1: L1.x State Diagram and Sub-State Definition

Configuring the PCIe LTR Registers for the L1 Substate Test Script

LPCIeDEV is SANBlaze’s PCIe device from the target proc file, which is needed in the setpci commands.

```
LPCIeDEV=`cat ${LPROCTARGET}|grep TargetName|awk -F'0000:' '{ print $2 }`
```

The following setpci commands are used to set the following registers:

```
#
# Set up the PCIe Express Capabilities to allow for L1.1 and L1.2
# This section may vary from drive vendor to vendor
#
# LTR Mechanism Enable
#
doCmd "setpci -s ${LPCIeDEV} CAP10+28.w=400:400"

#
# LTR Capability Register
#
doCmd "setpci -s ${LPCIeDEV} ECAP18+4.w=100f"
doCmd "setpci -s ${LPCIeDEV} ECAP18+6.w=100f"

#
# L1 PM Sbustates Extended Capability
#
doCmd "setpci -s ${LPCIeDEV} ECAP1e+a.w=4073"
```

Controlling CLKREQ#

As described in the previous section, the pin CLKREQ# available on NVMe U.2, U.3, EDSFF and M.2 connectors has been amended to provide a "wake up" signal to drives in the power states L1.1 and L1.2, as shown in Figure 1.

To verify the correct operation of the L1, L1.1 and L1.2 low power sub-states, the user must be able to monitor and control the CLKREQ# signal. This can be achieved using the SANBlaze RM5 NVMe test system and is available in the V10.5 software release.




Figure 2: L1.1 Waveforms Illustrating Upstream Port Initiated Exit




Figure 3: L1.1 Waveforms Illustrating Downstream Port Initiated Exit

CLKREQ# Signal

CLKREQ# is asserted while the device under test is in L0 (full power) state and may be asserted by the Upstream or Downstream device on the PCIe link to recover a device from L1.1 or L1.2 back to L0.Active.

Modes of CLKREQ# L1.1 L1.2 Testing

SANBlaze supports L1.1 and L1.2 testing on the SBExpress-RM5 and the upcoming SBExpress-DT5 NVMe test systems. There are currently five built-in scripts that work with the SANBlaze hardware, the operation of which is described below.

Based on the physical location of the device under test in the SBExpress-RM5 system, the tests operate in one of two modes. The scripts automatically determine the system configuration and run one of the two test modes based on system configuration and revision.

There are benefits to both modes of operation, and therefore it is recommended to test your device in each of the two modes, as described below.

Scripts for L1.1 L1.2 Sub-state Testing

The following scripts are available in the SANBlaze V10.5 software package, available from SANBlaze.

- PM_L1_Substate_Verify.sh
- PM_L0_Enable.sh
- PM_L1.1_Enable.sh
- PM_L1.2_Enable.sh
- PM_L1_Enable.sh

All scripts are variations of the first script PM_L1_Substate_Verify.sh, which walks the device under test from L0.Active to L1.Idle to L1.1 and L1.2. At each level, the scripts test the device under test for the expected power state transition by inspecting the PCIe bus link state and the current state of the CLKREQ# signal.

The variations are as follows:

- PM_L0_Enable.sh - Restores the device to L0.Active state.
- PM_L1.1_Enable.sh - Leaves the device in L1.1 state.
- PM_L1.2_Enable.sh - Leaves the device in L1.2 state.
- PM_L1_Enable.sh – Leaves the device in L1 (L1.0) state.

CLKREQ# Mode Definition

For the purpose of the scripts and the SANBlaze SBExpress-RM5 hardware, two test modes are defined. All slots of all SBExpress-RM5 systems can run Mode 1.

Mode 1 – SANBlaze “riser” controls CLKREQ#.

Mode 2 - CLKREQ# signal is controlled by an upstream PCIe device.

Mode 2 testing requires a custom firmware image, which is field upgradable and available by request from SANBlaze.

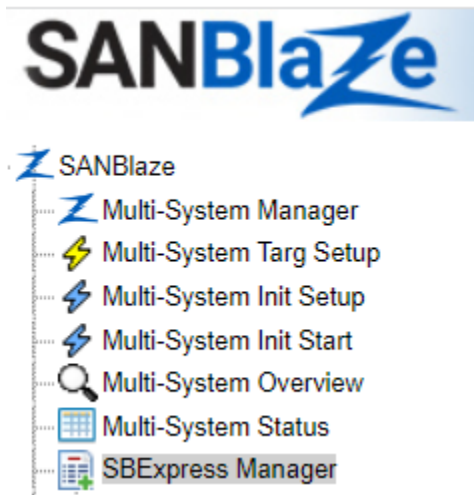
Staging the Mode 1 CLKREQ# Tests

In the first operating mode, the SANBlaze scripts control entering and exiting the L1.x sub-states by manually controlling CLKREQ#. This mode of operation verifies the sub-states and operates in all 16 slots of the SBExpress-RM5 test system in the manner described below.

Selecting "Test Manager" Testing

The Low Power Sub-State Tests are not considered part of the *Certified by SANBlaze* suite at this time, although they will be included in the *Certified by SANBlaze* Report, if selected manually. To select Test Manager testing:

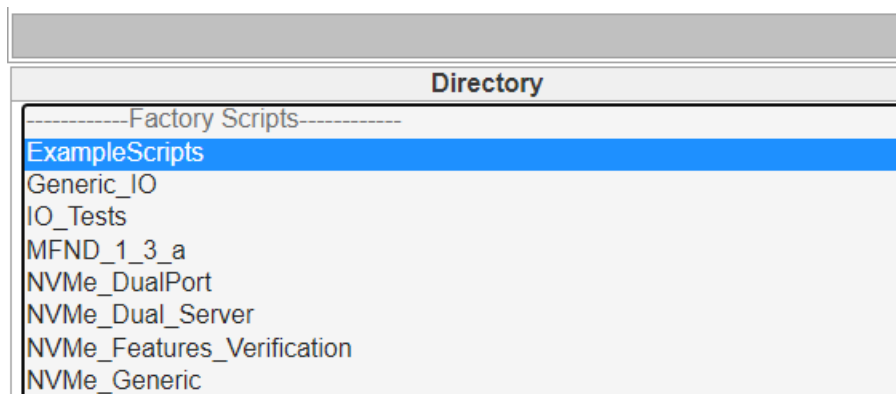
1. Select the SBExpress Manager Page from the left-hand menu:



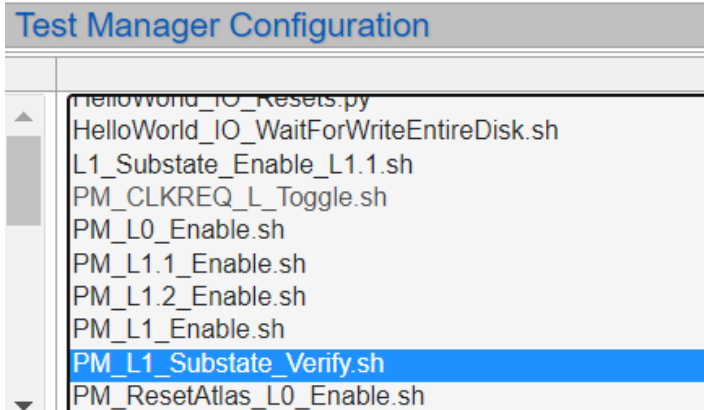
2. Select the "ShowTestManager" button, located near the bottom of the page:



3. Select "Example Scripts" on the left side of the page:



- Select the "PM_L1_Substate_Verify.sh" script on the right:



At this point, you may choose to assign the tests to all devices, set a "Pass Time" or a "Number of Passes". Pass time will be used as a dwell time between each sub-state transition, and the script runs the number of passes before declaring a "Pass".

For simplicity, leave the pass count at 1 and let the script set the dwell time.

- Select the "AddTests" button. The test is added to the selected device.

Controller	Namespace	Sequence #	# of Passes	Pass Time	Action
<input checked="" type="radio"/> 108 <input type="radio"/> All	<input checked="" type="radio"/> 1 <input type="radio"/> All	-1	1	1	AddTests

Your system will now be staged to run the low power sub-state testing on the selected device and the upper part of the page should show the test as "Idle" as shown below:

The screenshot shows the SANBlaze test management interface for a device with Vendor=NVMe, Product=Seagate FireCuda 530 ZP1000GM30013, SN=7VQ04A9G, Rev=SU6SM003. The test suite configuration shows:

- System Index: 1 All
- NVMe Initiator: 0 All
- Controller: 108 All
- Namespace: 1 All
- Loops: 1
- Report: Pass, Fail, Warn, Skip, All
- Include plots in new report:

 The test results table shows:

All	#	Seq	Name	State	Err/ Allowed	Pass/ Passes	Sec/ Pass	Start	End	Read Bytes	Write Bytes	Read I/Os	Write I/Os	% Done
<input checked="" type="checkbox"/>	1	1	PM_L1_Substate_Verify.sh	Idle	0/ 0	0/ 1	1			0	0	0	0	

- Selecting "Start" will run the test, placing the device under test in L0.Active state and then each of the L1.x sub-states as described above.

Monitoring the NVMe Sideband Signals

For the purpose of the demonstration of Mode1 testing, we are monitoring Link Statistics and the sideband signals PERST0#, CLKREQ#, PRSNT#, and DUALPORTEN# using a SerialTek Kodiak Gen5 PCIe/NMVe analyzer.

As seen in the trace, when the device under test deasserts CLKREQ# and shuts off the PCIe "Phys", you will see the system move to "Training". This is expected behavior for Mode1 testing.

When the system asserts CLKREQ#, the device under test will re-train and become active on the PCIe.

The script will monitor CLKREQ# and the state of the PCIe link and test each transition for expected behavior.

At the end of the test, the device will be given a "PASS, FAIL, or WARNING" status. Warning, for example, may be used if a device does not support a low power state that it claims to support.




Figure 4: SerialTek Kodiak Analyzer Results

SetFeatures NVMe Support

During the test, the "GetFeatures" and "SetFeatures" commands are executed in the following way. First a "GetFeatures" command is issued and the Number of Power States Supported (NPSS) value is read from the device.

The resultant NPSS value is sent to the device using SetFeatures to enable the lowest possible power state. Once the lowest power state is selected in this manner, the device is expected to enter and exit each low power state successfully.

Mode1 Testing Results

In mode1 CLKREQ# testing, the drive is expected to correctly respond to the GetFeatures request, and to accept the SetFeatures command to select the lowest power state. The drive is expected to behave in the manner below; if it does not, the drive will fail the test.

1. Drive starts in Full Power Mode L0.Active.
2. Drive responds to the GetFeatures command.
3. Drive accepts the SetFeatures command to select the lowest power state.
4. Script instructs the drive to L1.0 state using ASPMControl.
5. Script enables L1.1 state using L1PMControl.
6. Script checks that the drive deasserts CLKREQ#.
7. Script checks that PCIe link moves to "polling".
8. Script asserts CLKREQ# and verifies the drive returns to L0.Active state.
9. Steps 4 - 8 are repeated for L1.2 state.
10. Drive is returned to L0.Active and default values are restored to SetFeatures.

Staging the Mode 2 CLKREQ# Tests

In the second operating mode, the SANBlaze riser allows the PCIe parent of the device under test to control CLKREQ#. Staging the test is done by following the steps above for Mode1 testing.

The script determines if Mode2 testing is available on the given slot and performs Mode2 testing automatically.

Requirements for Mode2 Testing

To run Mode2 testing (Host Controls CLKREQ#), the SBExpress-RM5 system must be configured as follows:

- Device Under Test must be in slot 0 – 7.
- SANBlaze Riser supports passing CLKREQ# to host.
- SANBlaze M.2 adapter supports passing CLKREQ# to host.
- SBExpress-RM5 system must have B0 revision silicon.

The script will audit the system and run Mode1 if Mode2 is not available.

Execution of the Mode2 CLKREQ# Tests

In Mode2 CLKREQ# testing, the drive is expected to correctly respond to the GetFeatures request, and to accept the SetFeatures command to select the lowest power state. The drive is expected to behave in the manner below; if it does not, the drive will fail the test.

1. Drive starts in Full Power Mode L0.Active.
2. Drive responds to the GetFeatures command.
3. Drive accepts the SetFeatures command to select lowest power state.
4. Script instructs the drive to L1.0 state using ASPMControl.
5. Script enables L1.1 state using L1PMControl.
6. Script checks that drive and host have both deasserted CLKREQ#.
7. Script checks that PCIe link moves to "L1.Idle".
8. Script issues IO to the drive and verifies that the host asserts CLKREQ#.
9. Script expects the IOs to succeed.
10. Steps 4 - 8 are repeated for L1.2 state.
11. Drive is returned to L0.Active and default values are restored to SetFeatures.

The Kodiak PCIe analyzer trace will show the drive in L1.Idle and verify that the host is asserting CLKREQ# in response to IO requests as seen below.




Figure 5: The Kodiak PCIe analyzer trace verifies that the host is asserting CLKREQ#

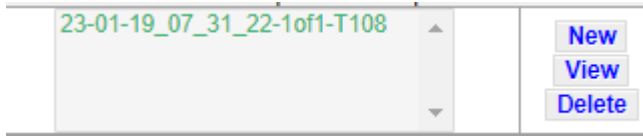
Interpreting and Publishing Test Results

Successful completion of the sub-state testing will be indicated on the SBExpress page, as shown here:

All	#	Seq	Name	State	Err/ Allowed	Pass/ Passes	Sec/ Pass	Start	End	Read Bytes	Write Bytes	Read I/Os	Write I/Os	% Done
<input checked="" type="checkbox"/>	1	1	PM_L1_Substate_Verify.sh	Passed	0/0	1/1	23	Jan19_07:30:55	Jan19_07:31:20	610304	0	25	0	<div style="width: 100%;"></div>

Selecting the link name of the test, in this case PM_L1_Substate_Verify.sh, will bring up a detailed log of the test execution.

Upon completion, the system will generate a test report which can be found on the SBExpress page here:



Selecting the name, and "View" will bring up the test report, which can then be exported or printed, with either a summary of the results, or the complete testing log.

Under the Covers for CLKREQ# Testing

Monitoring the test process using the Kodiak analyzer lets you see what is actually happening on the sideband and data path signals, but an analyzer is not necessary to run the sub-state tests.

Initiating Sub-States from the Command Line Interface (CLI)

In addition to the provided script, Sub-state transitions can be initialized and observed from the SANBlaze system CLI.

Use sb_sdb to Check Link State

sb_sdb is a tool which communicates over a serial bus to the PCIe switch. It can be used to query the switch for the status of the link. For example, the current link state of the device at -d 0 (first slot on the left of the enclosure) is L0.Active:

```
sb_sdb -d 0 -1
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR)
0x0BB0(LTSSM) State.Substate
  0 16(x10) 9e:00.0 1 0 4 4 0 0 0 0 1301h L0.Active
```

We will use sb_sdb to monitor the link status in this example of stepping the drive through the power sub-states.

Using SetFeatures to Enable Lowest Power State

First use Get and Set features to enable the lowest power state the device can support. Target number is 100 + slot number so our drive-in slot 0 is 100.

Using drive at slot 0 as our example = Target 100 = /iport0/target100

Next, find the number of power states supported and enable the lowest power state.

The command below will return the Number of Power States Supported:

```
cat /iport0/target100 |grep ^NPSS
NPSS=4
```

Highest power states are first as far as Power Features, so using the number from the command above will enable the lowest power state.

Use the NPSS number to send the "SetFeatures" command using the 4 from above for the -d data. For example:

```
io /iport0/target103 SetFeatures -fid 2 -d 4 -w
```

You have now used SetFeatures to tell the drive the lowest power state you want it to go into. It will still be in L0.Active, as seen below:

```
sb_sdb -d 0 -l
```

```
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR)
0x0BB0(LTSSM) State.Substate
  0 16(x10) 9e:00.0   1  0  4  4  0      0      0      0      1301h L0.Active
```

Turn off the System Watchdog

The SBExpress system has a "Watchdog" which will poll once per second unless disabled. This Watchdog will keep the drive alive, so for the purpose of the test, turn it off.

```
sbecho WatchdogDisabled=1 > /iport0/target100
```

Turn off the SMBus to the Drive

The SBExpress system reads data from the SMBus on the "Adapter" for the M.2 drive. The SMBus activity will keep the drive alive, so turn that off, as well.

```
sb_i2c2 -d 0 -U -f IOX_M2_ENABLE_I2C -w 0
```

Use ASPM to Enable Low Power States

Use the following ASPM command to enable L1.0:

```
sbecho ASPMControl=2 > /iport0/target100
```

You will see the drive in L1.Idle (L1.0 at this point).

```
sb_sdb -d 0 -1
```

```
Slot PSN(hex) PCIeDev Pres MRL Wid Spd Pwr 0x0FAC(TLP) 0x0FB0(DLLP) 0x0BF4(RERR)
0x0BB0(LTSSM) State.Substate
      0 16(x10) 9e:00.0 1 0 4 4 0 0 0 0 1906h L1.Idle06
```

You can see that you are not in L1.x because `clkreq_1` is still asserted:

```
sb_i2c2 -d 0 -f clkreq_1
```

```
INFO: System 01[03] SBExpress-RM5 SN=950A2050002 Rev=R01 i2c_main=i2c-8 i2c_mi=i2c-7
i2c_priv=i2c-5 SDB=/dev/ttyACM2 VLUN=0
```

```
INFO: 03[0x60] 0x05[07:06] clkreq_1 Def=0x01 Cur=0x00
```

Enabling L1.1 or L1.2

Use the following command to enable the L1 sub-state (8 = L1.1, 4 = L1.2):

```
sbecho L1PMControl=4 > /iport0/target100
```

At this point the drive will move to L1.2 substate, and will deassert `CLKREQ#`, as seen by the Kodiak Analyzer sideband trace below:

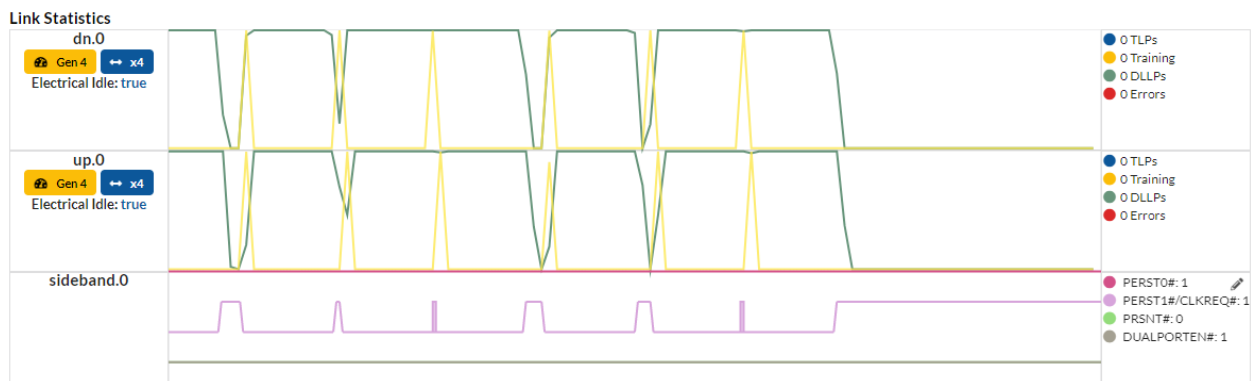


Figure 6: The Drive Deasserts `CLKREQ#`, as Shown in the Kodiak Analyzer Trace

You can also verify that `CLKREQ#` is disabled using the following command:

```
sb_i2c2 -d 3 -f clkreq_1
```

```
INFO: System 01[03] SBExpress-RM5 SN=950A2050002 Rev=R01 i2c_main=i2c-8 i2c_mi=i2c-7
i2c_priv=i2c-5 SDB=/dev/ttyACM2 VLUN=0
```

```
INFO: 03[0x60] 0x05[07:06] clkreq_1 Def=0x01 Cur=0x01
```

As shown above, `CLKREQ#` is now deasserted (0x01).

Initiating IO to the drive, or asserting CLKREQ# will bring the drive back online, for example by re-starting the Watchdog:

```
sbecho WatchdogDisabled=0 > /iport0/target100
```




Figure 7: Initiating IO to the Drive or Asserting CLKREQ# Brings the Drive Back Online

The behavior of the device under test will vary depending if the Host or the SBExpress system is controlling CLKREQ#; the drive will stay in L1 with respect to the upstream port or will move to Polling depending on Slot# (0 - 7) or (8 - 15).

You can manually bring the drive back to full power state by clearing the settings above:

```
sbecho L1PMControl=0 > /iport0/target100
sbecho ASPMControl=0 > /iport0/target100
```




Figure 8: Clear Settings to Manually Bring the Drive Back to Full Power State

Additional Capabilities of CLI Commands

While testing power states, these CLI programs are particularly useful:

- `sb_sdb` - Performs out of band communication to the PCIe Gen5 switch.
- `sb_i2c2` - Performs I/O to the system risers and adapters.

Some examples useful for power state testing follow:

<code>sb_sdb -d 0 -l</code>	Provides current link state information for the device at slot -d.
<code>sb_i2c2 -d 0 -i</code>	Provides information on the riser at slot -d.
<code>sb_i2c2 -d 0 -f ?</code>	Provides a list of "features" supported by the riser (e.g. <code>clkreq_l</code>).
<code>sb_i2c2 -d 0 -i -U</code>	Provides information on the m.2 adapter riser at slot -d.
<code>sb_i2c2 -d 0 -f ? -U</code>	Provides a list of "features" supported by the m.2 adapter.

Examples of `sb_i2c2` Commands

Assert PERST to the device at slot 0:

```
sb_i2c2 -d 0 -f PORT0_PERST_L -w 0
```

Deassert PERST:

```
sb_i2c2 -d 0 -f PORT0_PERST_L -w 1
```

Power off the device at slot 0:

```
sb_i2c2 -d 0 -f DISABLE_12V_L -w 0
```

Enable power:

```
sb_i2c2 -d 0 -f DISABLE_12V_L -w 1
```

Read and Write `CLKREQ_L`

Keep in mind that there may be multiple devices driving `CLKREQ_L` and because the signal is wired OR any drive asserting it will keep it asserted regardless of the other devices. The SANBlaze scripts start by deasserting `CLKREQ` on the chosen Riser and Adapter with the following commands:

Deassert `CLKREQ#` from riser :

```
sb_i2c2 -d 0 -f CLKREQ_L -w 1
```

Deassert `CLKREQ#` from adapter:

```
sb_i2c2 -d 0 -f CLKREQ_L -w 1 -U
```

The device under test will deassert `CLKREQ#` when going into L1 sub-states.

In Mode2 - Host drives `CLKREQ#`, the host will assert `CLKREQ` automatically when the host issues I/O to the drive.

In Mode1 - SANBlaze riser or adapter drives `CLKREQ#` to recover the drive to the L0.Active state.

Conclusion

SANBlaze provides a simple means of testing low power states L1.1 and L1.2 and provides automated scripts to validate the power states on NVMe drives capable of L1 sub-states.

Two methods are provided to validate that the device under test achieves the desired state.

- Mode1: SANBlaze riser controlling CLKREQ# transitions via SANBlaze scripts or customer scripts.
- Mode2: Host based CLKREQ# control which will automatically assert/deassert CLKREQ# based on data access to the device.

The automated tests can be added to an existing test suite and included in the final report for your device under test.

Testing can be done with a Kodiak Gen5 PCIe analyzer with an EDSFF or U.2 interposer for live monitoring of the CLKREQ# signal and link state, or these states can be read from the system CLI.

Minimum Software revision of V10.5 is required for L1 sub-state testing, Contact SANBlaze for more information on this capability.

Determining the Capability of your RM5 System

As stated above, a CLKREQ# capable firmware image is required for Mode2 testing. The image can be requested from SANBlaze.

For V10.5 and above, you can query your SANBlaze FW version using the `sb_flash` command, as follows:

```
sb_flash show
INFO: sb_flash sys=1 sdb port=/dev/ttyACM4 selected SBR Flash=0
INFO: Available Images
  Index ImageID Type Ver Mode UpLNK SSC/CFC Clock System Description
    1 03b60101 01 01 Base 30 CFC CC 03b6 Default CFC All Ports
    2 03b60201 02 01 Base 30 SSC SRIS 03b6 SRIS same clock all ports
    3 03b60301 03 01 Base 30 SSC SRIS 03b6 SRIS 2 different
.....clocks FPGA isolation
    4 03b60402 04 02 Base 70 CFC CC 03b6 CFC All ports
    5 03b60502 05 02 Base 30 CFC CC 03b6 CFC All ports clkreq_1 to
GPIO LED1
    6 03b60601 06 01 Base 30 SSC SRIS 03b6 SRIS same clock all ports
clkreq_1 to GPIO LED1
    7 03b60701 07 01 Base 30 SSC SRIS 03b6 SRIS 2 different clocks
FPGA isolation clkreq_1 to GPIO LED1
    8 03b60801 08 01 Base 70 CFC CC 03b6 CFC All ports clkreq_1
to GPIO LED1
INFO: Current Active Images:
  Index ImageID Type Ver Mode UpLNK SSC/CFC Clock System Description
Active: 0 03b60402 04 02 Base 30 CFC CC 03b6 CFC All ports
    1 03b60402 04 02 Base 30 CFC CC 03b6 CFC All ports
    2 03b60801 08 01 Base 30 CFC CC 03b6 CFC All ports clkreq_1
to GPIO LED1
    3 03b60801 08 01 Base 30 CFC CC 03b6 CFC All ports clkreq_1
to GPIO LED1
```

The image marked "Active" is the current running image. If the running image supports `clkreq_1`, mode2 (CLKREQ# from host) will be enabled on slots 0-7.

To change the image on the next boot, use:

```
sb_flash select N, where N is flash index 0-3
```

To replace one of the four images with one from the available images use:

```
sb_flash select 3 image 03b60801 show (ID from list of available images)
```

This will select flash bank 3 and program image 03b60801.

The image name above will indicate `clkreq_l` if `CLKREQ#` is supported in the lower 7 slots in a given image.

Note: A reboot is required after loading a new flash image or selecting an alternate flash image.

FAQs:

Q: I have an analyzer interposer between the SANBlaze EDSFF Gen5 "Riser" and the SANBlaze EDSFF to M.2 "Adapter". The PCIe establishes link, but the `CLKREQ#` signal doesn't seem to toggle when I run the test, and the test fails. If I take the analyzer out, the test passes.

A: Some of the early M.2 Adapters (Model 915) ran the `CLKREQ#` signal on a non-standard EDSFF pin. The adapter will work when plugged in directly, but the analyzer will not see the `CLKREQ#` signal. Contact SANBlaze to arrange an RMA for the 915 adapter.

Q: Does the 957 SANBlaze U.2 Riser and 923 U.2 to M.2 adapter have the issue above?

A: No, the 957 will accept a U.2 interposer and pass the signals through the analyzer interposer correctly.

Q: I have a SANBlaze RM5 0002 of 0003 revision motherboard, why can't I run the Mode2 `CLKREQ#` tests on slots 8-15?

A: `CLKREQ#` is a physical out-of-band signal that routes from the Device Under Test to the PCIe Gen5 bridge. `CLKREQ#` is physically routed to the bridge only from slots 0-7.

Q: I tried the `CLKREQ#` image on my A0 based RM5 motherboard and it worked. Why do I need to upgrade the system to B0?

A: There is an erratum for the A0 chip that specifically excludes L1.1 and L1.2 operation.

Q: I want to test L1.1 and L1.2 states on my A0 based system anyway. What can I expect?

A: When exiting L1.1 or L1.2 to return to L1.0, the A0 part may stop responding and leave the drive in a "polling" state. If you do 1000 iterations of the test, you are likely to see this issue.

Q: I want to test L1.1 and L1.2 states on my A0 based system anyway. If I see the issue above, what's the mechanism to recover my device?

A: You may issue the following command: `sbecho reset_atlas > /proc/vlun/nvmeto` clear the issue, but you will be resetting the entire PCIe subsystem and therefore all testing on all slots will be affected.