NVIDIA Grace CPU SuperChip

MGX System Quick Start Guide

DU-11600-001_03 | January 2024

Document History

DU-11600-001_03

Version	Date	Description of Change
01	November 9, 2023	Initial release
02	December 6, 2023	• Updated the Ubuntu installation steps. • Added the SLES installation steps
03	January 22, 2024	Updated the post-install sequence for Ubuntu.

Table of Contents

Chapter 1. Introduction

This guide provides information about how to unbox and set up a Supermicro® 2U MGX NVIDA® Grace CPU SuperChip system.

Figure 1. Front and Back View of the MGX System

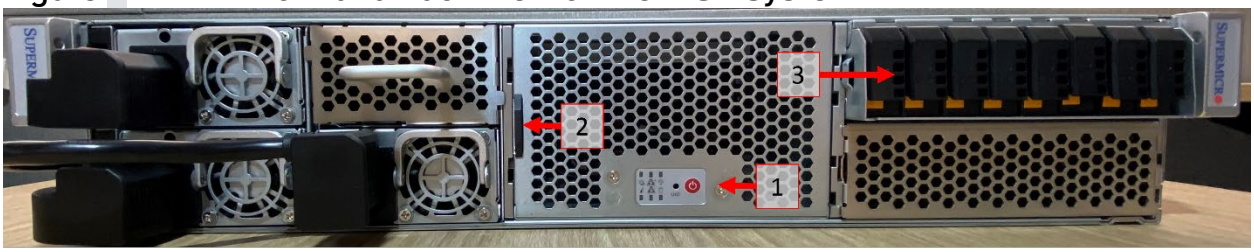

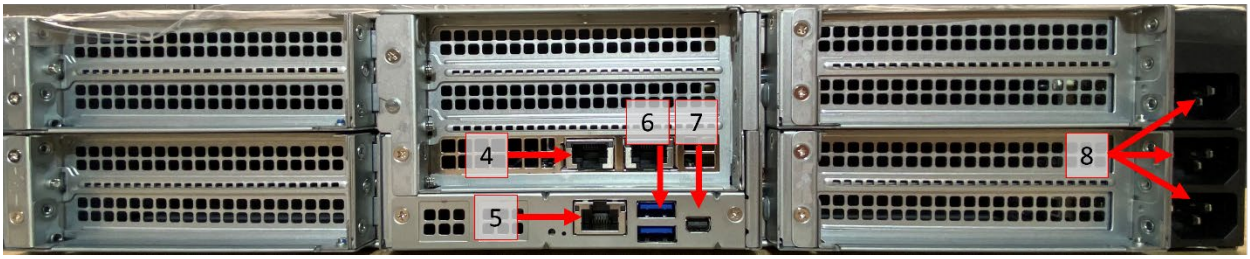

Table 1. Details

Item Number	Part Name
1	Front Panel Power
2	BMC Information Tag
3	NVMe Drive Bay
4	Host System RJ45 Ethernet Connection
5	BMC RJ45 Ethernet Connection
6	High Speed USB Ports
7	Mini Display Port
8	C13 Power Supply Cable Inserts

Chapter 2. Physical Install

Depending on location of installation, you need to install the server rail kit by following the manufacturer's directions, which can be found on the manufacturer's website. Install the rail kit into server rack using the rack manufacturer's recommended instructions.

2.1 Installing the Cable

Note: The numbers in parentheses correspond to the numbers in Figure 1.

1. The server comes with dual-port 10GBase-T Ethernet card.
2. Connect the (4) Host RJ45 Ethernet ports to the network that will be used to communicate with the server.
3. Connect the (5) BMC RJ45 Ethernet to the network that will be used to manage the Baseboard Management Console (BMC).
4. (Optional) Connect the Keyboard and Video Monitor to ports (6) and (7).

   This step is optional because the BMC interface provides virtual KVM access, but it can be useful to configure or view the BMC's IP address for the first time.
5. Connect three C13 power cables into the Power Supply Cable Inserts located on the rear of the server.

Chapter 3. First Boot

Note: The numbers in parentheses correspond to the numbers in Figure 1.

Now, you can boot the server by using the power button on the Front Panel (1). The server will initialize over the course of several minutes. After approximately five minutes, you will be able to log into the Baseboard Management Console (BMC) through its https interface.

- There are several methods to determine and set the IP address information you will use to access to the BMC.
- For your convenience, the manufacturer includes a tag (2) that includes the BMC's MAC address, which can be used to create a static DHCP assignment to assign an IP of your choice, or to correlate network traffic with the BMC's address.
- The BMC will default to requesting an IP address using DHCP.

Note: The boot process can take up to 20minute.

3.1 Method 1: Locate the BMC on Automatic DHCP Network

On a network that is already configured to use DHCP, you may use one of the following methods:

- **Option A**

  a. Open the web interface of your router, or the log for your DHCP server.

  b. Locate the list of issued IP addresses (**DHCP Leases**) and look for the BMC's MAC address on the list.

- **Option B**

  a. Request a static mapping from your router or DHCP server based on the server's BMC MAC address.

  b. Disconnect the server from power.

  c. Re-connect it and the BMC will automatically obtain this new address.

- **Option C**

  a. Connect the server to a network on the same subnet as your administration machine.

  b. Locate the server using the ARP table:

  - **Linux**: Run the `arp-scan –localnet` command.
  - **Windows PowerShell**: Run the `arp -a` command.

3.2 Method 2: View or Set the BMC IP Information Through a Local Display

With a keyboard and monitor attached to the server, you can determine and set the BMC IP settings and information through the System BIOS interface.

1. During initial boot, monitor for the BIOS options screen and follow instructions to enter the setup.
2. In the System BIOS menu, press **Tab** and select the **BMC** menu.
3. Select the **BMC network configuration** option and press enter.

Figure 1. The BMC network configuration Option

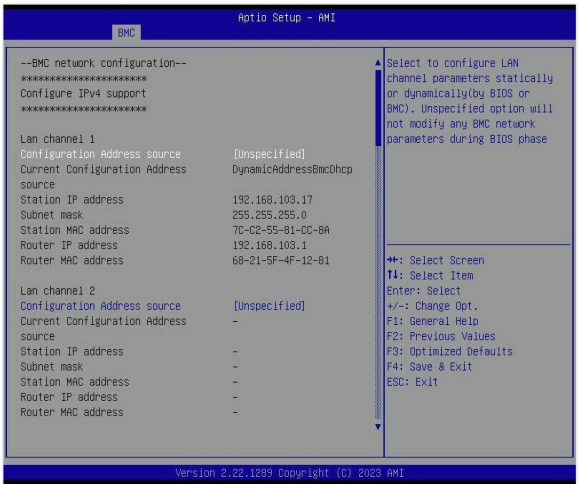

3.2.1 Setting a Static IP for the BMC through the System BIOS

For network configurations that require a static IP, this address can also be configured through the System BIOS settings.

1. Complete the instructions in “Method 2: View or Set the BMC IP Information “ on page 7.
2. In **BMC Network Configuration**, select **Configuration Address source** and then select **Static**”.

   This step allows you to alter the network and router information to correctly configure your BMC into your infrastructure.

Figure 2. Setting a Static IP Address

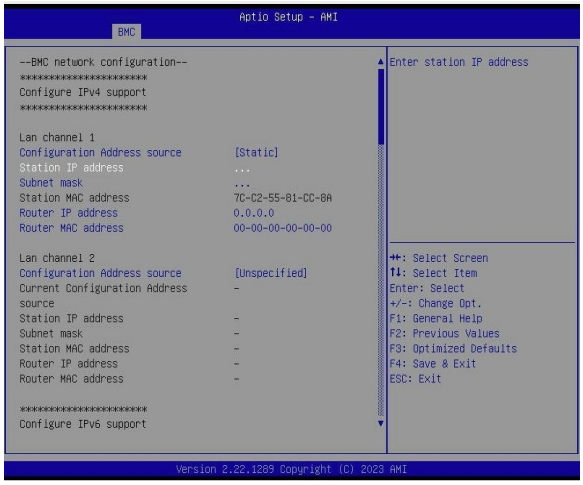

3.3 Accessing the BMC Network Interface

After you locate the IP address that is assigned to the server BMC, open a web browser on the same network and navigate to the IP address. For instance, if your network assigned 192.168.1.100 to your BMC’s MAC address, open a browser for a computer on the same network, and enter https://192.168.1.100 as the web address.

Note: To correctly load the page, https:// is required, and a certificate warning may be displayed. This is normal.

On the landing page, you will be prompted for a username and password:

- The default username from SuperMicro is **ADMIN** (case sensitive).
- The initial password is located on the BMC information tag (2) located on the front of the machine.

3.4 Installing the Operating System

Through the BMC's Remote Control menu, use **both** of the following interfaces to complete the operating system installation:

- The HTML5 Viewer, which is a virtual monitor that is like a physical monitor and that is connected to the box.
- The Serial over Lan (SOL) interface.

1. Write the ISO image in a bootable way onto the USB drive and place it into one of the USB ports on the server.
2. Launch the HTML5 Viewer window.
3. Use the on-screen interface to power cycle the server and begin a cold boot sequence.
   - If the system drive **is** completely blank, it will eventually boot to the mounted media,
   - If the system drive **is not** blank, watch the BIOS output and select the option to choose boot media.
4. Select the available USB boot item.

NVIDIA provides guides for each of the major OS distributions that are available in the following links:

- NVIDIA Grace Software with Ubuntu 22.04 Installation Guide.
- NVIDIA Grace Software with Red Hat Enterprise Linux 9 Installation Guide.
- NVIDIA Grace Software with SUSE Linux Enterprise Server 15 Installation Guide.

3.4.1 Installing Ubuntu on the Grace Server

1. Download Ubuntu 22.04.03 LTS from Downloading Ubuntu Server for ARM and write out bootable media.
2. After starting the server and selecting the boot device, the GRUB bootloader's Ubuntu installation menu will appear.
3. Select **Ubuntu Server with the HWE kernel**.

Figure 3. Ubuntu Server with the HWE Kernel Option

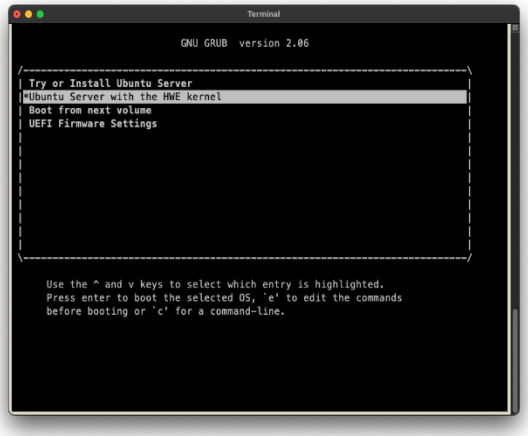

4. The HTML5 Viewer window will disappear and remain black for the rest of this process.

   This is a known issue that will be resolved in a future installer version that will incorporate newer versions of the kernel.

   To complete the installation, open the Serial over LAN Console (SoL Console)**.**
5. Complete the installation via the SoL viewer by selecting the options that are appropriate for you.
6. After the installation has completed, select **Reboot Now**.
7. After the system reboots, log in using the credentials that were specified during the installation.
8. Run the following commands to update the system and install the NVIDIA optimized Ubuntu kernel variant and reboot:

```
sudo DEBIAN_FRONTEND=noninteractive apt purge linux-image-$(uname -r) linux-headers-$(uname -r) linux-modules-$(uname -r) -y
sudo apt update
sudo apt install linux-nvidia-64k-hwe-22.04 -y
sudo reboot now
```

> **Note:** If linux-nvidia-64k-hwe-22.04 is not available, install linux-nvidia-64k-6.2 instead.

9. The system will reboot and the installation of Ubuntu on the Grace server is complete.

3.4.2 Installing Red Hat Enterprise Linux 9.3 on a Grace Server

1. After booting the ISO image, through the BMC's virtual media mount or through a physical media drive, the GRUB boot menu is displayed.
2. Select **Test this media & install Red Hat Enterprise Linux 9.3**.

Figure 4. Testing the Media

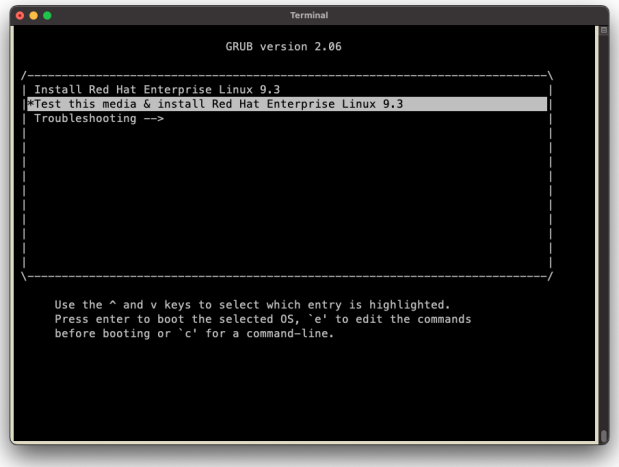

3. The HTML5 Viewer window will disappear and remain black for the rest of this process.
4. To complete the installation, use the SOL Console.
5. Follow the installer prompts to configure the manual installation.

> **Note:** There is a known GUI bug in the installer for RHEL version 9.2 that requires selecting "Use Text Mode" to successfully get through installation. This is not required for RHEL v 9.3 or above.

   a. From the Software Selection menu, select the 4k or 64k kernel from the Kernel Options submenu. See Section 5.4.3.1 Page Size for system configuration recommendations.

```
Begin the installation and allow the system to reboot.
```

6. After the system reboots, the installation of Red Hat Enterprise Linux 9.3 on the Grace server is complete.

3.4.3 Installing SUSE Linux Enterprise Server 15 SP5 on a Grace Server

1. After booting the ISO image, through the BMC's virtual media mount, or through a physical media drive, the GRUB splash menu is displayed.
2. To display the GRUB boot menu, press **t**.

Figure 5. SLES 15 SP5 GRUB Boot Menu

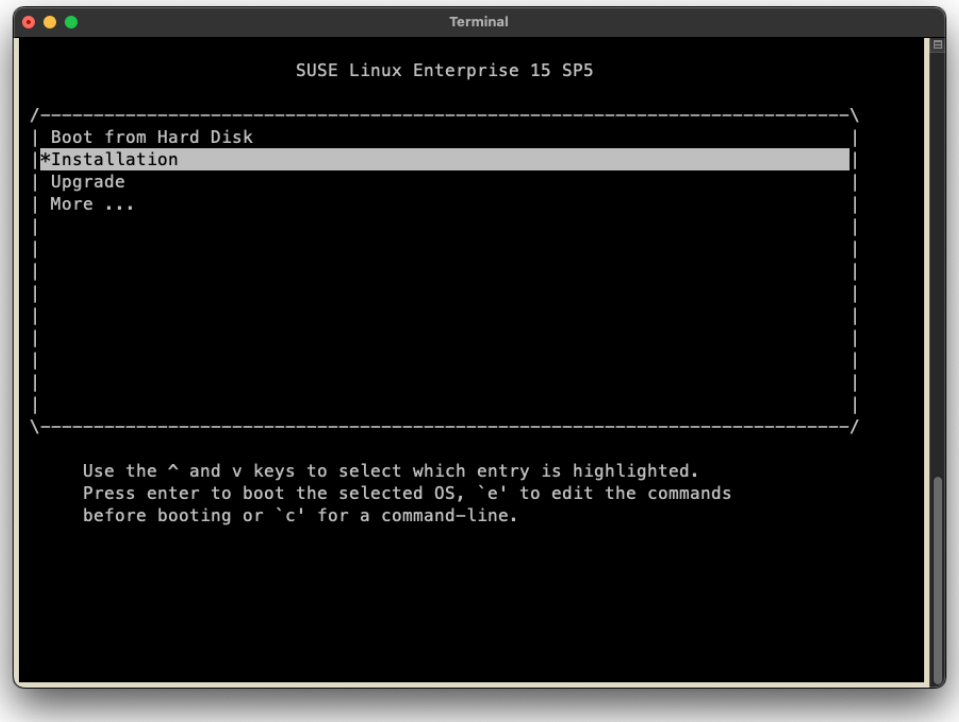

3. Highlight **Installation** and press **e** to edit the boot entry.
4. Append `modprobe.blacklist=ast` to the end of the list of kernel boot parameters.
5. Boot the entry by clicking **Ctrl-X** or pressing **F10**.
6. The HTML5 Viewer window will disappear and remain black for the rest of this process.
7. To complete the installation, use the SOL Console.
   a. Complete the installations by selecting the options that are appropriate for you.
   b. From the Installation Settings summary, select **Change menu.**
   c. Select the Software submenu to alter the installation and install the 64k kernel as the sole, default kernel for the system.
   d. From the software change pane, navigate to the **Search Phrase** box, and search for **kernel-64kb**.
   e. Select the **kernel-64kb** package.
   f. Return to the **Search Phrase** box and search for **kernel-default**.
   g. Deselect the kernel-default package.

h. Navigate to and select **Accept** to update the planned installation.
i. Begin the installation and allow the system to reboot.

8. After the system reboots, the installation of SUSE Linux Enterprise Server 15 SP5 on the Grace server is complete.

3.4.4 Multiple Linux Distributions: Compiling the Linux Kernel from Source

You can install almost any variety of Linux on the Grace server as long as the distribution is derived from the mainline Linux kernel.

3.4.4.1 Page Size

Grace supports 64K and 4K Linux kernel page sizes. To configure your Linux kernel with the page size that suits your business needs, change the following `kconfig` settings during the kernel compilation:

```
4K page size: CONFIG_ARM64_4K_PAGES=y
64K page size: CONFIG_ARM64_64K_PAGES=y
```

The 64K page size can benefit the applications that allocate a large amount of memory because there will be fewer page faults, better TLB hits, and efficiency.

Note: The recommended default value for the page size is 64K.

To build the kernel from the source:

1. Install the required minimum set of packages to build the kernel.

   ```
   sudo apt-get install git build-essential ncurses-dev xz-utils libssl-dev bc flex
   libelf-dev bison gcc-12-base
   ```

2. Download the desired version of the Linux Kernel Source from www.kernel.org.

   ```
   wget https://cdn.kernel.org/pub/linux/kernel/v6.x/linux-6.6.0.tar.gz
   ```

3. Extract the Source Code.

   ```
   tar -xvf linux-6.6.0.tar.gz
   ```

4. Copy an existing version of the NVIDIA kernel config file into the Linux Source directory as .config.

   ```
   cp -v /boot/config-*nvidia* .config
   ```

   You also can copy the distribution default config into the Source directory and **make menuconfig**. Ensure that you make the required changes in Table 1**Error! Reference source not found.**.

   ```
   cp -v /boot/config-$(uname -r) .config
   ```

5. Open the .config file in a text editor, search for the CONFIG_ARM64_4K_PAGES variable.
6. Make one of the following changes based on the page size configuration that you want.

Figure 6. Possible Changes Based on the Page Size

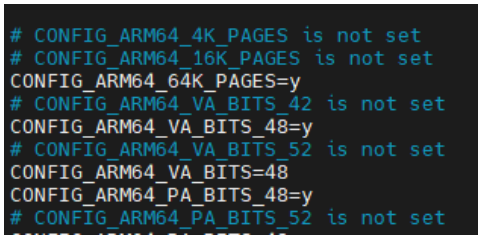

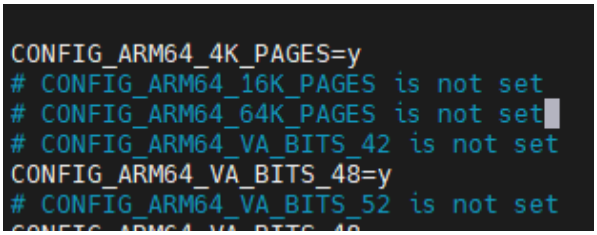

7. Save the changes to the .config file.
8. Run the following command from the main Source Code directory.

```
make
```

   a. If you are compiling the kernel on Ubuntu, you might see the following error that interrupts the build process:

```
No rule to make target 'debian/canonical-certs.pem
```

   b. To disable the conflicting security certificates, run the following commands:

```
./scripts/config --disable SYSTEM_TRUSTED_KEYS
./scripts/config --disable SYSTEM_REVOCATION_KEYS
```

9. To finish the building process, run the `make` command again.
10. After the building process is complete, to install the required modules and then the kernel, run the following commands.

```
sudo make modules_install
sudo make install
```

Chapter 4. Using the System

4.1 Grace Customized Build Stack

As Arm-Neoverse V2 and Arm v9 systems become more widespread, commonly distributed software will begin to incorporate architectural optimizations. Until that occurs, to take advantage of the accelerations available in Arm v9 architectures, you need to compile the build chain from the source.

Tip : Environment management tools, such as Modules or Spack, can control build environments and allow for multiple versions of build stacks.

4.1.1 Building GCC 12.3 From the Source

1. Download GCC using one of the following methods:
   - Download the source using a source mirror.
   - Run the following command to download directly to a host server that is connected to the internet.

     ```
     wget https://ftp.gnu.org/gnu/gcc/gcc-12.3.0/gcc-12.3.0.tar.gz
     ```
2. Untar the downloaded source into a directory on the target system where the user has proper file permissions.

   ```
   tar -xzvf gcc-12.3.0.tar.gz
   ```
3. Enter the source directory.

   ```
   cd gcc-12.3.0/
   ```
4. To download the packages and build GCC, run the following commands.

   ```
   ./contrib/download_prerequisites
   ```
5. To correctly configure the installation, run the following command.

   ```
   ./configure --disable-multilib --enable-shared --enable-languages=c,c++
   (--prefix can be used to change the default install location away from /usr/local/bin)
   ```
6. To build GCC from source, run the following command.

   ```
   make -j144
   ```

   You can exercise the cores in the Grace machine.

7. Install GCC.

```
sudo make install
```

8. Verify the GCC installation.

```
which gcc # reports /usr/local/bin/gcc or the location specified with --prefix
gcc --version # reports gcc (GCC) 12.3.0
```

4.2 Performance Benchmarking Guide

A selection of tests have been developed to stress the Grace Superchip architecture. Refer to the Grace CPU Benchmarking Guide.

Appendix A. Required and Recommended Kernel Configurations for Grace

Table 2. Required Bare Metal Configs

Kernel Config	Description
CONFIG_NR_CPUS=512	Supports the maximum Grace configuration.
CONFIG_NODES_SHIFT=6	Supports the maximum Grace configuration.
CONFIG_ARM_SMMU_V3_SVA=y	Support shared virtual addressing.
CONFIG_ARM64_PMEM=y	Support persistent memory.
CONFIG_ARM_SDE_INTERFACE=y	Support RAS notifications.
CONFIG_BLK_DEV_PMEM=m	Enable persistent memory block device.
CONFIG_DEVICE_MIGRATION=y	Enable device physical page migration.
CONFIG_DEVICE_PRIVATE=y	Supports unaddressable device memory; only required when using NVIDIA TRD.
CONFIG_HOTPLUG_PCI_PCIE=y	Supports the PCIe native hotplug.
CONFIG_PCIE_DPC=y	Supports downstream port containment.
CONFIG_PCIE_EDR=y	Enables the error disconnect recover support.
CONFIG_SPI_TEGRA210_QUAD=m	Support the QSPI controller.
CONFIG_TCG_TIS_SPI=m	Supports the TPM SPI interface.
CONFIG_MTD_SPI_NOR=y	Support the SPI NOR flash device.
CONFIG_IPMI_SSIF=m	Supports the SMBus interface to BMC.
arch/arm64/include/asm/irq.h: #if	Supports the maximum Grace configuration. Not required when kernel carries 721255b9826bd11c7a38b585905fc2dd0fb94e52

Kernel Config	Description

`defined(CONFIG_ARM_GIC_V3_ITS)`
`#define  NR_IRQS  (1 << 19)`
`#endif`	

Table 3 contains the recommended config settings that provide performance improvements for certain workloads.

Table 3. Recommended Bare Metal Configs

Kernel Config	Description
`CONFIG_ARM64_64K_PAGES=y`	Use 64K page size; required when using NVIDIA TRD.
`CONFIG_CPU_FREQ_DEFAULT_GOV_PERFORMANCE=y`	Set default CPU frequency governor to “performance”.
`CONFIG_CPU_FREQ_GOV_SCHEDUTIL=y`	Supports the schedutil CPU frequency governor.
`CONFIG_PREEMPT_DYNAMIC=y`	Allows dynamic preemption tuning using preempt.
`CONFIG_PREEMPT_NONE=y`	Default dynamic preemption tuning to `preempt=none` for throughput.
`CONFIG_DMABUF_HEAPS=y`	Enables DMA-BUF memory heaps.
`CONFIG_DMABUF_HEAPS_SYSTEM=y`	Enables the system dmabuf heap.
`CONFIG_DMI_SYSFS=y`	Enables the export of raw DMI table data.
`CONFIG_INIT_ON_ALLOC_DEFAULT_ON=n`	Disables heap memory zeroing on allocation by default.

Table 4 contains optional config settings that enable the performance tooling functions on Grace platforms.

Table 4. Bare Metal Configs: Performance Tools

Kernel Config	Description
`CONFIG_ARM_CORESIGHT_PMU_ARCH_SYSTEM_PMU=m`	Enables the ARM CoreSight PMU driver.
`CONFIG_SENS/ORS_ACPI_POWER=m`	Enables power telemetry through hwmon. Enable when sysfs endpoints for hardware power monitoring are not present.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

Arm

Arm, AMBA, and ARM Powered are registered trademarks of Arm Limited. Cortex, MPCore, and Mali are trademarks of Arm Limited. All other brands or product names are the property of their respective holders. ”Arm” is used to represent ARM Holdings plc; its operating company Arm Limited; and the regional subsidiaries Arm Inc.; Arm KK; Arm Korea Limited.; Arm Taiwan Limited; Arm France SAS; Arm Consulting (Shanghai) Co. Ltd.; Arm Germany GmbH; Arm Embedded Technologies Pvt. Ltd.; Arm Norway, AS, and Arm Sweden AB.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Copyright

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com