



NVIDIA DOCA

Glossary

Table of Contents

Chapter 1. Terms and Acronyms..... 1

Chapter 1. Terms and Acronyms

| Term | Description |
|-------------|--|
| ACE | AXI coherency extensions |
| ACPI | Advanced configuration and power interface |
| AMBA® | Advanced microcontroller bus architecture |
| ARB | Arbitrate |
| ATF | Arm-trusted firmware |
| AXI4 | Advanced eXtensible Interface 4 |
| BDF address | Bus, device, function address. This is the device's PCIe bus address to uniquely identify the specific device. |
| BERT | Boot error record table |
| BF_INST_DIR | The directory where the BlueField software is installed |
| BFB | BlueField bootstream |
| BMC | Board management controller |
| BSD | BlueField software distribution |
| BSP | BlueField support package |
| BUF | Buffer |
| CBS | Committed burst size |
| CHI | Coherent hub interface; Arm® protocol used over the BlueField Skymesh specification |
| CIR | Committed information rate |
| CL | Cache line |
| CMDQ | Command queue |
| CMO | Cache maintenance operation |
| COB | Collision buffer |
| DAT | Data |
| DEK | Data encryption key |
| DMA | Direct memory access |
| DOCA | DPU SDK |
| DOT | Device ownership transfer |

| Term | Description |
|------------|--|
| DPA | Data path accelerator; an auxiliary processor designed to accelerate data-path operations |
| DPCP | Direct packet control plane |
| DPDK | Data plane development kit |
| DPI | Deep packet inspection |
| DPU | Data processing unit, the third pillar of the data center with CPU and GPU |
| DVM | Distributed virtual memory |
| DW | Dword |
| EBS | Excess burst size |
| ECE | Enhanced connection establishment |
| ECPF | Embedded CPU physical function |
| EIR | Excess information rate |
| EMEM/EMI | External memory interface; block in the MSS which performs the actual read/write from the DDR device |
| eMMC | Embedded multi-media card |
| ESP | EFI system partition |
| ESP header | Encapsulating security payload |
| FIPS | Federal Information Processing Standards |
| FPGA | Field-programmable gate arrays |
| FS | File system |
| FW | Firmware |
| GDB | GNU debugger |
| GPT | GUID partition table |
| HART | HW thread; a logical DPA processing unit |
| HCA | Host-channel adapter |
| HNF | Home node interface |
| Host | <p>When referring to "the host" this documentation is referring to the server host. When referring to the Arm based host, the documentation will specifically call out "Arm host".</p> <ul style="list-style-type: none"> ▶ Server host OS refers to the Host Server OS (Linux or Windows) ▶ Arm host refers to the AARCH64 Linux OS which is running on the BlueField Arm Cores |
| HW | Hardware |
| hwmon | Hardware monitoring |
| IB | InfiniBand |

| Term | Description |
|-----------|--|
| ICM | Interface configuration memory |
| IKE | Internet key exchange |
| IPMB | Intelligent platform management bus |
| IPMI | Intelligent platform management interface |
| IR | Intermediate representation |
| KGDB | Kernel debugger |
| KGDBOC | Kernel debugger over console |
| LAT | Latency |
| LCRD | Link credit |
| LSO | Large send offload |
| LTO | Link-time optimization |
| MFT | Mellanox firmware tools |
| MIO | Memory-mapped I/O |
| MLNX_OFED | Mellanox OpenFabrics Enterprise Distribution |
| MSB | Most significant bit |
| MSS | Memory subsystem |
| MST | Mellanox software tools |
| NAT | Network address translation |
| NIC | Network interface card |
| NIST | National Institute of Standards and Technology |
| NS | Namespace |
| OCD | On-chip debugger |
| OOB | Out-of-band |
| OS | Operating system |
| OVS | Open vSwitch |
| PBS | Peak burst size |
| PCIe | PCI Express; Peripheral Component Interconnect Express |
| PF | Physical function |
| PIR | Peak information rate |
| PK | Public key |
| PKA | Public key accelerator |
| POC | Point of coherence |
| RD | Read |
| RDMA | Remote direct memory access |
| RegEx | Regular expression |
| REQ | Request |
| RES | Response |
| RN | Request node |

| Term | Description |
|------------|--|
| | RN-F – Fully coherent request node |
| | RN-D – IO coherent request node with DVM support |
| | RN-I – IO coherent request node |
| RNG | Random number generator/generation |
| RoCE | Ethernet and RDMA over converged Ethernet |
| RQ | Receive queue |
| RShim | Random shim |
| RTT | Round-trip time |
| RX | Receive |
| RXP | Regular expression processor |
| SA | Security association |
| SBSA | Server base system architecture |
| SDK | Software development kit |
| SF | Sub-function or scalable function |
| SG | Scatter-gather |
| SHA | Secure hash algorithm |
| SMMU | System memory management unit |
| SNAP | Storage-defined network-accelerated processing |
| SNP | Snooping |
| SPDK | Storage performance development kit |
| SQ | Send queue |
| SR-IOV | Single-root IO virtualization |
| STL | Stall |
| Sync event | Synchronization event |
| TBU | Translation buffer unit |
| TIR | Transport interface receive |
| TIS | Transport interface send |
| TLS | Transport layer security |
| TRB | Trail buffer |
| TSO | TCP send offload |
| TSO | Total store order |
| TX | Transmit |
| UDS | Unix domain socket |
| UEFI | Unified extensible firmware interface |
| UPVS | UEFI persistent variable store |
| VF | Virtual function |
| VFE | Virtio full emulation |

| Term | Description |
|----------------|-------------------------------|
| VM | Virtual machine |
| VMA | NVIDIA® Messaging Accelerator |
| VPI | Virtual protocol interconnect |
| VST | Virtual switch tagging |
| WorkQ or workq | Work queue |
| WQE | Work queue elements |
| WR | Write |
| WRDB | Write data buffer |
| XLIO | NVIDIA® Accelerated IO |

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 nor any of its direct or indirect subsidiaries and affiliates (collectively: "NVIDIA") make no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assume 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, and Mellanox are trademarks and/or registered trademarks of Mellanox Technologies Ltd. and/or NVIDIA Corporation in the U.S. and in other countries. The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2023 NVIDIA Corporation & affiliates. All rights reserved.