NVIDIA DriveOS 7.2 Features List

DA-13034-001
NVIDIA CONFIDENTIAL
| October 2026
| Prepared and Provided Under NDA

Table of Contents

Core Software Stack 3
GPU, Graphics, and Acceleration 4
PVA and VPU 5
Multimedia 6
Streaming and Inter-SoC Communication 7
Display 8
Virtualization and Guest OS 9
Networking and Communication 10
Boot, Storage, and Flashing 11
Developer Productivity and Tooling 12
Hardware Compatibility 13

Core Software Stack

- Upgrades Linux kernel to version 6.8, delivering the latest upstream security hardening, driver support, and CPU speculation mitigations.
- Updates CUDA to version 13.3, unlocking the latest GPU compute capabilities on Thor.
- Upgrades TensorRT™ (TRT) to version 11.0, improving inference performance and DNN model compatibility.
- Introduces TRT Edge LLM, enabling on-device large language model inference for speech, text, and AI assistant workloads.
- Introduces TRT Model Connect, a new open-source platform on GitHub providing broad AI model support on DRIVE AGX Thor hardware, enabling developers to bring pre-trained models onto Thor with reduced integration effort.
- Adds Nsight Deep Learning Designer 2025.5, enabling visual DNN optimization workflows directly in the developer environment.
- Updates cuDNN to version 9.23, improving deep learning kernel performance and framework compatibility.
- Upgrades the Nsight toolchain — Systems 2026.3, Graphics 2026.2, PerfSDK 2026.2 — delivering the latest profiling and debugging capabilities.
- Upgrades base OS to Ubuntu 24.04 (Noble) LTS, refreshing system libraries and userspace package baseline with long-term security support.
- Updates GCC to version 13.2, delivering the latest C/C++ compiler improvements, diagnostics, and security features.
- Upgrades Yocto to version 5.0 (Scarthgap), refreshing the embedded build system with current recipes and toolchain support.
- Adds Branch History Buffer and Straight-Line Speculation mitigations, hardening the kernel against speculative-execution side-channel vulnerabilities on Thor.

GPU, Graphics, and Acceleration

- Introduces GPU Hardware Scheduler with configurable per-context priorities (HIGH/MEDIUM/LOW), enabling concurrent AV, LLM, Cluster, and IVI workloads within a single Linux VM.
- Adds GPU Quality of Service for Guest OS, letting developers allocate guaranteed GPU timeslices per application domain — ADAS, AI, Cluster, IVI.
- Enables GPU availability and recovery mechanisms, allowing applications to recover from GPU hangs without a full system reboot.
- Introduces deterministic stride control for NvSciSyncObjects across DLA, PVA, and GPU, enabling pre-submission of workloads and long signal chains with minimal CPU involvement.
- Adds CUDA NvSciStreams Producer with RGBA fp16 and YUV420 support, unicast and multicast streaming across intra- and inter-process boundaries.
- Adds YUV422 DisplayPort output format support, broadening display pipeline compatibility for cockpit and visualization use cases.

PVA and VPU

- Updated PVA Algorithms Library with a new Brute Force Matcher for feature descriptor matching, enabling visual odometry, SLAM, and object tracking workloads on Thor's Programmable Vision Accelerator.
- Adds PVA signaler and waiter support in NvSciSync, enabling tight coordination between PVA and other compute engines.
- Enables persistent PVA engine power management, reducing latency for repeated PVA workloads by keeping the engine in a ready state.
- Adds real-time PVA utilization reporting in tegrastats, giving developers direct visibility into PVA load during runtime.

Multimedia

- Adds JPEG decode support via NVJPG engine and JPEG encoding MVP, enabling hardware-accelerated image codec operations for camera and recording pipelines.
- Introduces HEVC QP map encoding with macroblock-level quality control, giving developers fine-grained bitrate management for video recording workloads.
- Enables multichannel audio playback and recording via APE I2S, A2B microphone array integration with daisy-chain support for up to four arrays, and USB audio I/O.

Streaming and Inter-SoC Communication

- Enables PCIe C2C reconnection for both SOC suspend (SC7) and SOC shutdown, restoring chip-to-chip communication without a full reboot.
- Introduces consumer process re-attach and late-attach support in NvSciStream, allowing new consumers to join active multicast pipelines without disrupting existing producers.

Display

- Adds display mixed mode support for simultaneous MST and SST DisplayPort configurations, enabling multi-monitor cockpit setups from a single DP port.
- Enables display vertical sync polarity configurability, supporting adaptation to varied display hardware requirements.
- Integrates OpenWFD for standards-compliant display management.

Virtualization and Guest OS

- Adds Direct Drive Update (DDU) support, allowing developers to update a DRIVE AGX Thor target directly from their development host — see Developer Guide for workflow details.
- Adds NvSciIPC resource limits for Guest VM (up to 8192 channels), windowed watchdog monitoring, and CPU usage measurement, improving VM observability and reliability.
- Introduces platform configuration for PCIe server core elimination.
- Enables DPAUX interface configuration as I2C controller in Guest OS, expanding peripheral connectivity for LiDAR and other I2C devices.
- Adds hypervisor-based SoC reset interface for software-controlled resets without external intervention.

Networking and Communication

- Adds USXGMII Ethernet mode support and HW virtualized DMA channel networking, broadening automotive Ethernet coverage and providing isolated network channels per Guest VM.
- Introduces PTP boundary clock implementation and Ethernet concurrent timestamp API across 10 Gbps interfaces, enabling precise multi-sensor time synchronization.
- Adds MACsec VLAN tag placement support and MGBE MAC link status reporting for production Ethernet network configurations.

Boot, Storage, and Flashing

- Adds FFU via Bootburn for eMMC/UFS firmware updates with cryptographic integrity checks and anti-rollback protection.
- Introduces authenticated secondary IFS and DataFS (ISO 26262 / ISO 21434 certified), enabling class-key-based filesystem authentication separate from the secure boot key.
- Adds flexible partition and file system management with multi-level repartitioning and dual bootchain binary generation.
- Introduces storage configuration visualization tool and PCD storage configuration tooling, giving developers clear insight into partition layout and storage allocation.
- Adds board-specific parameterization via external JSON for Bootburn, preserving safety qualification while enabling platform customization without tool change.

Developer Productivity and Tooling

- Adds a curated set of DriveOS Skills via Additional Packages on NGC, equipping AI coding agents with step-by-step workflows to set up DRIVE AGX Thor DevKits and to build, validate, and deploy applications on DriveOS.
- Adds IDE integration with Docker support, enabling build, deploy, and debug workflows from VS Code without leaving the container environment.
- Introduces customizable device-tree property separation and board-specific JSON parameterization for Bootburn, reducing platform bring-up friction for custom hardware.
- Adds target file system customization with modular, feature-module-driven RFS assembly, making it easier to build minimal or purpose-specific root filesystems.
- Delivers pre-installed Ubuntu 24.04 Desktop RFS variant, enabling containerized development and Docker-based validation workflows on target.
- Adds storage lifetime estimation tool, giving developers actionable insight into storage health and projected wear across UFS and eMMC devices.

Hardware Compatibility

Adds software support for DRIVE AGX Thor SKU 10 and SKU 12 Developer Kit hardware revisions, including updated Infineon 10 Gb Ethernet PHY and Realtek Ethernet Switch devices — see Product Compatibility Notification for revision details.

Notice

The information provided in this specification is believed to be accurate and reliable as of the date provided. However, NVIDIA Corporation (“NVIDIA”) does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information. 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 publication supersedes and replaces all other specifications for the product that may have been previously supplied.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and other changes to this specification, at any time and/or to discontinue any product or service without notice. Customer should obtain the latest relevant specification 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. NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this specification.

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 these specifications will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do 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 specification. NVIDIA does not accept any 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 specification, 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 specification. 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 specification is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

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. 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 NVIDIA terms and conditions of sale for the product.

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.

OpenCL

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

Blackberry

BLACKBERRY, EMBLEM Design, QNX, AVIAGE, MOMENTICS, NEUTRINO and QNX CAR are the trademarks or registered trademarks of BlackBerry Limited, used under license, and the exclusive rights to such trademarks are expressly reserved.

Trademarks

NVIDIA and 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.

Copyright

© 2026 NVIDIA Corporation and affiliates. All rights reserved.