AGS8200
AI Server

Software User Guide

www.edge-core.com

Software User Guide

AGS8200

AI Server

E072024-CS-R01

How to Use This Guide

This guide includes detailed information on the AI Server, including how to install software components and set up the system. To deploy this device effectively and ensure trouble-free operation, you should first read the relevant sections in this guide.

Who Should Read this Guide?

This guide is for system technicians who are responsible for installing, maintaining, and troubleshooting this device.

How this Guide is Organized

The organization of this guide is based on software installation, monitoring hardware, and usage examples.

The guide includes these sections:

- Chapter 1 "Software Environment" — Lists the software package matrix and libraries.
- Chapter 2 "Software Installation" — Includes how to install software on the system, troubleshooting, and management access.
- Chapter 3 "Hardware Monitor" — Includes information on commands for monitoring system status.
- Chapter 4 "System Firmware Update" — Includes information on updating system firmware.
- Chapter 5 "Usage Examples" — Includes examples for network topology and using Dockers for model training.
- Chapter 6 "Scale-Out Network Reference Design" — Includes information on linking up all the scale-out ports of the AI Servers.
- Chapter 7 "References" — Includes a list of Intel® Habana Official Information provided as reference.

Conventions

The following conventions are used throughout this guide to show information:

Note: Emphasizes important information or calls your attention to related features or instructions.

Caution: Alerts you to a potential hazard that could cause loss of data, or damage the system or equipment.

Warning: Alerts you to a potential hazard that could cause personal injury.

Revision History

This section summarizes the changes in each revision of this guide.

July 2024 Revision

This is the first revision of this guide.

Contents

Figures

Tables

1 Software Environment

This chapter includes the following sections:

Software Package Matrix

Note: This user guide is based on the SynapseAI version 1.14.0. If you want to upgrade and know the latest software support matrix, refer to Intel® Habana's official document page. Habana Support Matrix

Table 1: Software Package Matrix

Component	Version
SynapseAI	1.14.0
Operating System ■ Version ■ Kernel	Ubuntu: ■ 22.04 ■ 5.4.0 and above
Python	Framework: ■ PyTorch: 3.10 ■ TensorFlow: 3.10
Kubernetes	1.25, 1.24
Docker	25.0.1
PyTorch	2.1.1
PyTorch Lightning/ Lightning	2.1.2
lightning-habana	1.3.0
DeepSpeed	Forked from 0.12.4 of the official DeepSpeed
TensorFlow	2.15.0
Intel Gaudi Horovod	Forked from 0.27.0 of the official Horovod
Open MPI	4.1.5
Libfabric	■ General support: 1.16.1 and above ■ Gaudi Direct with Verbs: 1.20.0
Optimum Habana	1.9
Transformers	4.34

Software Library List

Table 2: Library List

Software	Library List
System	gcc, cmake, lsof, curl, wget, ethtool, libelf-dev, [br] libbz2-dev, liblzma-dev, libibverbs-dev, librdmacm-dev, dkms, python3-dev, python3-pip, python3-venv, linux-modules-extra-5.15, linux-headers-5.15
TensorFlow	jemalloc, mesa-libGL, Python 3.10.9, libjemalloc2, protobuf-compiler, libgl1
PyTorch	unzip, libcurl4, moreutils, iproute2, libcairo2-dev, libglib2.0-dev, libselinux1-dev, libnuma-dev, libpcre2-dev, libatlas-base-dev, libjpeg-dev, liblapack-dev, libnuma-dev, google-perftools, numactl, libopenblas-dev
Network Manage	lldptool
HCCL	libfabric, hccl_ofi_wrapper, openmpi
Habana Driver	habanalabs-container-runtime, habanalabs-dkms, habanalabs-firmware, habanalabs-firmware-odm, habanalabs-firmware-tools, habanalabs-graph, habanalabs-rdma-core, habanalabs-qual, habanalabs-thunk

2 Software Installation

This chapter includes the following sections:

Software Installation Guide

This section provides a summary of the installation on the Intel® Gaudi® 2 device with the following steps.

For detailed instruction, please refer to Bare_Metal_Fresh_OS.

1. Operating System Installation.

   Following the Habana Support Matrix, the operating system used as an example is Ubuntu 22.04.

   Please prepare the OS image file and install with BIOS.

2. SynapseAI Installation.

   The user can install the SynapseAI software stack and driver by using the **habanalabs** installer.

   The commands below will download the installer and install all dependencies automatically.

```
wget -nv https://vault.habana.ai/artifactory/gaudi-installer/latest/
  habanalabs-installer.sh
chmod +x habanalabs-installer.sh
./habanalabs-installer.sh install --type base
```

3. Network Set-up.

   After installing SynapseAI, you can verify the Intel® Gaudi® 2 AI Accelerator status with the **hl-smi** command (See "Hardware Monitor"), and if you can see the OAM list, then the device is ready to train/inference with Intel® Gaudi® 2.

   For scale-out training/inferencing, you must set up the Intel® Gaudi® 2 AI Accelerator's network configuration by using **manage_network_ifs.sh** script to configure and link up all the network interfaces and check their status.

   Detailed information is described in the sections "Usage Examples" and "Scale-Out Network Reference Design".

Troubleshooting Commands

Uninstall SynapseAI Software Stack and Driver Before Update

When updating the SynapseAI software stack and driver, it is suggested to re-install the SynapseAI.

The user can use **habanalabs-installer.sh** to uninstall SynapseAI software.

```
root@habanas01:/# ./habanalabs-installer.sh uninstall
```

Uninstall habanalabs-dkms Manually

If the kernel versions have changed, **habanalabs-dkms** might have some problems when uninstalling it. You can uninstall it manually.

```
root@habanas01:/# dkms status
habanalabs-dkms/1.14.0-493, 5.15.0-94-generic, x86_64: installed.
root@habanas01:/# dkms uninstall habanalabs-dkms/1.14.0-493 -k 5.15.0-94-
  generic
root@habanas01:/# dkms remove habanalabs-dkms/1.14.0-493 -k 5.15.0-94-generic
```

Error When Installing habanalabs-rdma-core

If you meet an issue when installing **habanalabs-rdma-core**, add the following parameter and install again.

```
root@habanas01:/# export EXTRA_CMAKE_FLAGS="-DNO_PYVERBS=1"
```

Management Access

This section is used to introduce control plane login.

Console Switch

You can use a special key sequence to switch the console input and output to the CPU system or the BMC system.

Table 3: Console Switch Key Sequence

System	Key Sequence
CPU Linux system	Press ctrl+u and release, then press 1.
BMC system	Press ctrl+u and release, then press 2.

Management Login

You can use SSH or the console port to access the CPU system or BMC. The following table provides the default login information.

Table 4: Device Login Information

Item	CPU Linux System	BMC System
Hostname	ags8200	ast-2600-ags8200
Username	ubuntu	root
Password	admin	0penBmc

BMC IP Address

You can log in to the BMC using SSH or the console to change its IP address. The BMC IP address setting can be changed by IPMI tools under the CPU system. The following are some example commands.

```
root@ags8200:/home/ubuntu# ipmitool lan set 1 ipsrc static
root@ags8200:/home/ubuntu# ipmitool lan set 1 ipaddr 193.168.8.110
root@ags8200:/home/ubuntu# ipmitool lan set 1 netmask 255.255.0.0
root@ags8200:/home/ubuntu# ipmitool lan set 1 defgw ipaddr 193.168.110.254
```

3 Hardware Monitor

This chapter includes the following sections:

Hardware Monitor Commands

This section is used to introduce the HW monitor commands.

These commands can be used to observe the Intel® Gaudi® 2 AI Accelerator status.

System Management Interface Tool

hl-smi: Dump the current status of Intel® Gaudi® 2 AI Accelerators. For example, the usage, power consumption, and training process.

```
root@ags8200:/home/ubuntu# hl-smi
+-----------------------------------------------------------------------------+
| HL-SMI Version:                              hl-1.14.0-fw-48.0.1.0          |
| Driver Version:                                    1.14.0-9e8ecf8           |
|-------------------------------+----------------------+----------------------+
| AIP  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | AIP-Util  Compute M. |
|===============================+======================+======================|
|   0  HL-225              N/A  | 0000:33:00.0     N/A |                    0 |
| N/A   25C   N/A    77W / 600W |    768MiB / 98304MiB |     0%           N/A |
|-------------------------------+----------------------+----------------------+
|   1  HL-225              N/A  | 0000:9a:00.0     N/A |                    0 |
| N/A   22C   N/A    62W / 600W |    768MiB / 98304MiB |     0%           N/A |
|-------------------------------+----------------------+----------------------+
|   2  HL-225              N/A  | 0000:34:00.0     N/A |                    0 |
| N/A   25C   N/A    91W / 600W |    768MiB / 98304MiB |     0%           N/A |
|-------------------------------+----------------------+----------------------+
|   3  HL-225              N/A  | 0000:9b:00.0     N/A |                    0 |
| N/A   25C   N/A    88W / 600W |    768MiB / 98304MiB |     0%           N/A |
|-------------------------------+----------------------+----------------------+
|   4  HL-225              N/A  | 0000:4d:00.0     N/A |                    0 |
| N/A   26C   N/A    82W / 600W |    768MiB / 98304MiB |     0%           N/A |
|-------------------------------+----------------------+----------------------+
|   5  HL-225              N/A  | 0000:b3:00.0     N/A |                    0 |
| N/A   26C   N/A    97W / 600W |    768MiB / 98304MiB |     0%           N/A |
|-------------------------------+----------------------+----------------------+
|   6  HL-225              N/A  | 0000:4e:00.0     N/A |                    0 |
| N/A   23C   N/A    90W / 600W |    768MiB / 98304MiB |     0%           N/A |
|-------------------------------+----------------------+----------------------+
|   7  HL-225              N/A  | 0000:b4:00.0     N/A |                    0 |
| N/A   22C   N/A    90W / 600W |    768MiB / 98304MiB |     0%           N/A |
|-------------------------------+----------------------+----------------------+
| Compute Processes:                                               AIP Memory |
|  AIP       PID   Type   Process name                             Usage      |
|=============================================================================|
|   0        N/A   N/A    N/A                                      N/A        |
|   1        N/A   N/A    N/A                                      N/A        |
|   2        N/A   N/A    N/A                                      N/A        |
|   3        N/A   N/A    N/A                                      N/A        |
|   4        N/A   N/A    N/A                                      N/A        |
|   5        N/A   N/A    N/A                                      N/A        |
|   6        N/A   N/A    N/A                                      N/A        |
|   7        N/A   N/A    N/A                                      N/A        |
```

```
+==========================================================================+
```

Intelligent Platform Management Interface (IPMI)

ipmitool sensor: You can use ipmitool to show the peripheral device information that is controlled by the BMC.

```
root@ags8200:/home/ubuntu# ipmitool sensor
CPU0_FIVRA_Iout  | 45.400     | Amps       | ok    | na        | na        | na        | na        | na
  | na
CPU0_PVCCD_Iin   | 0.000      | Amps       | ok    | na        | na        | na        | na        | na
  | na
CPU0_PVCCIN_Iout | 49.940     | Amps       | ok    | na        | na        | na        | na        | na
  | na
CPU1_FIVRA_Iout  | 54.480     | Amps       | ok    | na        | na        | na        | na        | na
  | na
CPU1_PVCCD_Iin   | 0.234      | Amps       | ok    | na        | na        | na        | na        | na
  | na
CPU1_PVCCIN_Iout | 15.890     | Amps       | ok    | na        | na        | na        | na        | na
  | na
FAON_CPU0_Iout   | 6.810      | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
FAON_CPU1_Iout   | 6.810      | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
P54V_Iout        | 129.390    | Amps       | ok    | na        | na        | na        | na        | na
  | na
PSU1_54V_Iin     | 1.014      | Amps       | ok    | na        | na        | na        | 12.012    |
  14.040    | na
PSU1_54V_Iout    | 4.540      | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
PSU2_54V_Iin     | 1.170      | Amps       | ok    | na        | na        | na        | 12.012    |
  14.040    | na
PSU2_54V_Iout    | 4.540      | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
PSU3_54V_Iin     | 1.170      | Amps       | ok    | na        | na        | na        | 12.012    |
  14.040    | na
PSU3_54V_Iout    | 4.540      | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
PSU4_54V_Iin     | 1.092      | Amps       | ok    | na        | na        | na        | 12.012    |
  14.040    | na
PSU4_54V_Iout    | 4.540      | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
PSU5_54V_Iin     | 0.936      | Amps       | ok    | na        | na        | na        | 12.012    |
  14.040    | na
PSU5_54V_Iout    | 2.270      | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
PSU6_54V_Iin     | 1.014      | Amps       | ok    | na        | na        | na        | 12.012    |
  14.040    | na
PSU6_54V_Iout    | 4.540      | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
PSU7_Iin         | 0.780      | Amps       | ok    | na        | na        | na        | 12.012    |
  14.040    | na
PSU7_Iout        | 13.620     | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
PSU8_Iin         | 0.702      | Amps       | ok    | na        | na        | na        | 12.012    |
  14.040    | na
PSU8_Iout        | 11.350     | Amps       | ok    | na        | na        | na        | 70.370    |
  72.640    | na
```

```
SW12_0V8_Iout    | 83.990     | Amps       | ok    | na        | na        | na        | na        | na
  | na
SW34_0V8_Iout    | 83.990     | Amps       | ok    | na        | na        | na        | na        | na
  | na
Pwm_1            | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_2            | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_3            | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_4            | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_5            | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_6            | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_7            | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_8            | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_9            | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_10           | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_11           | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_12           | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_13           | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_14           | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_15           | 19.600     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_PSU1_54V_Fan | 30.184     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_PSU2_54V_Fan | 30.184     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_PSU3_54V_Fan | 30.184     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_PSU4_54V_Fan | 30.184     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_PSU5_54V_Fan | 30.184     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_PSU6_54V_Fan | 30.184     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_PSU7_Fan_1   | 16.072     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Pwm_PSU8_Fan_1   | 16.072     | unspecified | ok    | na        | na        | na        | na        | na
  | na
Fan_front_1      | 7154.000   | RPM        | ok    | na        | 1764.000  | 1960.000  | na        | na
  | na
Fan_front_2      | 7154.000   | RPM        | ok    | na        | 1764.000  | 1960.000  | na        | na
  | na
Fan_front_3      | 7154.000   | RPM        | ok    | na        | 1764.000  | 1960.000  | na        | na
  | na
Fan_front_4      | 6958.000   | RPM        | ok    | na        | 1764.000  | 1960.000  | na        | na
  | na
Fan_front_5      | 7154.000   | RPM        | ok    | na        | 1764.000  | 1960.000  | na        | na
  | na
Fan_front_6      | 7154.000   | RPM        | ok    | na        | 1764.000  | 1960.000  | na        | na
  | na
Fan_front_7      | 7350.000   | RPM        | ok    | na        | 1764.000  | 1960.000  | na        | na
  | na
```

```
Fan_front_8      | 7154.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_front_9      | 7154.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_front_10     | 7350.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_front_11     | 7154.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_front_12     | 7350.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_front_13     | 7154.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_front_14     | 7350.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_front_15     | 7350.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_1       | 8428.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_2       | 8232.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_3       | 8428.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_4       | 8428.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_5       | 8428.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_6       | 8624.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_7       | 8624.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_8       | 8624.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_9       | 8428.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_10      | 8428.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_11      | 8624.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_12      | 8820.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_13      | 8624.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_14      | 8624.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
Fan_rear_15      | 8624.000   | RPM       | ok    | na       | 1764.000  | 1960.000  | na       | na
  | na
PSU1_54V_Fan_Spe | 8968.000   | RPM       | ok    | na       | na        | na        | na       | na
  | na
PSU2_54V_Fan_Spe | 8968.000   | RPM       | ok    | na       | na        | na        | na       | na
  | na
PSU3_54V_Fan_Spe | 8968.000   | RPM       | ok    | na       | na        | na        | na       | na
  | na
PSU4_54V_Fan_Spe | 8968.000   | RPM       | ok    | na       | na        | na        | na       | na
  | na
PSU5_54V_Fan_Spe | 8968.000   | RPM       | ok    | na       | na        | na        | na       | na
  | na
PSU6_54V_Fan_Spe | 8968.000   | RPM       | ok    | na       | na        | na        | na       | na
  | na
PSU7_Fan_Speed_1 | 10030.000  | RPM       | ok    | na       | na        | na        | na       | na
  | na
PSU8_Fan_Speed_1 | 10148.000  | RPM       | ok    | na       | na        | na        | na       | na
  | na
CPU0_FIVRA_Pout  | 82.600     | Watts     | ok    | na       | na        | na        | na       | na
  | na
```

```
CPU0_PVCCD_Pin   | 0.000      | Watts      | ok    | na        | na        | na        | na        | na
  | na
CPU0_PVCCIN_Pout | 3009.000   | Watts      | ok    | na        | na        | na        | na        | na
  | na
CPU1_FIVRA_Pout  | 94.400     | Watts      | ok    | na        | na        | na        | na        | na
  | na
CPU1_PVCCD_Pin   | 0.000      | Watts      | ok    | na        | na        | na        | na        | na
  | na
CPU1_PVCCIN_Pout | 2065.000   | Watts      | ok    | na        | na        | na        | na        | na
  | na
FAON_CPU0_Pout   | 2053.200   | Watts      | ok    | na        | na        | na        | na        | na
  | na
FAON_CPU1_Pout   | 2053.200   | Watts      | ok    | na        | na        | na        | na        | na
  | na
PSU1_54V_Pin     | 224.200    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU1_54V_Pout    | 188.800    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU2_54V_Pin     | 259.600    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU2_54V_Pout    | 212.400    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU3_54V_Pin     | 271.400    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU3_54V_Pout    | 236.000    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU4_54V_Pin     | 236.000    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU4_54V_Pout    | 212.400    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU5_54V_Pin     | 212.400    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU5_54V_Pout    | 177.000    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU6_54V_Pin     | 224.200    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU6_54V_Pout    | 188.800    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU7_Pin         | 177.000    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU7_Pout        | 153.400    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU8_Pin         | 153.400    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU8_Pout        | 153.400    | Watts      | ok    | na        | na        | na        | 849.600   |
  896.800   | na
PSU_Power_Total  | 342.200    | Watts      | ok    | na        | na        | na        | 1699.200  |
  1805.400  | na
SW12_0V8_Pout    | 3009.000   | Watts      | ok    | na        | na        | na        | na        | na
  | na
SW34_0V8_Pout    | 3009.000   | Watts      | ok    | na        | na        | na        | na        | na
  | na
CPU0_FIVRA_Temp  | 38.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
CPU0_PVCCD_Temp  | 36.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
CPU0_PVCCIN_Temp | 41.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
CPU1_FIVRA_Temp  | 33.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
CPU1_PVCCD_Temp  | 32.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
CPU1_PVCCIN_Temp | 29.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
```

```
FAON_CPU0_Temp   | 39.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
FAON_CPU1_Temp   | 38.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
HSBP1_Temp       | 28.000     | degrees C  | ok    | na        | 7.000     | 12.000    | 52.000    |
  57.000    | na
HSBP2_Temp       | 27.000     | degrees C  | ok    | na        | 7.000     | 12.000    | 52.000    |
  57.000    | na
HSBP3_Temp       | 25.000     | degrees C  | ok    | na        | 7.000     | 12.000    | 52.000    |
  57.000    | na
LM75BD_MB        | 32.000     | degrees C  | ok    | na        | 0.000     | 5.000     | 110.000   |
  115.000   | na
P54V_Temp        | 57.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
PDB_Temp         | 26.000     | degrees C  | ok    | na        | 7.000     | 12.000    | 52.000    |
  57.000    | na
PSU1_54V_Temp    | 28.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
PSU2_54V_Temp    | 29.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
PSU3_54V_Temp    | 30.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
PSU4_54V_Temp    | 30.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
PSU5_54V_Temp    | 32.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
PSU6_54V_Temp    | 31.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
PSU7_Temp        | 25.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
PSU8_Temp        | 28.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
RISER1_Temp      | 28.000     | degrees C  | ok    | na        | 7.000     | 12.000    | 52.000    |
  57.000    | na
RISER2_Temp      | 27.000     | degrees C  | ok    | na        | 7.000     | 12.000    | 52.000    |
  57.000    | na
RISER3_Temp      | 26.000     | degrees C  | ok    | na        | 7.000     | 12.000    | 52.000    |
  57.000    | na
SW12_0V8_Temp    | 42.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
SW34_0V8_Temp    | 36.000     | degrees C  | ok    | na        | 5.000     | 10.000    | 60.000    |
  70.000    | na
SWB_Temp         | 34.000     | degrees C  | ok    | na        | 7.000     | 12.000    | 52.000    |
  57.000    | na
nvme0            | 30.876     | degrees C  | ok    | na        | 0.000     | 4.980     | 70.218    |
  80.178    | na
nvme1            | 30.876     | degrees C  | ok    | na        | 0.000     | 4.980     | 70.218    |
  80.178    | na
nvme2            | 31.872     | degrees C  | ok    | na        | 0.000     | 4.980     | 70.218    |
  80.178    | na
nvme3            | 31.872     | degrees C  | ok    | na        | 0.000     | 4.980     | 70.218    |
  80.178    | na
nvme4            | 31.872     | degrees C  | ok    | na        | 0.000     | 4.980     | 70.218    |
  80.178    | na
nvme5            | 31.872     | degrees C  | ok    | na        | 0.000     | 4.980     | 70.218    |
  80.178    | na
nvme6            | 30.876     | degrees C  | ok    | na        | 0.000     | 4.980     | 70.218    |
  80.178    | na
nvme7            | 31.872     | degrees C  | ok    | na        | 0.000     | 4.980     | 70.218    |
  80.178    | na
CPU0_FIVRA_Vout  | 1.794      | Volts      | ok    | na        | 1.248     | 1.404     | 2.184     | 2.340
  | na
CPU0_PVCCD_Vout  | 1.170      | Volts      | ok    | na        | 0.780     | 0.858     | 1.326     | 1.404
  | na
```

```
CPU0_PVCCIN_Vin  | 12.168     | Volts      | ok    | na        | 8.424     | 9.594     | 14.430    |
  15.600    | na
CPU0_PVCCIN_Vout | 1.638      | Volts      | ok    | na        | 1.248     | 1.404     | 2.184     | 2.340
  | na
CPU1_FIVRA_Vout  | 1.794      | Volts      | ok    | na        | 1.248     | 1.404     | 2.184     | 2.340
  | na
CPU1_PVCCD_Vout  | 1.170      | Volts      | ok    | na        | 0.780     | 0.858     | 1.326     | 1.404
  | na
CPU1_PVCCIN_Vin  | 12.090     | Volts      | ok    | na        | 8.424     | 9.594     | 14.430    |
  15.600    | na
CPU1_PVCCIN_Vout | 1.638      | Volts      | ok    | na        | 1.248     | 1.404     | 2.184     | 2.340
  | na
EHV_CPU0_Vout    | 1.794      | Volts      | ok    | na        | 1.248     | 1.404     | 2.184     | 2.340
  | na
EHV_CPU1_Vout    | 1.794      | Volts      | ok    | na        | 1.248     | 1.404     | 2.184     | 2.340
  | na
FAON_CPU0_Vout   | 1.092      | Volts      | ok    | na        | 0.702     | 0.780     | 1.170     | 1.326
  | na
FAON_CPU1_Vout   | 1.092      | Volts      | ok    | na        | 0.702     | 0.780     | 1.170     | 1.326
  | na
P54V_Vout        | 12.012     | Volts      | cr    | na        | 1.248     | 1.404     | 2.184     | 2.340
  | na
PSU1_54V_Vin     | 227.740    | Volts      | ok    | na        | 88.500    | 89.680    | 240.720   |
  264.320   | na
PSU1_54V_Vout    | 19.890     | Volts      | cr    | na        | 10.296    | 11.388    | 12.636    |
  13.962    | na
PSU2_54V_Vin     | 227.740    | Volts      | ok    | na        | 88.500    | 89.680    | 240.720   |
  264.320   | na
PSU2_54V_Vout    | 19.890     | Volts      | cr    | na        | 10.296    | 11.388    | 12.636    |
  13.962    | na
PSU3_54V_Vin     | 227.740    | Volts      | ok    | na        | 88.500    | 89.680    | 240.720   |
  264.320   | na
PSU3_54V_Vout    | 19.890     | Volts      | cr    | na        | 10.296    | 11.388    | 12.636    |
  13.962    | na
PSU4_54V_Vin     | 227.740    | Volts      | ok    | na        | 88.500    | 89.680    | 240.720   |
  264.320   | na
PSU4_54V_Vout    | 19.890     | Volts      | cr    | na        | 10.296    | 11.388    | 12.636    |
  13.962    | na
PSU5_54V_Vin     | 228.920    | Volts      | ok    | na        | 88.500    | 89.680    | 240.720   |
  264.320   | na
PSU5_54V_Vout    | 19.890     | Volts      | cr    | na        | 10.296    | 11.388    | 12.636    |
  13.962    | na
PSU6_54V_Vin     | 228.920    | Volts      | ok    | na        | 88.500    | 89.680    | 240.720   |
  264.320   | na
PSU6_54V_Vout    | 19.890     | Volts      | cr    | na        | 10.296    | 11.388    | 12.636    |
  13.962    | na
PSU7_Vin         | 230.100    | Volts      | ok    | na        | 88.500    | 89.680    | 240.720   |
  264.320   | na
PSU7_Vout        | 12.324     | Volts      | ok    | na        | 10.296    | 11.388    | 12.636    |
  13.962    | na
PSU8_Vin         | 228.920    | Volts      | ok    | na        | 88.500    | 89.680    | 240.720   |
  264.320   | na
PSU8_Vout        | 12.168     | Volts      | ok    | na        | 10.296    | 11.388    | 12.636    |
  13.962    | na
PVNN_MAIN_CPU0   | 1.009      | Volts      | ok    | na        | 0.696     | 0.970     | 1.029     | 1.303
  | na
PVNN_MAIN_CPU1   | 1.009      | Volts      | ok    | na        | 0.696     | 0.970     | 1.029     | 1.303
  | na
PVNN_PCH         | 0.902      | Volts      | ok    | na        | 0.627     | 0.853     | 0.941     | 1.166
  | na
PVPP_HBM_CPU0    | 0.000      | Volts      | cr    | na        | 8.428     | 2.548     | 2.744     |
  15.582    | na
PVPP_HBM_CPU1    | 0.000      | Volts      | cr    | na        | 8.428     | 2.548     | 2.744     |
  15.582    | na
```

```
SW12_0V8_Vin      | 12.090     | Volts      | ok    | na        | 8.424     | 9.594     | 14.430    |
  15.600    | na
SW12_0V8_Vout     | 0.780      | Volts      | cr    | na        | 1.248     | 1.404     | 2.184     | 2.340
  | na
SW34_0V8_Vin      | 12.090     | Volts      | ok    | na        | 8.424     | 9.594     | 14.430    |
  15.600    | na
SW34_0V8_Vout     | 0.780      | Volts      | cr    | na        | 1.248     | 1.404     | 2.184     | 2.340
  | na
VBATT             | 2.862      | Volts      | ok    | na        | 2.097     | 2.391     | 3.606     | 3.900
  | na
VCC1V05_PCH       | 1.049      | Volts      | ok    | na        | 0.735     | 1.000     | 1.078     | 1.362
  | na
VCC1V8_PCH        | 1.823      | Volts      | ok    | na        | 1.254     | 1.725     | 1.862     | 2.332
  | na
VCC3V3            | 3.312      | Volts      | ok    | na        | 2.313     | 3.136     | 3.469     | 4.292
  | na
VCC3V3_RISER      | 3.312      | Volts      | ok    | na        | 2.313     | 2.999     | 3.606     | 4.292
  | na
VCC3V3_SB         | 3.312      | Volts      | ok    | na        | 2.313     | 3.214     | 3.391     | 4.292
  | na
VCC5V             | 5.130      | Volts      | ok    | na        | 3.510     | 4.752     | 5.238     | 6.480
  | na
VCC12V_CPU0_DIMM  | 12.838     | Volts      | ok    | na        | 8.428     | 4.214     | 14.994    |
  15.582    | na
VCC12V_CPU1_DIMM  | 12.838     | Volts      | ok    | na        | 8.428     | 4.214     | 14.994    |
  15.582    | na
VCC12V_HSBP       | 12.740     | Volts      | nc    | na        | 8.428     | 11.074    | 12.642    |
  15.582    | na
VCC12V_RISER      | 12.936     | Volts      | nc    | na        | 8.428     | 11.074    | 12.642    |
  15.582    | na
```

lm-sensors

sensors: After installing the SynapseAI software stack and driver, you can show all Intel® Gaudi® 2 AI Accelerator temperature and power information by the **lm-sensor** tool. The tools also display other HW temperature and power information that are supported by the Linux kernel.

```
root@ags8200:/home/ubuntu# sensors
HL225-pci-d300
Adapter: PCI adapter
ADC 54V:           54.09 V  (highest = +54.50 V)
On Chip 0:         +30.4°C  (crit = +93.0°C, highest = +30.7°C)
On Chip 1:         +30.3°C  (crit = +93.0°C, highest = +30.6°C)
On Chip 2:         +30.3°C  (crit = +93.0°C, highest = +30.5°C)
On Chip 3:         +29.7°C  (crit = +93.0°C, highest = +29.9°C)
HBM 0 Temp:        +30.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 1 Temp:        +30.0°C  (crit = +125.0°C, highest = +35.0°C)
HBM 2 Temp:        +32.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 3 Temp:        +30.0°C  (crit = +125.0°C, highest = +37.0°C)
HBM 4 Temp:        +30.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 5 Temp:        +31.0°C  (crit = +125.0°C, highest = +35.0°C)
On Chip TD 0:      +30.1°C  (crit = +88.0°C, highest = +30.1°C)
On Chip TD 1:      +30.2°C  (crit = +88.0°C, highest = +30.2°C)
On Chip TD 2:      +30.2°C  (crit = +88.0°C, highest = +30.2°C)
On Chip TD 3:      +29.9°C  (crit = +88.0°C, highest = +30.1°C)
On Board 2 Top:    +31.7°C  (crit = +125.0°C, highest = +31.7°C)
On Board 01 Top:   +30.5°C  (crit = +80.0°C, highest = +30.5°C)
```

```
On Board 01 Bot:   +31.0°C  (crit = +80.0°C, highest = +31.0°C)
On Board 23 Top:   +31.0°C  (crit = +80.0°C, highest = +31.0°C)
On Board 23 Bot:   +32.0°C  (crit = +80.0°C, highest = +32.0°C)
CPLD Temp:         +34.8°C  (crit = +90.0°C, highest = +34.8°C)
VRM1 Temp:         +31.0°C  (highest = +31.0°C)
VRM2 Temp:         +39.0°C  (highest = +39.0°C)
54V Power Draw:    57.12 W  (highest = 104.80 W)
ADC 12V1 Current:   2.54 A  (highest =  +2.78 A)

HL225-pci-4f00
Adapter: PCI adapter
ADC 54V:           54.28 V  (highest = +54.69 V)
On Chip 0:         +29.9°C  (crit = +93.0°C, highest = +30.1°C)
On Chip 1:         +30.1°C  (crit = +93.0°C, highest = +30.4°C)
On Chip 2:         +31.0°C  (crit = +93.0°C, highest = +31.0°C)
On Chip 3:         +29.9°C  (crit = +93.0°C, highest = +30.1°C)
HBM 0 Temp:        +30.0°C  (crit = +125.0°C, highest = +35.0°C)
HBM 1 Temp:        +31.0°C  (crit = +125.0°C, highest = +35.0°C)
HBM 2 Temp:        +30.0°C  (crit = +125.0°C, highest = +35.0°C)
HBM 3 Temp:        +31.0°C  (crit = +125.0°C, highest = +36.0°C)
HBM 4 Temp:        +31.0°C  (crit = +125.0°C, highest = +40.0°C)
HBM 5 Temp:        +32.0°C  (crit = +125.0°C, highest = +36.0°C)
On Chip TD 0:      +30.2°C  (crit = +88.0°C, highest = +30.2°C)
On Chip TD 1:      +30.8°C  (crit = +88.0°C, highest = +30.8°C)
On Chip TD 2:      +30.9°C  (crit = +88.0°C, highest = +30.9°C)
On Chip TD 3:      +30.2°C  (crit = +88.0°C, highest = +30.3°C)
On Board 2 Top:    +31.0°C  (crit = +125.0°C, highest = +31.0°C)
On Board 01 Top:   +30.5°C  (crit = +80.0°C, highest = +30.5°C)
On Board 01 Bot:   +31.5°C  (crit = +80.0°C, highest = +31.5°C)
On Board 23 Top:   +30.5°C  (crit = +80.0°C, highest = +30.5°C)
On Board 23 Bot:   +31.5°C  (crit = +80.0°C, highest = +31.5°C)
CPLD Temp:         +34.0°C  (crit = +90.0°C, highest = +34.0°C)
VRM1 Temp:         +33.0°C  (highest = +33.0°C)
VRM2 Temp:         +38.0°C  (highest = +38.0°C)
54V Power Draw:    54.42 W  (highest = 105.34 W)
ADC 12V1 Current:   2.33 A  (highest =  +2.57 A)

power_meter-acpi-0
Adapter: ACPI interface
power1:      350.00 W  (interval =   1.00 s)

HL225-pci-a400
Adapter: PCI adapter
ADC 54V:           54.21 V  (highest = +54.62 V)
On Chip 0:         +29.8°C  (crit = +93.0°C, highest = +30.1°C)
On Chip 1:         +29.3°C  (crit = +93.0°C, highest = +29.6°C)
On Chip 2:         +30.3°C  (crit = +93.0°C, highest = +30.5°C)
On Chip 3:         +29.6°C  (crit = +93.0°C, highest = +29.9°C)
HBM 0 Temp:        +30.0°C  (crit = +125.0°C, highest = +35.0°C)
HBM 1 Temp:        +31.0°C  (crit = +125.0°C, highest = +37.0°C)
HBM 2 Temp:        +29.0°C  (crit = +125.0°C, highest = +37.0°C)
HBM 3 Temp:        +29.0°C  (crit = +125.0°C, highest = +36.0°C)
HBM 4 Temp:        +31.0°C  (crit = +125.0°C, highest = +39.0°C)
HBM 5 Temp:        +29.0°C  (crit = +125.0°C, highest = +38.0°C)
On Chip TD 0:      +30.1°C  (crit = +88.0°C, highest = +30.1°C)
On Chip TD 1:      +30.6°C  (crit = +88.0°C, highest = +30.6°C)
On Chip TD 2:      +30.6°C  (crit = +88.0°C, highest = +30.6°C)
On Chip TD 3:      +30.2°C  (crit = +88.0°C, highest = +30.2°C)
On Board 2 Top:    +31.4°C  (crit = +125.0°C, highest = +31.4°C)
On Board 01 Top:   +31.0°C  (crit = +80.0°C, highest = +31.0°C)
On Board 01 Bot:   +31.5°C  (crit = +80.0°C, highest = +31.5°C)
On Board 23 Top:   +30.5°C  (crit = +80.0°C, highest = +31.0°C)
On Board 23 Bot:   +31.5°C  (crit = +80.0°C, highest = +31.5°C)
CPLD Temp:         +34.8°C  (crit = +90.0°C, highest = +34.8°C)
VRM1 Temp:         +32.0°C  (highest = +32.0°C)
```

```
VRM2 Temp:         +39.0°C  (highest = +39.0°C)
54V Power Draw:    70.79 W  (highest = 121.58 W)
ADC 12V1 Current:   2.33 A  (highest =  +2.75 A)

nvme-pci-ce00
Adapter: PCI adapter
Composite:    +31.9°C  (low  = -273.1°C, high = +79.8°C)
                       (crit = +82.8°C)
Sensor 1:     +31.9°C  (low  = -273.1°C, high = +65261.8°C)
Sensor 2:     +40.9°C  (low  = -273.1°C, high = +65261.8°C)

HL225-pci-2300
Adapter: PCI adapter
ADC 54V:           54.15 V  (highest = +54.59 V)
On Chip 0:         +33.5°C  (crit = +93.0°C, highest = +33.8°C)
On Chip 1:         +33.6°C  (crit = +93.0°C, highest = +33.6°C)
On Chip 2:         +33.7°C  (crit = +93.0°C, highest = +34.0°C)
On Chip 3:         +33.6°C  (crit = +93.0°C, highest = +33.6°C)
HBM 0 Temp:        +34.0°C  (crit = +125.0°C, highest = +37.0°C)
HBM 1 Temp:        +35.0°C  (crit = +125.0°C, highest = +37.0°C)
HBM 2 Temp:        +32.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 3 Temp:        +36.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 4 Temp:        +34.0°C  (crit = +125.0°C, highest = +35.0°C)
HBM 5 Temp:        +35.0°C  (crit = +125.0°C, highest = +38.0°C)
On Chip TD 0:      +34.2°C  (crit = +88.0°C, highest = +34.2°C)
On Chip TD 1:      +34.1°C  (crit = +88.0°C, highest = +34.1°C)
On Chip TD 2:      +34.2°C  (crit = +88.0°C, highest = +34.2°C)
On Chip TD 3:      +33.9°C  (crit = +88.0°C, highest = +33.9°C)
On Board 2 Top:    +34.4°C  (crit = +125.0°C, highest = +34.4°C)
On Board 01 Top:   +34.0°C  (crit = +80.0°C, highest = +34.5°C)
On Board 01 Bot:   +35.0°C  (crit = +80.0°C, highest = +35.0°C)
On Board 23 Top:   +34.0°C  (crit = +80.0°C, highest = +34.0°C)
On Board 23 Bot:   +34.5°C  (crit = +80.0°C, highest = +34.5°C)
CPLD Temp:         +38.5°C  (crit = +90.0°C, highest = +38.5°C)
VRM1 Temp:         +36.0°C  (highest = +36.0°C)
VRM2 Temp:         +42.0°C  (highest = +42.0°C)
54V Power Draw:    55.31 W  (highest = 104.86 W)
ADC 12V1 Current:   2.32 A  (highest =  +2.71 A)

nvme-pci-a200
Adapter: PCI adapter
Composite:    +31.9°C  (low  = -273.1°C, high = +79.8°C)
                       (crit = +82.8°C)
Sensor 1:     +31.9°C  (low  = -273.1°C, high = +65261.8°C)
Sensor 2:     +39.9°C  (low  = -273.1°C, high = +65261.8°C)

nvme-pci-4d00
Adapter: PCI adapter
Composite:    +31.9°C  (low  = -273.1°C, high = +79.8°C)
                       (crit = +82.8°C)
Sensor 1:     +31.9°C  (low  = -273.1°C, high = +65261.8°C)
Sensor 2:     +40.9°C  (low  = -273.1°C, high = +65261.8°C)

coretemp-isa-0000
Adapter: ISA adapter
Package id 0:  +30.0°C  (high = +90.0°C, crit = +98.0°C)
Core 0:        +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 1:        +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 2:        +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 3:        +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 4:        +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 5:        +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 6:        +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 7:        +28.0°C  (high = +90.0°C, crit = +98.0°C)
Core 8:        +27.0°C  (high = +90.0°C, crit = +98.0°C)
```

```
Core 9:        +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 10:       +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 11:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 12:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 13:       +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 14:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 15:       +30.0°C  (high = +90.0°C, crit = +98.0°C)
Core 16:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 17:       +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 18:       +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 19:       +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 20:       +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 21:       +28.0°C  (high = +90.0°C, crit = +98.0°C)
Core 22:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 23:       +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 24:       +28.0°C  (high = +90.0°C, crit = +98.0°C)
Core 25:       +24.0°C  (high = +90.0°C, crit = +98.0°C)
Core 26:       +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 27:       +24.0°C  (high = +90.0°C, crit = +98.0°C)
Core 28:       +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 29:       +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 30:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 31:       +25.0°C  (high = +90.0°C, crit = +98.0°C)

nvme-pci-1e00
Adapter: PCI adapter
Composite:    +30.9°C  (low  = -273.1°C, high = +79.8°C)
                       (crit = +82.8°C)
Sensor 1:     +30.9°C  (low  = -273.1°C, high = +65261.8°C)
Sensor 2:     +39.9°C  (low  = -273.1°C, high = +65261.8°C)

HL225-pci-a700
Adapter: PCI adapter
ADC 54V:           54.09 V  (highest = +54.50 V)
On Chip 0:         +33.7°C  (crit = +93.0°C, highest = +33.9°C)
On Chip 1:         +33.6°C  (crit = +93.0°C, highest = +33.9°C)
On Chip 2:         +33.6°C  (crit = +93.0°C, highest = +33.6°C)
On Chip 3:         +33.5°C  (crit = +93.0°C, highest = +33.5°C)
HBM 0 Temp:        +33.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 1 Temp:        +32.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 2 Temp:        +34.0°C  (crit = +125.0°C, highest = +36.0°C)
HBM 3 Temp:        +34.0°C  (crit = +125.0°C, highest = +39.0°C)
HBM 4 Temp:        +33.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 5 Temp:        +36.0°C  (crit = +125.0°C, highest = +40.0°C)
On Chip TD 0:      +33.2°C  (crit = +88.0°C, highest = +33.2°C)
On Chip TD 1:      +33.8°C  (crit = +88.0°C, highest = +33.8°C)
On Chip TD 2:      +33.2°C  (crit = +88.0°C, highest = +33.2°C)
On Chip TD 3:      +33.1°C  (crit = +88.0°C, highest = +33.1°C)
On Board 2 Top:    +34.3°C  (crit = +125.0°C, highest = +34.3°C)
On Board 01 Top:   +34.0°C  (crit = +80.0°C, highest = +34.0°C)
On Board 01 Bot:   +35.0°C  (crit = +80.0°C, highest = +35.0°C)
On Board 23 Top:   +34.0°C  (crit = +80.0°C, highest = +34.0°C)
On Board 23 Bot:   +34.5°C  (crit = +80.0°C, highest = +34.5°C)
CPLD Temp:         +36.5°C  (crit = +90.0°C, highest = +36.5°C)
VRM1 Temp:         +36.0°C  (highest = +36.0°C)
VRM2 Temp:         +42.0°C  (highest = +42.0°C)
54V Power Draw:    62.10 W  (highest = 113.47 W)
ADC 12V1 Current:   2.39 A  (highest =  +2.68 A)

HL225-pci-5200
Adapter: PCI adapter
ADC 54V:           54.12 V  (highest = +54.56 V)
On Chip 0:         +34.2°C  (crit = +93.0°C, highest = +34.2°C)
On Chip 1:         +34.0°C  (crit = +93.0°C, highest = +34.0°C)
On Chip 2:         +34.1°C  (crit = +93.0°C, highest = +34.1°C)
```

```
On Chip 3:         +33.6°C  (crit = +93.0°C, highest = +33.9°C)
HBM 0 Temp:        +36.0°C  (crit = +125.0°C, highest = +39.0°C)
HBM 1 Temp:        +33.0°C  (crit = +125.0°C, highest = +37.0°C)
HBM 2 Temp:        +34.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 3 Temp:        +32.0°C  (crit = +125.0°C, highest = +36.0°C)
HBM 4 Temp:        +34.0°C  (crit = +125.0°C, highest = +36.0°C)
HBM 5 Temp:        +34.0°C  (crit = +125.0°C, highest = +36.0°C)
On Chip TD 0:      +33.4°C  (crit = +88.0°C, highest = +33.4°C)
On Chip TD 1:      +33.7°C  (crit = +88.0°C, highest = +33.8°C)
On Chip TD 2:      +33.4°C  (crit = +88.0°C, highest = +33.4°C)
On Chip TD 3:      +33.2°C  (crit = +88.0°C, highest = +33.2°C)
On Board 2 Top:    +34.1°C  (crit = +125.0°C, highest = +34.1°C)
On Board 01 Top:   +33.5°C  (crit = +80.0°C, highest = +33.5°C)
On Board 01 Bot:   +34.5°C  (crit = +80.0°C, highest = +35.0°C)
On Board 23 Top:   +33.5°C  (crit = +80.0°C, highest = +33.5°C)
On Board 23 Bot:   +34.5°C  (crit = +80.0°C, highest = +34.5°C)
CPLD Temp:         +35.8°C  (crit = +90.0°C, highest = +36.0°C)
VRM1 Temp:         +36.0°C  (highest = +36.0°C)
VRM2 Temp:         +42.0°C  (highest = +42.0°C)
54V Power Draw:    56.39 W  (highest = 109.32 W)
ADC 12V1 Current:   2.38 A  (highest =  +2.84 A)

HL225-pci-d000
Adapter: PCI adapter
ADC 54V:           54.21 V  (highest = +54.62 V)
On Chip 0:         +33.6°C  (crit = +93.0°C, highest = +33.9°C)
On Chip 1:         +33.4°C  (crit = +93.0°C, highest = +33.7°C)
On Chip 2:         +33.4°C  (crit = +93.0°C, highest = +33.4°C)
On Chip 3:         +33.1°C  (crit = +93.0°C, highest = +33.1°C)
HBM 0 Temp:        +33.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 1 Temp:        +33.0°C  (crit = +125.0°C, highest = +35.0°C)
HBM 2 Temp:        +32.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 3 Temp:        +33.0°C  (crit = +125.0°C, highest = +34.0°C)
HBM 4 Temp:        +32.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 5 Temp:        +34.0°C  (crit = +125.0°C, highest = +38.0°C)
On Chip TD 0:      +32.6°C  (crit = +88.0°C, highest = +32.6°C)
On Chip TD 1:      +33.1°C  (crit = +88.0°C, highest = +33.1°C)
On Chip TD 2:      +32.3°C  (crit = +88.0°C, highest = +32.3°C)
On Chip TD 3:      +32.2°C  (crit = +88.0°C, highest = +32.2°C)
On Board 2 Top:    +33.9°C  (crit = +125.0°C, highest = +33.9°C)
On Board 01 Top:   +33.0°C  (crit = +80.0°C, highest = +33.0°C)
On Board 01 Bot:   +34.0°C  (crit = +80.0°C, highest = +34.0°C)
On Board 23 Top:   +33.0°C  (crit = +80.0°C, highest = +33.0°C)
On Board 23 Bot:   +33.5°C  (crit = +80.0°C, highest = +33.5°C)
CPLD Temp:         +36.2°C  (crit = +90.0°C, highest = +36.2°C)
VRM1 Temp:         +36.0°C  (highest = +36.0°C)
VRM2 Temp:         +41.0°C  (highest = +41.0°C)
54V Power Draw:    66.13 W  (highest = 116.75 W)
ADC 12V1 Current:   2.43 A  (highest =  +2.68 A)

nvme-pci-cf00
Adapter: PCI adapter
Composite:    +31.9°C  (low  = -273.1°C, high = +79.8°C)
                       (crit = +82.8°C)
Sensor 1:     +31.9°C  (low  = -273.1°C, high = +65261.8°C)
Sensor 2:     +40.9°C  (low  = -273.1°C, high = +65261.8°C)

HL225-pci-2000
Adapter: PCI adapter
ADC 54V:           54.31 V  (highest = +54.72 V)
On Chip 0:         +31.3°C  (crit = +93.0°C, highest = +31.6°C)
On Chip 1:         +31.4°C  (crit = +93.0°C, highest = +31.6°C)
On Chip 2:         +31.3°C  (crit = +93.0°C, highest = +31.6°C)
On Chip 3:         +31.1°C  (crit = +93.0°C, highest = +31.4°C)
HBM 0 Temp:        +33.0°C  (crit = +125.0°C, highest = +40.0°C)
```

```
HBM 1 Temp:        +31.0°C  (crit = +125.0°C, highest = +39.0°C)
HBM 2 Temp:        +34.0°C  (crit = +125.0°C, highest = +38.0°C)
HBM 3 Temp:        +31.0°C  (crit = +125.0°C, highest = +36.0°C)
HBM 4 Temp:        +31.0°C  (crit = +125.0°C, highest = +39.0°C)
HBM 5 Temp:        +31.0°C  (crit = +125.0°C, highest = +38.0°C)
On Chip TD 0:      +31.0°C  (crit = +88.0°C, highest = +31.0°C)
On Chip TD 1:      +31.4°C  (crit = +88.0°C, highest = +31.4°C)
On Chip TD 2:      +31.1°C  (crit = +88.0°C, highest = +31.2°C)
On Chip TD 3:      +31.1°C  (crit = +88.0°C, highest = +31.1°C)
On Board 2 Top:    +31.3°C  (crit = +125.0°C, highest = +31.3°C)
On Board 01 Top:   +30.5°C  (crit = +80.0°C, highest = +30.5°C)
On Board 01 Bot:   +31.5°C  (crit = +80.0°C, highest = +31.5°C)
On Board 23 Top:   +31.0°C  (crit = +80.0°C, highest = +31.0°C)
On Board 23 Bot:   +31.5°C  (crit = +80.0°C, highest = +31.5°C)
CPLD Temp:         +33.2°C  (crit = +90.0°C, highest = +33.2°C)
VRM1 Temp:         +32.0°C  (highest = +32.0°C)
VRM2 Temp:         +38.0°C  (highest = +39.0°C)
54V Power Draw:    39.48 W  (highest =  90.53 W)
ADC 12V1 Current:   2.61 A  (highest =  +2.78 A)

nvme-pci-a300
Adapter: PCI adapter
Composite:    +31.9°C  (low  = -273.1°C, high = +79.8°C)
                       (crit = +82.8°C)
Sensor 1:     +31.9°C  (low  = -273.1°C, high = +65261.8°C)
Sensor 2:     +40.9°C  (low  = -273.1°C, high = +65261.8°C)

nvme-pci-4e00
Adapter: PCI adapter
Composite:    +31.9°C  (low  = -273.1°C, high = +79.8°C)
                       (crit = +82.8°C)
Sensor 1:     +31.9°C  (low  = -273.1°C, high = +65261.8°C)
Sensor 2:     +39.9°C  (low  = -273.1°C, high = +65261.8°C)

coretemp-isa-0001
Adapter: ISA adapter
Package id 1:  +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 0:        +24.0°C  (high = +90.0°C, crit = +98.0°C)
Core 1:        +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 2:        +21.0°C  (high = +90.0°C, crit = +98.0°C)
Core 3:        +24.0°C  (high = +90.0°C, crit = +98.0°C)
Core 4:        +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 5:        +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 6:        +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 7:        +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 8:        +23.0°C  (high = +90.0°C, crit = +98.0°C)
Core 9:        +23.0°C  (high = +90.0°C, crit = +98.0°C)
Core 10:       +23.0°C  (high = +90.0°C, crit = +98.0°C)
Core 11:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 12:       +23.0°C  (high = +90.0°C, crit = +98.0°C)
Core 13:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 14:       +21.0°C  (high = +90.0°C, crit = +98.0°C)
Core 15:       +24.0°C  (high = +90.0°C, crit = +98.0°C)
Core 16:       +27.0°C  (high = +90.0°C, crit = +98.0°C)
Core 17:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 18:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 19:       +23.0°C  (high = +90.0°C, crit = +98.0°C)
Core 20:       +22.0°C  (high = +90.0°C, crit = +98.0°C)
Core 21:       +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 22:       +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 23:       +23.0°C  (high = +90.0°C, crit = +98.0°C)
Core 24:       +24.0°C  (high = +90.0°C, crit = +98.0°C)
Core 25:       +26.0°C  (high = +90.0°C, crit = +98.0°C)
Core 26:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
Core 27:       +25.0°C  (high = +90.0°C, crit = +98.0°C)
```

```
Core 28:       +23.0°C  (high = +90.0°C, crit = +98.0°C)
Core 29:       +24.0°C  (high = +90.0°C, crit = +98.0°C)
Core 30:       +22.0°C  (high = +90.0°C, crit = +98.0°C)
Core 31:       +23.0°C  (high = +90.0°C, crit = +98.0°C)

nvme-pci-1f00
Adapter: PCI adapter
Composite:    +30.9°C  (low  = -273.1°C, high = +79.8°C)
                       (crit = +82.8°C)
Sensor 1:     +30.9°C  (low  = -273.1°C, high = +65261.8°C)
Sensor 2:     +38.9°C  (low  = -273.1°C, high = +65261.8°C)
```

Manage Network Interfaces Script

manage_network_ifs.sh: This script can be used as reference to bring up, take down, set IPs, unset IPs, and check the status of the Intel® Gaudi® 2 AI Accelerator network interfaces.

```
root@ags8200:/home/ubuntu# /opt/habanalabs/qual/gaudi2/bin/
  manage_network_ifs.sh --status
accel0
3 ports up (8, 22, 23)
accel1
3 ports up (8, 22, 23)
accel2
3 ports up (8, 22, 23)
accel3
3 ports up (8, 22, 23)
accel4
3 ports up (8, 22, 23)
accel5
3 ports up (8, 22, 23)
accel6
3 ports up (8, 22, 23)
accel7
3 ports up (8, 22, 23)
```

For detailed information, please refer to manage_network_ifs.sh. You must link up all ports of all Intel® Gaudi® 2 AI Accelerators to do the scale-out training or the network test.

4 System Firmware Update

This chapter provides the system firmware update process. The latest firmware will be made available to fix various issues. You can follow the procedures to check the firmware version and update the firmware.

This chapter includes the following sections:

BIOS Update

This section shows how to check the BIOS version and update the firmware.

Check Version

Under CPU system:

```
dmidecode -t bios
```

Update

Under BMC:

```
# Update primary BIOS image
fwupd_util.sh bios primary <bios_file.bin>

# Update secondary BIOS image
fwupd_util.sh bios secondary <bios_file.bin>
```

Apply

Reboot the CPU system. The first boot-up after updating the BIOS takes more time. It may take a few minutes.

Under CPU system:

```
reboot
```

BMC Update

This section shows how to check the BMC version and update the firmware.

Check Version

Under BMC:

```
cat /etc/os-release
```

Update

Under BMC:

```
# transfer the BMC image file to BMC /tmp/ first

# Update primary BMC image
fwupd_util.sh bmc primary /tmp/<bmc.bin>

# Update secondary BMC image
fwupd_util.sh bmc secondary /tmp/<bmc.bin>
```

Apply

Update primary BMC image will reboot automatically.

The secondary BMC image will only work if the primary BMC image is damaged.

Main CPLD Update

This section shows how to check the main CPLD version and update the firmware.

Check Version

Under BMC:

```
show_version.sh
```

Update

Under BMC:

```
gpioset 0 115=1
gpioset 0 117=0
gpioset 0 203=0
gpioset 0 140=0
fwupd_util.sh cpld system <cpld_file.vme>
```

Apply

Before applying a CPLD update, it is suggested that the CPU system be powered off to avoid disk data loss.

Under BMC:

```
power_util.sh chassis cycle
```

PCIE CPLD Update

This section shows how to check the PCIE CPLD version and update the firmware.

Check Version

Under BMC:

```
show_version.sh
```

Update

Under BMC:

```
# The user must set GPIO before updating, the config is under BMC
# Now user can update PCIE CPLD from BMC
gpioset 0 115=1
gpioset 0 117=0
gpioset 0 203=0
gpioset 0 140=0

ispvm syscpld <cpld_file.vme>
```

Apply

Before applying a PCIE CPLD update, it is suggested that the CPU system be powered off to avoid disk data loss.

Under BMC:

```
power_util.sh chassis cycle
```

HSBP CPLD Update

This section shows how to check the HSBP CPLD version and update the firmware.

Check Version

Under BMC:

```
show_version.sh
```

Update

Under BMC:

```
# Update HSBP1 CPLD
fwupd_util.sh cpld HSBP1 <cpld_file.hex>

# Update HSBP2 CPLD
fwupd_util.sh cpld HSBP2 <cpld_file.hex>

# Update HSBP3 CPLD
fwupd_util.sh cpld HSBP3 <cpld_file.hex>
```

Apply

Before applying an HSBP CPLD update, it is suggested that the CPU system be powered off to avoid disk data loss.

Under BMC:

```
power_util.sh chassis cycle
```

5 Usage Examples

This chapter includes the following sections:

Network Environment Example

This section is used to introduce how to configure network topology for scale-out training.

Please refer to "Configuration" for scale-out configuration.

Figure 1: Network Topology

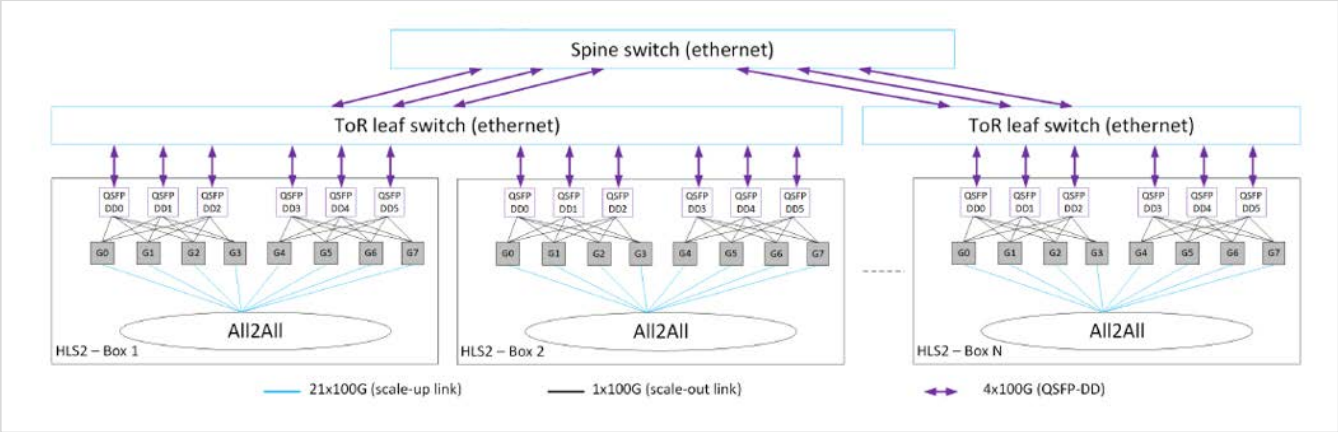

Each Intel® Gaudi® 2 AI Accelerator has 21 scale-up ports and 3 scale-out ports running at 100 Gbps.

The scale-up ports are used to connect to seven other Intel® Gaudi® 2 AI Accelerators inside a single AGS8200 box in an all-to-all manner.

When training with scale-out, please ensure that the configuration of each switch is set up before starting training.

Note: Link failure between the ToR leaf switch and any ports of Intel® Gaudi® 2 AI Accelerators will cause a task failure during the training task.

Link failure between leaf and spine will only reduce the performance.

Multi-Server Training Setup

For the scale-out learning, it is required to train the model in the Docker. Ensure that Docker is installed and set up, and then please refer to "Start Training with Docker".

To communicate with other devices via the network, the SSH public key needs to be shared among all devices with learning attendance.

Detailed configurations are described in PyTorch model and TensorFlow model.

Docker SSH Port Setup

By default, the Habana Docker uses port 22 for SSH. If the default port is occupied then the port number can also be changed.

Run the following commands to configure the selected port number, **port 3022** in the example below:

```
sed -i 's/#Port 22/Port 3022/g' /etc/ssh/sshd_config
sed -i 's/#PermitRootLogin prohibit-password/PermitRootLogin yes/' /etc/ssh/
  sshd_config
service ssh restart
```

Password-less SSH Setup

It is necessary to set up password-less SSH between all connected servers used in scale-out training. Follow the steps below:

Note: In the default environment, the SSH public key is configured to be shared with both devices, **habanas01** and **habanas02**.

1. Enter the following in all the servers' Docker sessions:

```
mkdir ~/.ssh
cd ~/.ssh
ssh-keygen -t rsa -b 4096
```

   a. Copy **id_rsa.pub** contents from every server's Docker to every other server's Docker's **~/.ssh/authorized_keys** (all public keys need to be in all hosts' authorized_keys):

```
cat id_rsa.pub > authorized_keys
vi authorized_keys
```

   b. Copy the contents from inside to other systems.

   c. Paste all hosts' public keys in all hosts' **authorized_keys** file.

2. On each system, add all hosts (including itself) to **known_hosts**. The IP addresses used below are just for illustration:

```
ssh-keyscan -p 3022 -H 10.3.124.124 >> ~/.ssh/known_hosts
ssh-keyscan -p 3022 -H 10.3.124.175 >> ~/.ssh/known_hosts
```

Start Training with Docker

In the Intel® Gaudi® 2 system, Docker is recommended for model training, and therefore this guide introduces how to access Docker.

The followng example is based on SynapseAI 1.14.0.

Note: The following steps need the root privilege. Please use **sudo** or gain root privilege for the following steps.

To install the SynapseAI driver or using another method without Docker, please refer to Habana Official Installation for detailed steps.

The following provides examples of how to enter the Docker of PyTorch or TensorFlow.

The learning examples used for PyTorch and TensorFlow can be download from Habana's model reference git.

These models are **ImageNet 2012** with Resnet structure that could be used to classify the object in the dataset.

Prerequisite: Prepare and Set Up Dataset

In the following examples, the ImageNet 2012 dataset needs to be organized as per PyTorch/TensorFlow requirements.

Please refer to Pytorch and TensorFlow reference model for data preprocessing.

Note: In our example, the ImageNet 2012 datasets are set up in advance in the file system with the path **/data**. These datasets are preprocessed, and therefore the models can be trained directly with these datasets. If there is one model that requires another dataset, it is necessary to set up the dataset manually.

Entering PyTorch Docker

```
docker run -v /data:/data -v /home/habana:/home/habana -v /root/.ssh:/root/
  .ssh -it --runtime=habana \
-e HABANA_VISIBLE_DEVICES=all -e
  OMPI_MCA_btl_vader_single_copy_mechanism=none \
-w /home/habana/Model-References/PyTorch/computer_vision/classification/
  torchvision \
--cap-add=sys_nice --net=host \
--ipc=host vault.habana.ai/gaudi-docker/1.14.0/ubuntu22.04/habanalabs/
  pytorch-installer-2.1.1:latest
```

Note: -v can mount a local host folder to the Docker, and the SSH information of the host, the model reference, and the dataset of imagenet are mounted with this command.

-e can setup the environment variable in Docker, and all the Intel® Gaudi® 2 AI Accelerators are setup with this command. If it is necessary to configure the device as visible, please refer to this official page.

-w can configure the initial position after accessing the Docker.

vault.habana.ai/gaudi-docker/1.14.0/ubuntu22.04/habanalabs/pytorch-installer-2.1.1:latest is the Docker image that preconfigures the PyTorch environment.

Entering TensorFlow Docker

```
docker run -v /data:/data -v /home/habana:/home/habana -v /root/.ssh:/root/
  .ssh -it --runtime=habana \
-e HABANA_VISIBLE_DEVICES=all -e
  OMPI_MCA_btl_vader_single_copy_mechanism=none \
-w /home/habana/Model-References/TensorFlow/computer_vision/Resnets/
  resnet_keras \
--cap-add=sys_nice --net=host \
--ipc=host vault.habana.ai/gaudi-docker/1.14.0/Ubuntu22.04/habanalabs/
  tensorflow-installer-tf-cpu-$2.15.0:latest
```

Note: -v can mount a local host folder to the Docker, and the SSH information of the host, the model reference, and the dataset of ImageNet are mounted with this command.

-e can setup the environment variable in Docker, and all the Intel® Gaudi® 2 AI Accelerators are setup with this command. If it is necessary to configure the device as visible, please refer to this official page.

-w can configure the initial position after accessing into the Docker.

vault.habana.ai/gaudi-docker/1.14.0/ubuntu22.04/habanalabs/tensorflow-installer-tf-cpu-2.15.0:latest is the Docker image that preconfigures the TensorFlow environment.

Start PyTorch Training Example

In this section, we train resnet50 with ImageNet on the PyTorch framework.

This model is built by Habana Model Reference.

Training in 1 Accelerator

ResNet50, lazy mode, BF16 mixed precision, batch size 256, custom learning rate, Habana dataloader (with hardware decode support on Gaudi® 2), 1 Intel® Gaudi® 2 AI Accelerator on a single server:

```
export PYTHON=/usr/bin/python3.10

$PYTHON -u train.py --dl-worker-type HABANA --batch-size 256 --model resnet50
  --device hpu \
--workers 8 --print-freq 20 --dl-time-exclude False --deterministic \
--data-path /data/pytorch/imagenet/ILSVRC2012 --epochs 90 --autocast  --lr
  0.1 \
--custom-lr-values 0.1 0.01 0.001 0.0001 --custom-lr-milestones 0 30 60 80
```

Scale-Up Training in 8 Accelerators

ResNet50, lazy mode, BF16 mixed precision, batch size 256, custom learning rate, 8 Intel® Gaudi® 2 AI Accelerators on a single server, print-frequency 1, and include dataloading time throughput computation:

```
export MASTER_ADDR=localhost
export MASTER_PORT=12355
export PYTHON=/usr/bin/python3.10

mpirun -n 8 --bind-to core --map-by socket:PE=6 --rank-by core --report-
  bindings --allow-run-as-root \
python train.py --data-path=/data/pytorch/imagenet/ILSVRC2012 --
  model=resnet50 --device=hpu \
--batch-size=256 --epochs=90 --print-freq=1 --output-dir=. --seed=123 --
  autocast \
--custom-lr-values 0.275 0.45 0.625 0.8 0.08 0.008 0.0008 --custom-lr-
  milestones 1 2 3 4 30 60 80 \
--deterministic --dl-time-exclude=False
```

Scale-Out Training in 2 Devices with 16 Accelerators

Run training on 16 Intel® Gaudi® 2 AI Accelerators.

```
export MASTER_ADDR=172.20.0.205
export MASTER_PORT=12355
export PYTHON=/usr/bin/python3.10
export HCCL_SOCKET_IFNAME=ens7f1np1

mpirun --allow-run-as-root --mca plm_rsh_args -p3022 --bind-to core --map-by
  ppr:4:socket:PE=6 -np 16 \
--mca btl_tcp_if_include 172.20.0.0/16 --merge-stderr-to-stdout --prefix
  $MPI_ROOT -H 172.20.0.205:8,172.20.0.206:8 \
-x PYTHONPATH -x MASTER_ADDR -x RDMAV_FORK_SAFE=1 -x FI_EFA_USE_DEVICE_RDMA=1
  -x MASTER_PORT \
$PYTHON -u /home/habana/Model-References/PyTorch/computer_vision/
  classification/torchvision/train.py --batch-size=256 \
--model=resnet50 --device=hpu --workers=8 --print-freq=100 --epochs=40 -ebe 4
  --data-path=/data/pytorch/imagenet/ILSVRC2012 \
--dl-time-exclude=False --dl-worker-type=HABANA --autocast --optimizer=lars -
  -label-smoothing=0.1 --lars-weight-decay=0.0001 \
--lars_base_learning_rate=13 --lars_warmup_epochs=7 --lars_decay_epochs=41
```

Start TensorFlow Training Example

In this section, we train resnet50 with the ImageNet dataset on a TensorFlow framework.

This model is built by Habana Model Reference.

Prerequisites

Some prerequisites are required to set up your system to run this model on Intel® Gaudi® 2.

Install the required packages using pip:

```
python3 -m pip install -r requirements.txt
```

Set up the environment variables:

```
export PYTHONPATH=/home/habana/Model-References:$PYTHONPATH
export PYTHON=/usr/bin/python3.10
```

Training in 1 Accelerator

One Intel® Gaudi® 2 AI Accelerator, batch 256, 90 epochs, BF16 precision, SGD, Gaudi® 2 with media acceleration:

```
$PYTHON resnet_ctl_imagenet_main.py -dt bf16 -dlit bf16 -te 90 -ebe 90 -bs 256
  --jpeg_data_dir /data/tensorflow/imagenet --enable_tensorboard
```

Scale-Up Training in 8 Accelerators

Eight Intel® Gaudi® 2 AI Accelerators on 1 server, batch 256, 40 epochs, BF16 precision, LARS, Gaudi® 2 with media acceleration:

```
mpirun --allow-run-as-root --bind-to core -np 8 --map-by socket:PE=6 --merge-
  stderr-to-stdout \
  $PYTHON resnet_ctl_imagenet_main.py \
    --dtype bf16 \
    --data_loader_image_type bf16 \
    --use_horovod \
    -te 40 \
    -ebe 40 \
    -bs 256 \
    --optimizer LARS \
    --base_learning_rate 9.5 \
    --warmup_epochs 3 \
    --lr_schedule polynomial \
    --label_smoothing 0.1 \
    --weight_decay 0.0001 \
    --single_l2_loss_op \
```

```
    --jpeg_data_dir /data/tensorflow/imagenet \
    --enable_tensorboard
```

Scale-Out Training in 2 Devices with 16 Accelerators

Run training on 16 Intel® Gaudi® 2 AI Accelerators - Horovod:

```
mpirun \
 --allow-run-as-root --mca plm_rsh_args -p3022 \
 --bind-to core \
 --map-by socket:PE=6 -np 16 \
 --mca btl_tcp_if_include <interface_name> \
 --tag-output --merge-stderr-to-stdout --prefix $MPI_ROOT \
 -H 172.20.0.205:8,172.20.0.206:8 \
 -x HABANA_LOGS \
 -x PYTHONPATH -x HCCL_SOCKET_IFNAME=<interface_name> \
    $PYTHON resnet_ctl_imagenet_main.py \
      -dt bf16 \
      -dlit bf16 \
      -bs 256 \
      -te 40 \
      -ebe 40 \
      --use_horovod \
      --data_dir /data/tensorflow/imagenet/tf_records \
      --optimizer LARS \
      --base_learning_rate 13 \
      --warmup_epochs 7 \
      --momentum 0.9 \
      --lars_decay_epochs 41 \
      --lr_schedule polynomial \
      --label_smoothing 0.1 \
      --weight_decay 0.0001 \
      --single_l2_loss_op \
      --enable_tensorboard
```

6 Scale-Out Network Reference Design

This chapter includes the following sections:

Configuration

Link up all the scale-out ports of the AGS8200 servers.

Use **manage_network_ifs.sh** script to link up all the ports and check their status. All the ports must be linked up before tests or training.

```
root@habanas01:/# ./manage_network_ifs.sh --up
up : [########################################] 100%
24 habanalabs network interfaces were toggled up
root@habanas01:/# ./manage_network_ifs.sh --status
accel0
3 ports up (1, 8, 9)
accel1
3 ports up (1, 8, 9)
accel2
3 ports up (1, 8, 9)
accel3
3 ports up (1, 8, 9)
accel4
3 ports up (1, 8, 9)
accel5
3 ports up (1, 8, 9)
accel6
3 ports up (1, 8, 9)
accel7
```

Link up all the ports on the leaf switch.

1. Set the breakout mode.

```
admin@sonic:~$ sudo config interface breakout Ethernet0 4x100G -y -f
admin@sonic:~$ sudo config interface breakout Ethernet4 4x100G -y -f
admin@sonic:~$ sudo config interface breakout Ethernet8 4x100G -y -f
...
...
```

2. Add the ports to a VLAN.

```
admin@sonic:~$ sudo config vlan add 1
admin@sonic:~$ sudo config vlan member add 1 Ethernet0
admin@sonic:~$ sudo config vlan member add 1 Ethernet1
admin@sonic:~$ sudo config vlan member add 1 Ethernet2
...
...
```

3. Enable the interfaces.

```
admin@sonic:~$ config interface startup Ethernet0
admin@sonic:~$ config interface startup Ethernet1
admin@sonic:~$ config interface startup Ethernet2
...
...
```

7 References

This chapter includes the following sections:

Intel® Habana References

In this section, a quick navigation for Intel® Habana Official Information is provided as reference.

The information includes the base knowledge, training and porting the model, platform update, monitoring and orchestration solutions.

With these references, this section aims to make it more convenient to search related information for architecture developers, IT administrators and Cloud maintainers.

Table 5: Intel® Habana References

Habana Official Tutorials	Description	Titles
Habana Base Knowledge	Base Knowledge for Gaudi	■ Document Server ■ Tutorial Videos ■ Software Environment Information
Model Training Guide	Gaudi compatible framework	■ Pytorch Training ■ Tensorflow Training
Habana Model Reference	Habana verified model	■ Model Reference
SynapseAI Software Support	SynapseAI Install and Update	■ Stack and Driver Installation ■ Platform Upgrade and Full System Installation
Management and Monitoring	Detailed management tool	■ Management and Monitoring
Orchestration	Orchestration Solution	■ Orchestration ■ Kubernetes