



Overview

Table of contents

BlueField DPU

BlueField SuperNIC

The BMC node enables remote power cycling, board environment monitoring, NVIDIA® BlueField® SoC temperature monitoring, board power and consumption monitoring, and individual interface resets. The BMC also supports the ability to push a bootstream to the BlueField. It is recommended to manage NVIDIA® BlueField® using Redfish commands. However, IPMI commands and sysfs monitoring infrastructure are available as well .

Warning

Make sure to log into the BMC first and change the [global default password](#) to prevent malicious attackers from hacking your system.

The procedures described in this manual assume that you have already installed and powered on your device according to the instructions in BlueField's [specific hardware guide](#).

- Support for IPMI 2.0 (v1.1) Standards
 - Thermal control – access to all relevant temperature sensors, fan control
 - System management – power state control, power on/off, reboot/reset
 - Environmental monitoring – voltage/current/power
 - Serial over LAN (SOL)
 - RMCP/RMCP+
 - Event log management
 - Event alerting
 - VLAN support
- Support for DMTF Standards
 - Redfish specification (DSP0266)
 - Network Controller Sideband Interface (NC-SI) Specification (DMTF DSP0222)

- Support for BMC image update

BlueField DPU

The NVIDIA® BlueField® DPU is a data center infrastructure on a chip that combines a high-speed networking interface with powerful, software-programmable Arm cores, enabling breakthrough networking, storage, and security performance. The BlueField DPU offloads, accelerates, and isolates a broad range of software-defined infrastructure services which traditionally ran on the host's CPU, overcoming performance and scalability bottlenecks, and eliminating security threats in modern data centers.

BlueField DPUs transform traditional computing environments to secure and accelerated data centers, allowing organizations to efficiently run data-driven, cloud-native applications alongside legacy applications. By decoupling the data center infrastructure from business applications, BlueField DPUs enhance data center security, streamline operations, and reduce total cost of ownership.

The BlueField DPU contains a programmable CPU based on Arm cores, a state-of-the-art NVIDIA® ConnectX®, and an enhanced set of security, storage, and networking accelerators that can be configured to perform multiple software-defined, hardware-accelerated functions. With a BlueField DPU, a software-defined network, and/or software-defined storage solution can be deployed and offloaded from the main host CPU in the server. Similarly, other dedicated services (e.g., distributed firewall, deep packet inspection, malware detection) can run on the BlueField DPU and can be accelerated with zero CPU overheads.

The BlueField DPU resembles a server embedded within the server itself, creating a secure environment where an infrastructure stack can operate independently from the primary (i.e., host) CPU, effectively isolating it from the untrusted tenant applications.

This is the recommended mode for utilizing the DPU in which software running on the host CPU has no direct access to the DPU. For instance, in a scenario where a cloud service provider is responsible for managing both networking and storage in a cloud infrastructure stack, it can establish an isolated environment within the DPU.

BlueField SuperNIC

The NVIDIA® BlueField® SuperNIC is the world's most advanced network accelerator, designed for supercharging hyperscale generative AI workloads. It delivers deterministic, isolated performance, with secure cloud multi-tenancy. Featured on the Spectrum-X networking platform, NVIDIA integrates BlueField-3 SuperNICs across its accelerated systems to enable peak AI workload efficiency. Powered by the NVIDIA DOCA software, the SuperNIC offers up to 400Gb/s connectivity between GPU servers, with features like RoCE adaptive routing, direct data placement (DDP), and programmable congestion

control. With its unique HHHL form factor and low-power platform, the BlueField-3 SuperNIC fits most enterprise-class servers.

 **Info**

To read more about the BlueField-3 features and benefits, refer to [this page](#).

 **Info**

To read more about the BlueField-2 features and benefits, refer to [this page](#).

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 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 2025, NVIDIA. PDF Generated on 11/20/2025