Aerial System Scripts

Table of contents

System Configuration Validation Script

Included in the release package is a script that checks and displays key system configuration settings that are important for running the Aerial cuBB SDK.

> $ pip3 install psutil $ cd $cuBB_SDK/cuPHY/util/cuBB_system_checks $ sudo -E python3 ./cuBB_system_checks.py

The output of `cuBB_system_checks.py` may differ slightly between bare-metal and container versions of the environment. The script helps to retrieve the software-component versions and hardware configuration. Refer to the Release Manifest in the *cuBB Release Notes* to ensure the correct software-component versions are installed. Below is an example output on a bare-metal platform:

> *# To get the system or ptp info, the command has to run on the host.* $ sudo python3 cuBB_system_checks.py --sys -----General---------- Hostname : devkit-1 IP address : 192.168.1.100 Linux distro : "Ubuntu 22.04.3 LTS" Linux kernel version : 5.15.0-1042-nvidia -----System---------- Manufacturer : GIGABYTE Product Name : E251-U70-00 Base Board Manufacturer : GIGABYTE Base Board Product Name : MU71-SU0-00 Chassis Manufacturer : GIGABYTE Chassis Type : Rack Mount Chassis Chassis Height : Unspecified Processor : Intel(R) Xeon(R) Gold 6240R CPU @ 2.40GHz Max Speed : 4000 MHz Current Speed : 2400 MHz $ sudo python3 cuBB_system_checks.py -----General---------- Hostname : devkit-1 IP address : 192.168.1.100 Linux distro : "Ubuntu 22.04.3 LTS" Linux kernel version : 5.15.0-1042-nvidia -----Kernel Command Line---------- Audit subsystem : audit=0 Clock source : clocksource=tsc HugePage count : hugepages=16 HugePage size : hugepagesz=1G CPU idle time management : idle=poll Max Intel C-state : intel_idle.max_cstate=0 Intel IOMMU : intel_iommu=off IOMMU : iommu=off Isolated CPUs : isolcpus=2-21 Corrected errors : mce=ignore_ce Adaptive-tick CPUs : nohz_full=2-21 Soft-lockup detector disable : nosoftlockup Max processor C-state : processor.max_cstate=0 RCU callback polling : rcu_nocb_poll No-RCU-callback CPUs : rcu_nocbs=2-21 TSC stability checks : tsc=reliable -----CPU---------- ----- CPU cores : 24 Thread(s) per CPU core : 1 CPU MHz: : N/A CPU sockets : 1 ----- Environment variables---------- CUDA_DEVICE_MAX_CONNECTIONS : N/A cuBB_SDK : N/A -----Memory---------- HugePage count : 16 Free

HugePages : 16 HugePage size : 1048576 kB Shared memory size : 47G -----Nvidia GPUs---------- GPU driver version : 535.54.03 CUDA version : 12.2 GPU0 GPU product name : NVIDIA A100-PCIE-40GB GPU persistence mode : Enabled Current GPU temperature : 27 C GPU clock frequency : 1410 MHz Max GPU clock frequency : 1410 MHz GPU PCIe bus id : 00000000:B6:00.0 -----GPUDirect topology---------- GPU0 NIC0 NIC1 CPU Affinity NUMA Affinity GPU NUMA ID GPU0 X PIX PIX 0-23 N/A N/A NIC0 PIX X PIX NIC1 PIX PIX X Legend: X = Self SYS = Connection traversing PCIe as well as the SMP interconnect between NUMA nodes (e.g., QPI/UPI) NODE = Connection traversing PCIe as well as the interconnect between PCIe Host Bridges within a NUMA node PHB = Connection traversing PCIe as well as a PCIe Host Bridge (typically the CPU) PXB = Connection traversing multiple PCIe bridges (without traversing the PCIe Host Bridge) PIX = Connection traversing at most a single PCIe bridge NV# = Connection traversing a bonded set of # *NVLinks* NIC Legend: NIC0: mlx5_0 NIC1: mlx5_1 -----Mellanox NICs---------- -------- NIC0 NIC product name : ConnectX6DX NIC part number : MCX623106AE-CDA_Ax NIC PCIe bus id : 0000:b5:00.0 NIC FW version : 22.39.2048 FLEX_PARSER_PROFILE_ENABLE : 4 PROG_PARSE_GRAPH : True(1) ACCURATE_TX_SCHEDULER : True(1) CQE_COMPRESSION : AGGRESSIVE(1) REAL_TIME_CLOCK_ENABLE : True(1) -----Mellanox NIC Interfaces---------- Interface0 Name : ens6f0 Network adapter : mlx5_0 PCIe bus id : 0000:b5:00.0 Ethernet address : b8:ce:f6:33:fd:ee Operstate : up MTU : 1514 RX flow control : off TX flow control : off PTP hardware clock : 2 QoS Priority trust state : pcp PCIe MRRS : 4096 bytes Interface1 Name : ens6f1 Network adapter : mlx5_1 PCIe bus id : 0000:b5:00.1 Ethernet address : b8:ce:f6:33:fd:ef Operstate : up MTU : 1500 RX flow control : off TX flow control : off PTP hardware clock : 3 QoS Priority trust state : pcp PCIe MRRS : 512 bytes -----Linux PTP---------- ptp4l.service - Precision Time Protocol (PTP) service Loaded: loaded (/lib/systemd/system/ptp4l.service; enabled; vendor preset: enabled) Active: active (running) since Tue 2022-09-27 00:05:26 UTC; 1 day 7h ago Docs: man:ptp4l Main PID: 1594 (ptp4l) Tasks: 1 (limit: 94581) Memory: 840.0K CGroup: /system.slice/ptp4l.service 1594 /usr/sbin/ptp4l -f /etc/ptp.conf Sep 27 00:05:26 dc6-devkit-18 systemd[1]: Started Precision Time Protocol (PTP) service. Sep 27 00:05:26 dc6-devkit-18 taskset[1594]: ptp4l[127.145]: selected /dev/ptp2 as PTP clock Sep 27 00:05:27 dc6-devkit-18 taskset[1594]: ptp4l[127.162]: port 1: INITIALIZING to LISTENING on INIT_COMPLETE Sep 27 00:05:27 dc6-devkit-18 taskset[1594]: ptp4l[127.162]: port 0: INITIALIZING to LISTENING on INIT_COMPLETE

Sep 27 00:05:27 dc6-devkit-18 taskset[1594]: ptp4l[127.186]: port 1: new foreign master b8cef6.fffe.33fe16-1 Sep 27 00:05:27 dc6-devkit-18 taskset[1594]: ptp4l[127.436]: selected best master clock b8cef6.fffe.33fe16 Sep 27 00:05:27 dc6-devkit-18 taskset[1594]: ptp4l[127.436]: assuming the grand master role Sep 27 00:05:27 dc6-devkit-18 taskset[1594]: ptp4l[127.436]: port 1: LISTENING to GRAND_MASTER on RS_GRAND_MASTER phc2sys.service - Synchronize system clock or PTP hardware clock (PHC) Loaded: loaded (/lib/systemd/system/phc2sys.service; enabled; vendor preset: enabled) Active: active (running) since Tue 2022-09-27 00:05:26 UTC; 1 day 7h ago Docs: man:phc2sys Main PID: 1598 (sh) Tasks: 2 (limit: 94581) Memory: 5.4M CGroup: /system.slice/phc2sys.service 1598 /bin/sh -c /usr/sbin/phc2sys -s /dev/ptp$(ethtool -T $(lshw -c network -businfo | grep b5:00.0 | awk '{print $2}') | grep PTP | awk '{print $4}') -c CLOCK_REALTIME -n 24 -O 0 -R 256 -u 256 1897 /usr/sbin/phc2sys -s /dev/ptp2 -c CLOCK_REALTIME -n 24 -O 0 -R 256 -u 256 Sep 28 07:16:46 dc6-devkit-18 phc2sys[1897]: [112407.124] CLOCK_REALTIME rms 10 max 34 freq +7048 +/- 25 delay 1765 +/- 8 Sep 28 07:16:47 dc6-devkit-18 phc2sys[1897]: [112408.140] CLOCK_REALTIME rms 10 max 27 freq +7031 +/- 39 delay 1765 +/- 8 Sep 28 07:16:49 dc6-devkit-18 phc2sys[1897]: [112409.155] CLOCK_REALTIME rms 9 max 27 freq +7044 +/- 30 delay 1764 +/- 7 Sep 28 07:16:50 dc6-devkit-18 phc2sys[1897]: [112410.171] CLOCK_REALTIME rms 9 max 24 freq +7041 +/- 17 delay 1765 +/- 8 Sep 28 07:16:51 dc6-devkit-18 phc2sys[1897]: [112411.188] CLOCK_REALTIME rms 9 max 28 freq +7036 +/- 21 delay 1766 +/- 7 Sep 28 07:16:52 dc6-devkit-18 phc2sys[1897]: [112412.203] CLOCK_REALTIME rms 9 max 22 freq +7055 +/- 21 delay 1766 +/- 7 Sep 28 07:16:53 dc6-devkit-18 phc2sys[1897]: [112413.219] CLOCK_REALTIME rms 9 max 24 freq +7038 +/- 20 delay 1764 +/- 8 Sep 28 07:16:54 dc6-devkit-18 phc2sys[1897]: [112414.235] CLOCK_REALTIME rms 9 max 23 freq +7041 +/- 19 delay 1763 +/- 7 Sep 28 07:16:55 dc6-devkit-18 phc2sys[1897]: [112415.251] CLOCK_REALTIME rms 9 max 22 freq +7043 +/- 11 delay 1763 +/- 8 Sep 28 07:16:56 dc6-devkit-18 phc2sys[1897]: [112416.267] CLOCK_REALTIME rms 10 max 24 freq +7052 +/- 20 delay 1762 +/- 7 Sep 28 07:16:57 dc6-devkit-18 phc2sys[1897]: [112417.283] CLOCK_REALTIME rms 10 max 30 freq +7035 +/- 39 delay 1765 +/- 8 -----Software Packages---------- cmake : N/A docker /usr/bin : 24.0.7 gcc /usr/bin : 11.4.0 git-lfs : N/A MOFED : N/A meson : N/A ninja : N/A ptp4l /usr/sbin : 3.1.1-3 -----Loaded Kernel Modules---------- GDRCopy : gdrdrv GPUDirect RDMA : N/A Nvidia : nvidia -----Non-persistent settings----------

```
----- VM swappiness : vm.swappiness = 60 VM zone reclaim mode :
vm.zone_reclaim_mode = 0 -----Docker images-------------------------------
```

Checking the NIC Status

To query back the Mellanox NIC firmware settings initialized with the script above, use these commands:

```
$ sudo mlxconfig -d /dev/mst/mt4125_pciconf0 q | grep
"CQE_COMPRESSION\|PROG_PARSE_GRAPH\
\|FLEX_PARSER_PROFILE_ENABLE\|REAL_TIME_CLOCK_ENABLE\|ACCURATE_TX_SCHEI
# FLEX_PARSER_PROFILE_ENABLE 4 # PROG_PARSE_GRAPH True(1) #
ACCURATE_TX_SCHEDULER True(1) # CQE_COMPRESSION AGGRESSIVE(1) #
REAL_TIME_CLOCK_ENABLE True(1)
```

To check the current status of a NIC port, use this command:

```
$ sudo mlxlink -d /dev/mst/mt4125_pciconf0
```

Alternatively, you can use the System Configuration Validation Script to obtain a full list of configuration settings.