



NVIDIA Grace Platform Support Software Patches and Configurations

Document History

TB-11481-001_06

Version	Date	Authors	Description of Change
01	August 10, 2023		Initial release
02	September 7, 2023		<ul style="list-style-type: none">• Indicated TPM patch is now upstream.• Added additional arm_cspmu and mm patches.• Added CONFIG_DEVICE_MIGRATION and CONFIG_DEVICE_PRIVATE.• Clarified that CONFIG_ARM64_64K_PAGES is conditionally required for some NVIDIA GPU Driver versions.
03	September 8, 2023		<ul style="list-style-type: none">• Revised conditional requirements for CONFIG_ARM64_64K_PAGES and CONFIG_DEVICE_PRIVATE.• Added CONFIG_INIT_ON_ALLOC_DEFAULT_ON.
04	November 2, 2023		<ul style="list-style-type: none">• Reformatted document to categorize content.• Removed duplicate entries, added an additional iommu and dma-mapping patch.• Added CONFIG_SENSORS_ACPI_POWER .
05	December 18, 2023		<ul style="list-style-type: none">• Added contact information for NVIDIA Enterprise Experience.
06	April 18, 2024		<ul style="list-style-type: none">• Sorted tables by kernel version.• Added a performance patch table.• Added additional mm, pci, and thermal fix patches.• Added additional performance tools patches.• Updated CONFIG_TCG_TIS_SPI and CONFIG_SPI_TEGRA210_QUAD from modules to builtins .

Table of Contents

Introduction	4
Platform Software Patches and Configurations	4
Linux Kernel Patches	4
Upstream Bare Metal Linux Kernel	4
Linux Kernel Configs	11

Introduction

NVIDIA® Grace™ systems are composed of multiple hardware components that require support across different software subsystems. While several of these components have had adequate support in their respective upstream domains for some time, there is a subset that requires specific updates to remain reliably functional on Grace platforms. This guide provides information about the known patches and configuration settings that are required (or recommended) for interoperability with the Grace platform and might be helpful for customers who want to create their own custom software, for example a Linux kernel, for these systems. This guide is applicable to all platforms that use the NVIDIA Grace system.

For questions about the content in this document, contact [NVIDIA Enterprise Experience](#).

Platform Software Patches and Configurations

This section provides information about software patches and configuration settings that are required or recommended for the Grace platform.

Linux Kernel Patches

This section provides information about the Linux Kernel patches that support the Grace platform.

Upstream Bare Metal Linux Kernel

The tables in this section list Linux kernel patches that are *upstream*, which means that the patches are accepted into the [main Linux kernel branch](#).



Note: There might be circumstances where additional, dependent patches are required to support the patches listed in these tables, for example, when the patch listed is part of a larger series.

The Git description is pulled directly from the main Linux kernel Git log and is intended to help with searches and comparisons. The description might contain spelling and grammatical errors.

[Table 1](#) contains patches that enable functions and are required for bare metal support on the Grace platform.

Table 1. Bare Metal Support: Enablement

Git Commit	Git Description	Minimum Linux Kernel Release
4c207e7121fa92b66bf1896bf8ccb9edfb0f9731	PCI: Mark some NVIDIA GPUs to avoid bus reset	v5.13
6937b7dd434962377e00efc04adac0390c287199	PCI: Add support for ACPI _RST reset method	v5.15
a9c38c5d267cb94871dfa2de5539c92025c855d7	dma-mapping: remove bogus test for pfn_valid from dma_map_resource	v5.16
bd2fdedbf2bac27f4a2ac16b84ab9b9e5f67006c	i2c: tegra: Add the ACPI support	v5.17
f89d2cc3967af9948ffc58e4cc9a1331f1c4971a	spi: tegra210-quad: use devm call for cdata memory	
50665d58db052b04f640fd54de1632aeecd4fc77	i2c: tegra: use i2c_timings for bus clock freq	
5f84e73f9a8f14b95115b0eb2080da6d9fa7a82e	gpio: tegra186: Add IRQ per bank for Tegra241	
d1056b771ddbe119f733fd980c854d1bcf33d13e	gpio: tegra186: Add support for Tegra241	
d7583c8a57485da19feb6dd85573763a8c5ec1d1	i2c: tegra: Add SMBus block read function	v5.18
ca0acb511c21738b32386ce0f85c284b351d919e	device property: Add fwnode_irq_get_byname	
a263a84088f689bf0c1552a510b25d0bcc45fcae	i2c: smbus: Use device_*(*) functions instead of of_*(*)	
20e8ef5c7ffaa4c1d5f65727c536ffaf37078d5f	docs: firmware-guide: ACPI: Add named interrupt doc	
ac982578e7d340dc4f4fd243f4a4b24787d28c3f	spi: tegra210-quad: use device_reset method	
ea23f0e148b82e5bc6c6c814926f53133552f0f3	spi: tegra210-quad: add new chips to compatible	
75a1b44a54bd97500e524cf42e8c81cc632672b3	spi: tegra210-quad: add acpi support	
1b8342cc4a387933780c50f0cf51c94455be7d11	spi: tegra210-quad: combined sequence mode	
b76134178168b5104851b3c72d9b1092b7414ff9	spi: tegra210-quad: Multi-cs support	v6.0
0ddf10a319f14ed626dab3a2d390e5059e20cea5	arm64: tegra: Enable Tegra SPI & QSPI in deconfig	v6.1
721255b9826bd11c7a38b585905fc2dd0fb94e52	genirq: Use a maple tree for interrupt descriptor management	v6.5

[Table 2](#) contains patches that resolve critical issues and hardware errata.

Table 2. Bare Metal Support: Critical Fixes

Git Commit	Git Description	Minimum Linux Kernel Release
5bf71889ad9a4d39b7665c105a005c5a33d730ba	i2c: tegra: Set ACPI node as primary fwnode	v6.2
9f855779a3874eee70e9f6be57b5f7774f14e510	i2c: tegra: Fix PEC support for SMBUS block read	v6.4
35727af2b15d98a2dd2811d631d3a3886111312e	irqchip/gicv3: Workaround for NVIDIA erratum T241-FABRIC-4	
4327a6137ed43a091d900b1ac833345d60f32228	drm/ast: Fix ARM compatibility	

[Table 3](#) contains patches that resolve faults in enablement patches and other issues that have been discovered while testing the Grace platform with various workloads.

Table 3. Bare Metal Support: Fixes

Git Commit	Git Description	Minimum Linux Kernel Release
8777dd9dff4020bba66654ec92e4b0ab6367ad30	spi: tegra210-quad: Fix combined sequence	v6.1
ae4b3c1252f0fd0951d2f072a02ba46cac8d6c92	spi: tegra210-quad: Don't initialise DMA if not supported	
2197aa6b0aa236b9896a09b9d08d6924d18b84f6	spi: tegra210-quad: Fix duplicate resource error	
2e3468778dbe3ec389a10c21a703bb8e5be5cfbc	mm: remember young/dirty bit for page migrations	
101ca8d05913b7d1e6e8b9dd792193d4082fff86	rtc: efi: Enable SET/GET WAKEUP services as optional	v6.2
ef3691683d7bfd0a2acf48812e4ffe894f10bfa8	KVM: arm64: GICv4.1: Fix race with doorbell on VPE activation/deactivation	
05da178ce0aa152f0592e10fa874054187f3621b	ACPI/IORT: Update SMMUv3 DeviceID support	
047ee71ae4f412d8819e39e4b08c588fa299cfc2	spi: tegra210-quad: Fix validate combined sequence	v6.3
2449d436681d40bc63ec2c766fd51b632270d8a7	spi: tegra210-quad: Fix iterator outside loop	
f7482d8285b638be87a594a30edaaaf1341135c1a	spi: tegra210-quad: set half duplex flag	

Git Commit	Git Description	Minimum Linux Kernel Release
010338d729c1090036eb40d2a60b7b7bce2445b8	arm64: kaslr: don't pretend KASLR is enabled if offset < MIN_KIMG_ALIGN	
7c7b962938dda6a9cd095de557ee5250706ea88	mm: take a page reference when removing device exclusive entries	
c0e0421a60bf468e88cf569fbd727346b138ed04	ACPI: processor: Reorder acpi_processor_driver_init()	
c43198af05cffa5de8d4f356c40ce4bdca066272	thermal: core: Introduce thermal_cooling_device_present()	
790930f44289c8209c57461b2db499fcc702e0b3	thermal: core: Introduce thermal_cooling_device_update()	
22c52fa5155a2f48aedb0f675903b20457285a27	ACPI: processor: thermal: Update CPU cooling devices on cpufreq policy changes	
b57841fb0b564c61508222e885ac8f30a2811089	thermal: core: Drop excessive lockdep_assert_held() calls	
bba5065963f8ade14f3caa5b0f5b4a53d3054dfd	PCI/AER: Configure ECRC only if AER is native	
67a142dc9eb96a5cc018e5db62390665eb5f038c	spi: Add TPM HW flow flag	v6.4
967ca91a996f82219f2883e9e53d8e20df49025a	spi: tegra210-quad: Enable TPM wait polling	
3e35d303ab7d22c4b6597e56ba46ee7cc61f3a5a	arm64: module: rework module VA range selection	v6.5
a86a42ac2bd652fdc7836a9d880c306a2485c142	tpm_tis-spi: Add hardware wait polling	v6.6
d5afb4b47e13161b3f33904d45110f9e6463bad6	iommu/arm-smmu-v3: Fix soft lockup triggered by arm_smmu_mm_invalidate_range	
416ef04fe00c5f2f6fb8e13d8dbe1b5a0a274f83	mm/page_alloc: fix min_free_kbytes calculation regarding ZONE_MOVABLE	
081488051d28d32569ebb7c7a23572778b2e7d57	mm/mglru: fix underprotected page cache	v6.7
5095a2b23987d3c3c47dd16b3d4080e2733b8bb9	mm/mglru: try to stop at high watermarks	
8aa420617918d12d1f5d55030a503c9418e73c2c	mm/mglru: respect min_ttl_ms with memcgs	
4376807bf2d5371c3e00080c972be568c3f8a7d1	mm/mglru: reclaim offlined memcgs harder	
db744ddd59be798c2627efbfc71f707f5a935a40	PCI/MSI: Prevent MSI hardware interrupt number truncation	v6.8
6d4266675279b38c301243f3a4fac4a511b03246	PCI/ASPM: Update save_state when configuration changes	v6.9

[Table 4](#) contains optional patches that are known to offer improved performance on Grace platforms.

Table 4. Bare Metal Support: Performance

Git Commit	Git Description	Minimum Linux Kernel Release
a2ee7581afd59015b8f9ae01fad131aed9f26f01	ACPI: thermal: Add Thermal fast Sampling Period (_TFP) support	v6.8
310293a2b94197f3d75e65ab22672287a7938a00	ACPI: processor: reduce CPUFREQ thermal reduction pctg for Tegra241	
ccb45b34d44016b91fa75646741d317d6d6fdeea	ACPI: arm64: export acpi_arch_thermal_cpufreq_pctg()	
7dc7c5ef6463111991002f24c0aea08afe86f2cc	mm: allow deferred splitting of arbitrary anon large folios	
372cbd4d5a0665bf7e181c72f5e40e1bf59b0b08	mm: non-pmd-mappable, large folios for folio_add_new_anon_rmap()	
3485b88390b0af9e05dc2c3f57e9936f41e159a0	mm: thp: introduce multi-size THP sysfs interface	
19eaf44954df64f9bc8dec398219e15ad0811497	mm: thp: support allocation of anonymous multi-size THP	
b6aab3384cafbfa151c53d3b5f7e1f8d073aadf03	selftests/mm/kugepaged: restore thp settings at exit	
00679a183ac6d2584723cfc2a2c07c8285f802dc	selftests/mm: factor out thp settings management	
4f5070a5e40db2e9dbf5ffff4ec678d6fbb338d5c	selftests/mm: support multi-size THP interface in thp_settings	
9f0704eae8a4edc8dca9c8a297f798d505a4103a	selftests/mm/khugepaged: enlighten for multi-size THP	
12dc16b38463a671bc91dc2df10f3a014a27ff3b	selftests/mm/cow: generalize do_run_with_thp() helper	
c0f79103322c322ea9342d52c2d81528b7b56232	selftests/mm/cow: add tests for anonymous multi-size THP	
6280d7317ccae19c776a3b6cf9848c964f958091	mm: clarify the spec for set_ptes()	v6.9
2bdba9868a4ffcb1492db7272f34b54387910177	mm: thp: batch-collapse PMD with set_ptes()	
583ceaaa339960e673ac0029f323bb1c6ffc96d7	mm: introduce pte_advance_pfn() and use for pte_next_pfn()	
c1bd2b4028ae5b4d2ada64b31c40cc44cdf00972	arm64/mm: convert pte_next_pfn() to pte_advance_pfn()	

Git Commit	Git Description	Minimum Linux Kernel Release
506b586769ecef8c83fff64de227e7fa84b7be42	x86/mm: convert pte_next_pfn() to pte_advance_pfn()	
fb23bf6bd288db3187c27b971e558a3e9f70ae96	mm: tidy up pte_next_pfn() definition	
532736558e8ef2865eae1d84b52dda4422cac810	arm64/mm: convert READ_ONCE(*ptep) to ptep_get(ptep)	
659e193027910a5d3083e34b488ab459d2ef5082	arm64/mm: convert set_pte_at() to set_ptes(..., 1)	
cbb0294fdd72a5f63ec59fad5c0a98d63bd572fc	arm64/mm: convert ptep_clear() to ptep_get_and_clear()	
5a00bfd6a52cf31e93d5f1b734087deb32a3cffa	arm64/mm: new ptep layer to manage contig bit	
d9d8dc2bd3fb2689309f704fe85e6dde2b1bd73a	arm64/mm: dplit __flush_tlb_range() to elide trailing DSB	
4602e5757bcceb231c3a13c36c373ad4a750eddb	arm64/mm: wire up PTE_CONT for user mappings	
311a6cf29690bb8295327bad0e76e0ad48cadcc4	arm64/mm: implement new wrprotect_ptes() batch API	
6b1e4efb6f5499ae8f9f5cdda7502285a0edbf51	arm64/mm: implement new [get_and_]clear_full_ptes() batch APIs	
c6ec76a2ebc5829e5826b218d2e1475ec11b333e	mm: add pte_batch_hint() to reduce scanning in folio_pte_batch()	
fb5451e5f72b31002760083a99fbb41771c4f1ad	arm64/mm: implement pte_batch_hint()	
b972fc6afba002319fe23bc698ce6431ee43868c	arm64/mm: __always_inline to improve fork() perf	
f0c2264958e18bc7bc35b567d51b99461e4de34f	arm64/mm: automatically fold contpte mappings	

[Table 5](#) contains optional patches that enable functions, or resolve faults, with performance tooling on Grace platforms.

Table 5. Bare Metal Support: Performance Tools

Git Commit	Git Description	Minimum Linux Kernel Release
1c5d62f5dd979dbe7f4b16a5bd3c64fbecb7d92a	ACPICA: Add support for ARM Performance Monitoring Unit Table.	v5.19
6251d38059ae22304ede4f3748af9f795bdbf4fd	ACPI: ARM Performance Monitoring Unit Table (APMT) initial support	v6.2
e37dfd65731dc4f001fa7dfa7f705e6840017d5a	perf: arm_cspmu: Add support for ARM CoreSight PMU driver	v6.2
84481be7167eba1957c8718c8b044a47965fdecb	perf: arm_cspmu: Add support for NVIDIA SCF and MCF attribute	
39522031798dbb53b169b95d4ab25b53301d4eaf	ACPI: APMT: Fix kerneldoc and indentation	
1830902eb896824ca313a50f3486645c2df21327	perf: arm_cspmu: Fix modular builds due to missing MODULE_LICENSE(s)	
e72dbf9085b56bbc19ff332f82adec1891077637	perf: arm_cspmu: Fix build failure on x86_64	
a91bbd5c9984a2b15e68aad7a79c2809fbd10fbe	perf: arm_cspmu: Fix module cyclic dependency	
16e15834659e9f5c05b9f12da6e86d76165c60a3	perf: arm_cspmu: Fix variable dereference warning	v6.4
f462eb1d0df39499eff2407b2581ddf2fe89e744	arm64: defconfig: Enable ARM CoreSight PMU driver	
225d757012e0afa673d8c862e6fb39ed2f429b4d	perf: arm_cspmu: Set irq affinity only if overflow interrupt is used	v6.5
71e0cb32d5fc61468e83ed962379af71bba8237e	perf/arm_cspmu: Fix event attribute type	
7e51d05e43f19f394b3b0df01a22225143f9c5b5	perf: arm_cspmu: Add missing MODULE_DEVICE_TABLE	
87b3b6d53efccbb64396675b35839180ef14ef40	ACPI/APMT: Don't register invalid resource	
d2e3bb51281875be23cb4726a59b03d0a53eb0d3	perf/arm_cspmu: Decouple APMT dependency	
f9bd34e3753ea8f1433a3ba70f03a165a1416f98	perf/arm_cspmu: Clean up ACPI dependency	
3095e90eee5ea2d5658cab90b6da9c6d5d0a3bdf	coresight: etm4x: Allocate and device assign 'struct etmv4_drvdata' earlier	v6.6
4e3b9a6eae987c80330e5253754dab35acc2a63b	coresight: etm4x: Drop iomem 'base' argument from etm4_probe()	
5a1c7097472fcde5745654e3a59f55140903d9cc	coresight: etm4x: Drop pid argument from etm4_probe()	

Git Commit	Git Description	Minimum Linux Kernel Release
73d779a03a76ac3fe26832cba3c9ad04194af595	coresight: etm4x: Change etm4_platform_driver driver for MMIO devices	
3a2888aa1f962c55ca36119aeb67355c7bf54e4	coresight: platform: acpi: Ignore the absence of graph	
134124acb57f8ad59634d0e4530812205bf55250	coresight: etm4x: Add ACPI support in platform driver	
1aa3d0274a4aac338ee45a3dfc3b17c944bcc2bc	arm_pmu: acpi: Add a representative platform device for TRBE	
81e5ee471609848ee1ebf3beb2a46788113fe0eb	arm_pmu: acpi: Refactor arm_spe_acpi_register_device()	
15c7ef7341a2e54cfa12ac502c65d6fd2cce2b62	perf: arm_cspmu: Reject events meant for other PMUs	v6.7
a16afcc58a8c5ebc65c852faf001f8f61f05e4ef	perf cs-etm: Fix incorrect or missing decoder for raw trace	
17f8b216e02654a0b37127736ce78b32ccaa867b	coresight: trbe: Enable ACPI based TRBE devices	v6.8
4277f035d227e829133df284be7e35b7236a5b0f	coresight: trbe: Add a representative coresight_platform_data for TRBE	

Linux Kernel Configs

This section provides information about the Linux Kernel config settings for the Grace platform.

[Table 6](#) contains config settings that enable function and are required for bare metal support on the Grace platform.

Table 6. Bare Metal Configs: Required

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.

Kernel Config	Description
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 GPU Driver.
CONFIG_HOTPLUG_PCI_PCIE=y	Supports the PCIe native hotplug.
CONFIG_IOMMU_DEFAULT_PASSTHROUGH=n	Disable IOMMU translation bypass for DMA. Refer to <i>Input-Output Memory Management Unit Passthrough</i> in the NVIDIA Grace Performance Tuning Guide for more information.
CONFIG_PCIE_DPC=y	Supports downstream port containment.
CONFIG_PCIE_EDR=y	Enables the error disconnect recover support.
CONFIG_SPI_TEGRA210_QUAD=y	Support the QSPI controller.
CONFIG_TCG_TIS_SPI=y	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 defined(CONFIG_ARM_GIC_V3_ITS) #define NR_IRQS (1 << 19) #endif	Supports the maximum Grace configuration. Not required when the kernel carries 721255b9826bd11c7a38b585905fc2dd0fb94e52 .

[Table 7](#) contains the recommended config settings that provide performance improvements for certain workloads.

Table 7. Bare Metal Configs: Recommended

Kernel Config	Description
CONFIG_ARM64_64K_PAGES=y	Use 64K page size; required when using NVIDIA GPU Driver.
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.

Kernel Config	Description
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 8](#) contains optional config settings that enable the performance tooling functions on Grace platforms.

Table 8. Bare Metal Configs: Performance Tools

Kernel Config	Description
CONFIG_ARM_CORESIGHT_PMU_ARCH_SYSTEM_PMU=m	Enables the ARM CoreSight PMU driver.
CONFIG_SENSORS_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 NON-INFRINGEMENT, 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

© 2024 NVIDIA Corporation & Affiliates. All rights reserved.