

NVIDIA Introduces New Family of BlueField DPUs to Bring Breakthrough Networking, Storage and Security Performance to Every Data Center

New Data-Center-Infrastructure-on-a-Chip Software — DOCA SDK — Enables Developers to Create New DPU-Accelerated Infrastructure Applications; VMware Announces NVIDIA BlueField-2 Support as Part of Redefinition of Hybrid Cloud Architecture

GTC — NVIDIA today announced a new kind of processor — DPUs, or data processing units — supported by DOCA, a novel data-center-infrastructure-on-a-chip architecture that enables breakthrough networking, storage and security performance.

NVIDIA founder and CEO Jensen Huang revealed the company's three-year DPU roadmap in today's [GPU Technology Conference keynote](#). It features the new [NVIDIA® BlueField®-2 family](#) of DPUs and NVIDIA DOCA™ software development kit for building applications on DPU-accelerated data center infrastructure services.

“The data center has become the new unit of computing,” said Huang. “DPUs are an essential element of modern and secure accelerated data centers in which CPUs, GPUs and DPUs are able to combine into a single computing unit that's fully programmable, AI-enabled and can deliver levels of security and compute power not previously possible.”

Optimized to offload critical networking, storage and security tasks from CPUs, BlueField-2 DPUs enable organizations to transform their IT infrastructure into state-of-the-art data centers that are accelerated, fully programmable and armed with “zero-trust” security features to prevent data breaches and cyberattacks.

A single BlueField-2 DPU can deliver the same data center services that could consume up to 125 CPU cores. This frees up valuable CPU cores to run a wide range of other enterprise applications.

Widespread Adoption of NVIDIA DPUs

Leading server manufacturers worldwide — including [ASUS](#), [Atos](#), [Dell Technologies](#), Fujitsu, [GIGABYTE](#), H3C, [Inspur](#), [Lenovo](#), [Quanta/QCT](#) and [Supermicro](#) — have plans to integrate NVIDIA DPUs into their enterprise server offerings.

These commitments from system providers are complemented by extensive support from software infrastructure partners, including:

- [VMware](#) announced substantial work underway with NVIDIA as part of its recently announced Project Monterey initiative to support BlueField-2 DPUs with VMware Cloud Foundation.
- [Red Hat](#) plans to offer support for BlueField-2 DPUs with Red Hat Enterprise Linux and Red Hat OpenShift, components of Red Hat's open hybrid cloud portfolio, which is used by 95 percent of the Fortune 500.
- [Canonical](#) announced support of BlueField-2 DPUs and DOCA in its Ubuntu Linux platform, the most popular operating system among public clouds.
- [Check Point Software Technologies](#), a leading cybersecurity provider, is integrating BlueField-2 DPUs into its technologies, which more than 100,000 organizations worldwide use to protect themselves from cyberattacks.

NVIDIA DPU Portfolio:

NVIDIA's current [DPU lineup](#) includes two PCIe products:

- The **NVIDIA BlueField-2 DPU**, which features all of the capabilities of the NVIDIA Mellanox® ConnectX®-6 Dx SmartNIC combined with powerful Arm cores. Fully programmable, it delivers data transfer rates of 200 gigabits per second and accelerates key data center security, networking and storage tasks, including isolation, root trust, key management, RDMA/RoCE, GPUDirect, elastic block storage, data compression and more.
- The **NVIDIA BlueField-2X DPU**, which includes all the key features of a BlueField-2 DPU enhanced with an NVIDIA Ampere GPU's AI capabilities that can be applied to data center security, networking and storage tasks. Drawing from NVIDIA's third-generation Tensor Cores, it is able to use AI for real-time security analytics, including identifying abnormal traffic, which could indicate theft of confidential data, encrypted traffic analytics at line rate, host introspection to identify malicious activity, and dynamic security orchestration and automated response.

NVIDIA DOCA Software Development Kit

The new [NVIDIA DOCA SDK](#) enables developers to build applications on DPU-accelerated data center infrastructure services, much like the [NVIDIA CUDA®](#) programming model enables developers to build GPU-accelerated applications.

DOCA provides developers a comprehensive, open platform for building software-defined, hardware-accelerated networking,

storage, security and management applications running on the BlueField family of DPUs.

DOCA is fully integrated into [NVIDIA NGC™](#), a software catalog offering a convenient, containerized software environment for third-party application providers to leverage advanced DPU data-center-accelerated services and to develop, certify and distribute applications to customers.

Availability

BlueField-2 DPUs are sampling now and expected to be featured in new systems from leading server manufacturers in 2021. BlueField-2X DPUs are under development and are also expected to become available in 2021.

DOCA is [available](#) for early access partners now.

About NVIDIA

NVIDIA's (NASDAQ: NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined modern computer graphics and revolutionized parallel computing. More recently, GPU deep learning ignited modern AI — the next era of computing — with the GPU acting as the brain of computers, robots and self-driving cars that can perceive and understand the world. More information at <http://nvidianews.nvidia.com/>.

Certain statements in this press release including, but not limited to, statements as to: the benefits, performance, features and availability of DPUs and the NVIDIA DOCA SDK; DPUs as an essential element of modern and secure accelerated data centers; the ability of CPUs, GPUs and DPUs to combine into a single computing unit that is fully programmable, AI-enabled and can deliver levels of security and compute power not previously thought possible; and leading server manufacturers' plans to integrate DPUs into their enterprise server offerings are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

© 2020 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, BlueField, ConnectX, CUDA, Mellanox, NGC and NVIDIA NGC are trademarks and/or registered trademarks of NVIDIA Corporation and/or Mellanox Technologies in the U.S. and other countries. All other trademarks and copyrights are the property of their respective owners.

Kristin Bryson
Enterprise Data Center, AI/DL
+1-203-241-9190
kbryson@nvidia.com