Next-Generation Networking for the Next Wave of AI

White Paper

November 2023

Table of Contents

List of Figures

List of Tables

Abstract

In the era of generative AI, the pivotal role of GPU-accelerated computing cannot be overstated. Modern AI workloads, being computationally intensive with unique characteristics, demand specialized network infrastructure for optimal performance. NVIDIA, a leader in accelerated computing and AI, has developed Spectrum-X, an Ethernet networking platform that's designed to enhance the performance and efficiency of AI clouds. At the heart of the Spectrum-X platform is the BlueField-3 SuperNIC, a novel network accelerator that works seamlessly with the NVIDIA Spectrum™-4 Ethernet switch. This combination forms the foundation of an accelerated computing fabric for AI. NVIDIA has strategically integrated the BlueField-3 SuperNICs into its flagship AI systems, enhancing their performance while ensuring deterministic and isolated performance for tenant jobs. This white paper dives into how the BlueField-3 SuperNIC takes Ethernet networking for AI to the next level within NVIDIA's Spectrum-X platform.

November 2023

Peak AI Performance Demands Optimized Networking

Making its debut in November 2022, OpenAI's ChatGPT is regarded by many as the ”iPhone moment of AI”, attracting over 100 million users worldwide in just two months. As organizations rapidly deploy AI tools, research powerhouse McKinsey assesses that the impact of generative AI on productivity could add trillions of dollars of value to the global economy across dozens of use cases. With significant opportunities on the horizon, enterprises worldwide are racing to find the right infrastructure to build generative AI models and applications. This marks the genesis of a new computing era, characterized by accelerated computing and generative AI.

Figure 1. A New Computing Era

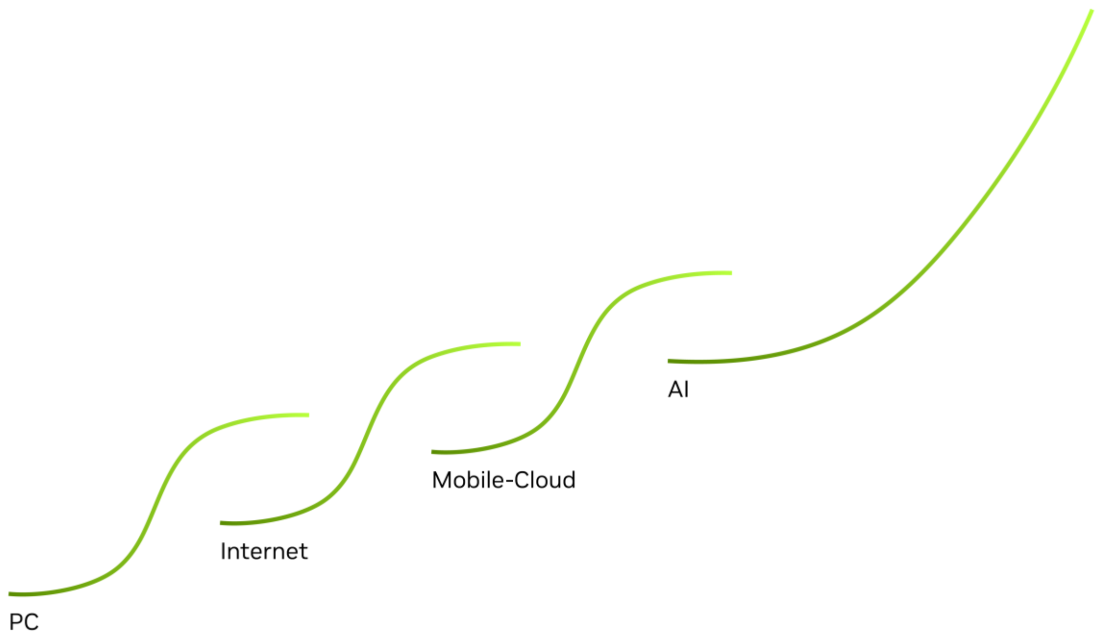

Accelerated computing plays a crucial role in AI, providing the specialized hardware and software needed for the underlying infrastructure. Modern AI relies heavily on the parallel processing strengths of GPU-accelerated computing. This enables faster model training, handles larger model sizes, and enables real-time inference. These accelerators are extensively used in AI applications, allowing researchers and developers to process large datasets and train complex neural networks more efficiently, thus propelling AI forward.

Generative AI workloads, which operate at data center scale, exhibit unique characteristics, requiring high-speed and low-latency network connectivity between GPU servers for optimal performance. The following table highlights the key characteristics of both traditional cloud networks and networks for AI.

Table 1. Key Characteristics of traditional Cloud Networks and Networks for AI

Traditional Ethernet-based Clouds	AI Computing Ethernet Networking
Loosely coupled applications	Distributed tightly coupled processing
Low bandwidth TCP flows and utilization	High bandwidth RoCE flows and utilization
High jitter tolerance	Low jitter tolerance
Heterogeneous traffic, statistical multi-pathing	Bursty network capacity, elephant flows

Examining GPU-accelerated systems running massively-distributed workloads reveals significant differences from CPU-based systems, with implications for the underlying network infrastructure.

Table 2. Key Features of General-Purpose CPU Systems and GPU-Accelerated Systems

Traditional Ethernet-based Clouds	AI Computing Ethernet Networking
General-purpose processor handles a wide range of tasks	Specialized processor designed for parallel computation
Usually ships with two CPUs with a few dozens of cores in total	Systems with four to eight GPUs each with tens of thousands of cores
Scale-out to a few dozen nodes per workload	Workloads operate at data center-scale, and up to tens of thousands of GPUs
CPU-centric network I/O	GPU-centric network I/O

Traditional Ethernet, designed to offer broad interoperability and serve loosely coupled applications, falls short in delivering the necessary network performance and efficiency required to meet the rigorous demands of computationally intensive distributed AI workloads. To achieve peak performance, organizations require a new class of network infrastructure, capable of handling the specific characteristics of generative AI workloads while also supporting the multi-tenant nature of modern cloud computing.

Optimized Ethernet Networking for AI

NVIDIA Spectrum-X stands as the world's first Ethernet networking platform optimized for AI. It achieves this through the tight coupling of NVIDIA Spectrum-4, a cutting-edge Ethernet switch, and NVIDIA BlueField-3 SuperNIC, a powerful network accelerator. This integration enhances networking performance for AI by up to 1.6X, while delivering consistent and predictable performance, alongside fine-grained security, in multi-tenant clouds. Spectrum-X is supercharged by NVIDIA's acceleration software and software development kits (SDKs), empowering developers to build software-defined, cloud-native AI applications.

Figure 2. NVIDIA Spectrum-X Networking Platform

The following diagrams illustrate the varying levels of network efficiency for a set of Large Language Model (LLM) training workloads. In traditional Ethernet setups, workload placement within the cluster greatly impacts network efficiency, which is calculated as a percentage of peak bandwidth. For example, when a job is run on several nodes that are all situated within the same rack, it performs better than if the same task was distributed across nodes in different racks. In contrast, Spectrum-X consistently delivers high-performance levels, regardless of workload distribution, outperforming standard Ethernet environments by as much as 60% in terms of efficiency.

Figure 3 **NVIDIA Spectrum-X Achieves Consistent and Predictable Performance for LLM Collectives**

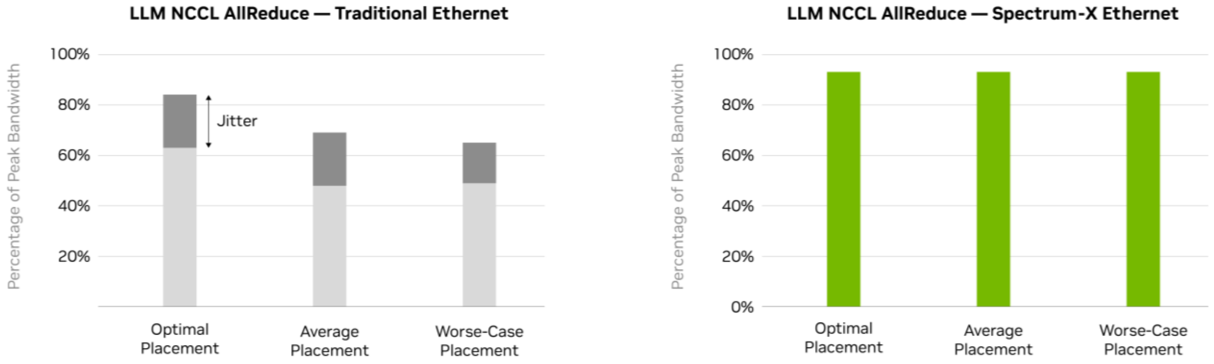

The BlueField-3 SuperNIC plays a pivotal role within the Spectrum-X platform. This advanced network accelerator is purpose-built to supercharge hyperscale AI workloads. Drawing on the robust capabilities of the BlueField-3 networking platform, the BlueField-3 SuperNIC is designed for network-intensive, massively parallel computing. It delivers extremely fast and efficient communication between GPU servers, providing the Spectrum-X platform with a host of benefits:

- Peak AI Workload Efficiency: BlueField-3 accelerates AI training and inference performance by delivering higher network bandwidth and lower latency.
- Secure Multi-Tenant AI Cloud: BlueField-3 provides secure, high-performance networking solutions for multi-tenant AI clouds.
- Consistent and Predictable Performance: BlueField-3 ensures that each job and tenant's performance is isolated, predictable, and unaffected by other network activities.
- Enhanced Power Efficiency: With a sub-75-watt power envelope, BlueField-3 optimizes power consumption in AI data centers.
- Extensible Network Infrastructure: BlueField-3 offers advanced programmability, making it flexible and adaptable to serve a myriad of network environments.
- Broad OEM Server Support: BlueField-3 fits seamlessly into most enterprise-class servers.

Subsequent sections discuss the key BlueField-3 SuperNIC innovations within the Spectrum-X platform, and their role in optimizing Ethernet networks for AI:

1. 400Gb/s RoCE (RDMA over Converged Ethernet) networking
2. RoCE adaptive routing and packet reordering
3. Advanced congestion control
4. Secure and accelerated VPC networking for AI computing
5. Power-efficient, low-profile design

Figure 4 NVIDIA BlueField-3 SuperNIC

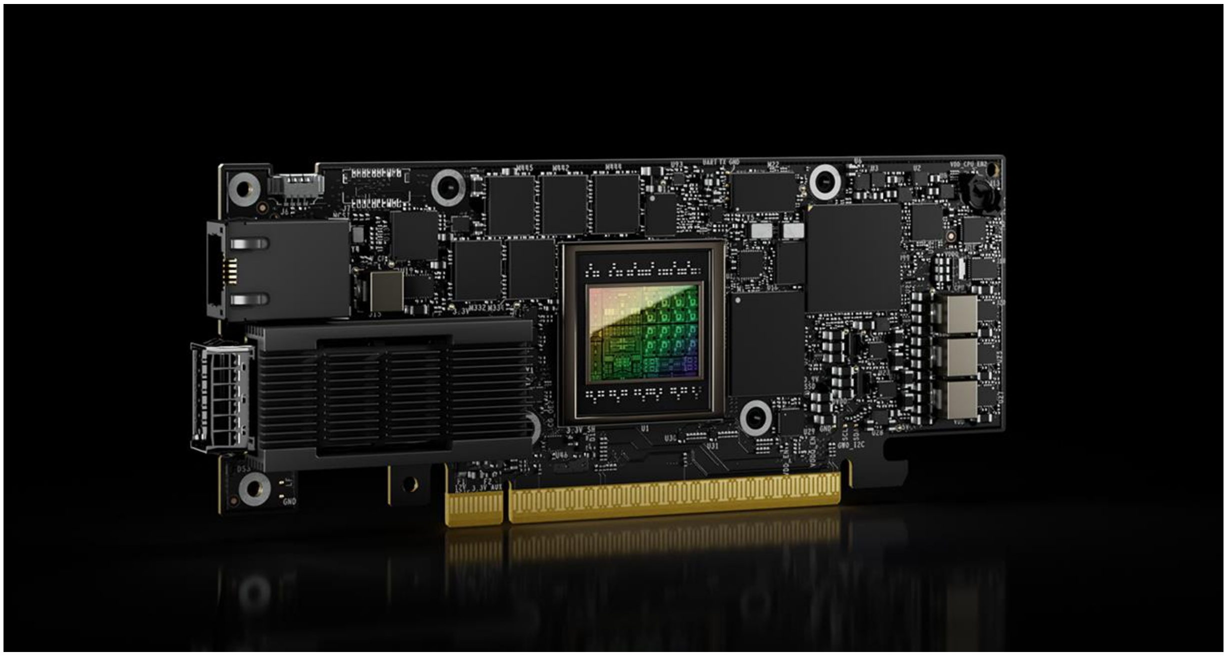

400Gb/s RoCE Networking

Large language AI models are trained on vast amounts of data, from which they learn intricate patterns, relationships, and nuances of language comprehension. The effective processing of these data sets relies on the tight coordination of NVIDIA GPUs, which are distributed across hundreds of physical nodes. Bridging the gap between computational power and data processing for Generative AI, the network bandwidth available for effective GPU communication often determines time-to-market for an AI-powered product offering or a business process.

The BlueField-3 SuperNIC addresses this connectivity challenge. It provides 400Gb/s line-speed, low-latency network connectivity between GPUs. Equipped with best-in-class, in-hardware RDMA over Converged Ethernet (RoCE) acceleration, and GPUDirect RDMA technology, the BlueField-3 SuperNIC enables the ideal Ethernet network fabric for GPU-to-GPU communication.

Further exploring the technical enhancements for AI acceleration, Remote Direct Memory Access (RDMA), and its Ethernet derivative, RoCE, play a crucial role in accelerating AI workloads by streamlining data transfers, reducing latency, and maximizing network throughput. GPUDirect RDMA allows data to be moved directly between GPU memory, without involving the CPU, resulting in lower latency and reduced CPU overhead. For AI training, GPUDirect RDMA ensures that GPUs can communicate efficiently and at high bandwidth, enabling parallel processing and scaling of AI training tasks across multiple GPUs and nodes.

Figure 5. BlueField-3 SuperNIC with GPUDirect RoCE Enables Direct GPU-to-GPU Communication

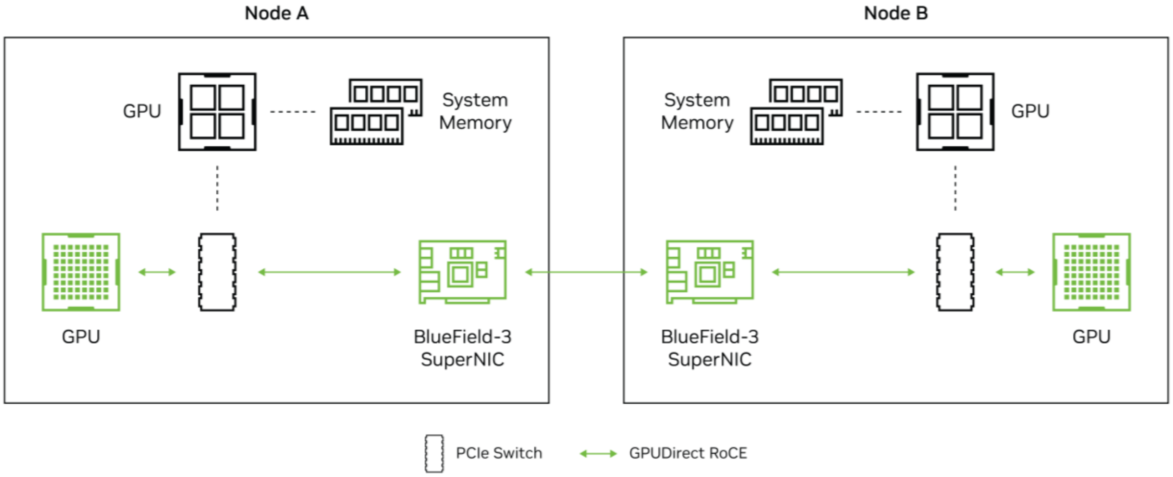

The preceding figure illustrates how GPUDirect RDMA on BlueField-3 SuperNICs enables direct communication across physical nodes. This approach eliminates the role of the CPUs and the need for buffer copies of data via the system memory, resulting in a 10X performance improvement. BlueField-3 SuperNICs deliver 400Gb/s RoCE network connectivity, a networking protocol synonymous with accelerated Ethernet networking for AI.

RoCE Adaptive Routing (AR) and Packet Reordering

With the proliferation of generative AI clouds, which can scale to thousands of nodes, traditional IP routing techniques, such as equal cost multipath (ECMP), have demonstrated inefficiencies. These are particularly noticeable with 'elephant flows' between the same pairs of GPU nodes, typical with AI training. Elephant flows have a voracious appetite for bandwidth, capable of saturating the entire network bandwidth and often persisting over extended durations. ECMP is designed to distribute traffic evenly across available paths using a hashing mechanism. However, using ECMP for routing AI traffic often leads to network congestion due to inadequate load balancing, necessitating more tailored approaches to ensure efficient network operation.

Drawing on the innovative capabilities of the NVIDIA Spectrum-4 Ethernet switch and BlueField-3 SuperNIC, NVIDIA Spectrum-X delivers end-to-end RoCE adaptive routing (AR). This state-of-the-art technology adjusts how data is routed, dynamically modifying the routing offlows, significantly enhancing network efficiency. Instead of statistically distributing flows between available paths based on the characteristics of the traffic, the Spectrum-4 switch performs packet spraying on all available paths for every flow.

This innovation guarantees optimal routing for AI's elephant flows, since every such flow is effectively routed on a packet-by-packet basis by utilizing all available route options.

With packets belonging to the same flow now routed through different paths, it is very likely that they will reach their destination out-of-order. To address this challenge effectively, the BlueField-3 SuperNICs and Spectrum-4 operate in close coordination. Spectrum-X employs a unique technique that seamlessly reorders data packets on the receiving ends at 400Gb/s speed. This technology is crucial for enabling packet-by-packet routing on multiple paths in high-speed networks.

The following figure illustrates the impact of RoCE Adaptive Routing on RDMA applications. It highlights issues with static, hash-based forwarding, such as uplink port collisions, longer flow completion times, reduced bandwidth, and fairness problems. These issues are resolved with adaptive routing. In the ECMP scenario, some flows have significantly longer completion times and lower bandwidth due to collisions, while RoCE Adaptive Routing ensures consistent completion times and similar peak bandwidth for all flows.

Figure 6. Flow Completion Time with ECMP Static Routing and RoCE Adaptive Routing

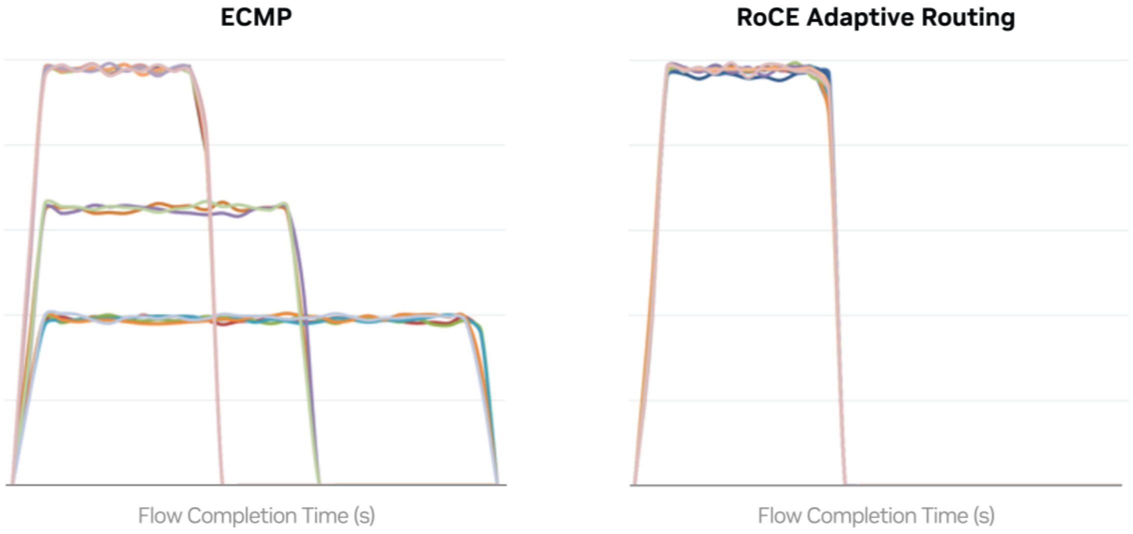

The following figure illustrates the process of end-to-end adaptive routing and out-of-order packet delivery. Data packets are sent from nodes A and B to nodes C and D, respectively. The first step illustrates how those packets are transmitted from the sending nodes and sent into the switch network. Following this, the receiving switch evenly distributes packets on all four available routes. The illustration makes it evident that packets of the same flow—marked in green and purple—are routed along different paths, which is the root cause of out-of-order receipt on the receiving ends. The final part of the figure shows how the BlueField-3 SuperNICs on the receiving nodes successfully reorder the packets (D, C, B, A à A, B, C, D).

Figure 7 **End-to-end RoCE Adaptive Routing and Packet Reordering**

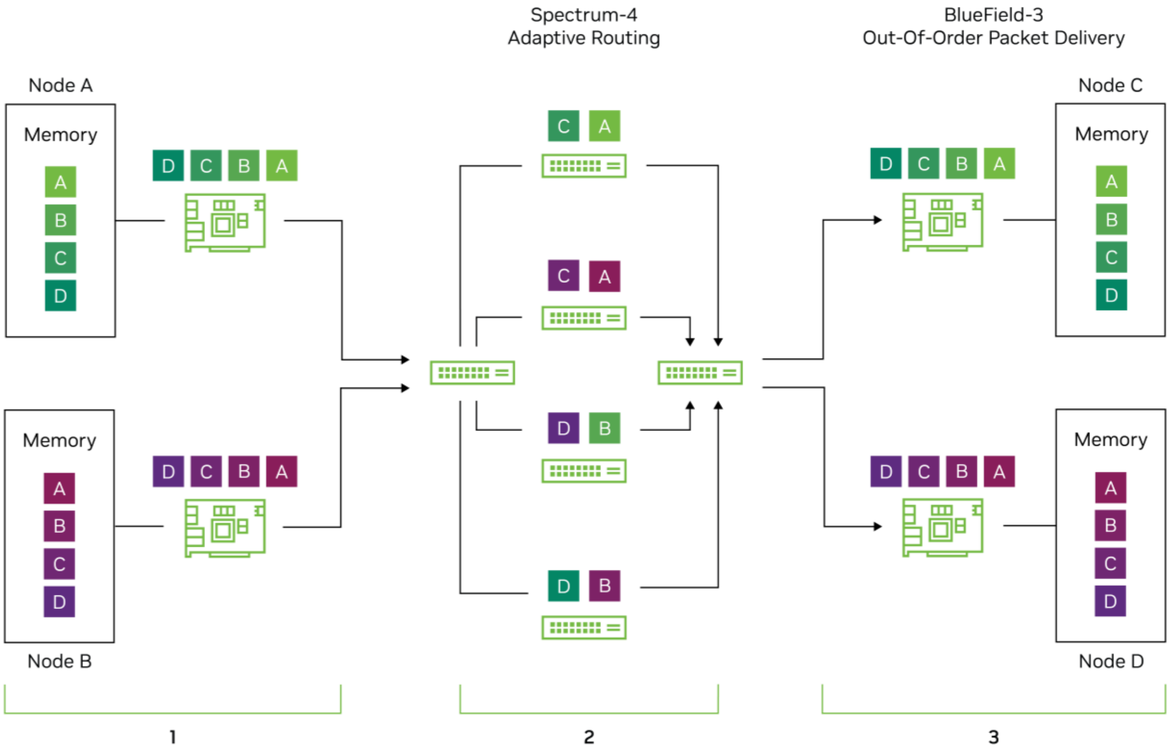

RoCE adaptive routing and packet reordering, enabled through the seamless integration of the Spectrum-4 switch and BlueField-3 SuperNIC, are essential to achieving high network performance and efficiency, which accelerates Ethernet AI workloads.

Advanced Congestion Control

Ethernet networks, by design, are prone to congestion. This can hinder the efficient transfer of data and slow down AI training and inference. A key factor in building optimal networks for AI, is how they manage and prevent congestion. Before diving into the solutions tailored for AI, it's worth noting how traditional Ethernet cloud networks handle the issue. In typical TCP/IP communication, the TCP protocol uses control mechanisms such as flow control and sliding windows to ensure the sender does not overwhelm or overload the receiver with data.

In contrast, networks designed for AI have their own unique set of challenges. They utilize RoCE for GPU-to-GPU communication. RoCE performance thrives in highly reliable, low-latency network environments, which means they require sophisticated congestion control to manage network traffic during congestion events effectively. Furthermore, AI clouds are especially prone to congestion due to their multi-tenant nature. When multiple tenants run jobs concurrently, congestion can swiftly accumulate due to the actions of a single tenant's job. This congestion can have a cascading effect, increasing latency and constraining the available network bandwidth available for AI tasks.

In addition, AI model training exhibits a distinct traffic pattern characterized by its highly bursty nature. This phenomenon arises from the intricacies of AI collective operations, where many GPU nodes collaborate to efficiently distribute the computing workload. Through a series of highly correlated communication flows, the participating nodes exchange data such as model parameters or aggregated results, triggering intensive bursts of network traffic. The following figure illustrates this bursty network traffic, as observed through sampling a top-of-rack switch port in an AI test cluster. This unique nature of traffic renders conventional congestion control methods inefficient, especially when these traffic bursts can occur as frequently as every few milliseconds.

Figure 8 AI Model Training Exhibits a Distinct Traffic Pattern Characterized by Network Traffic Bursts

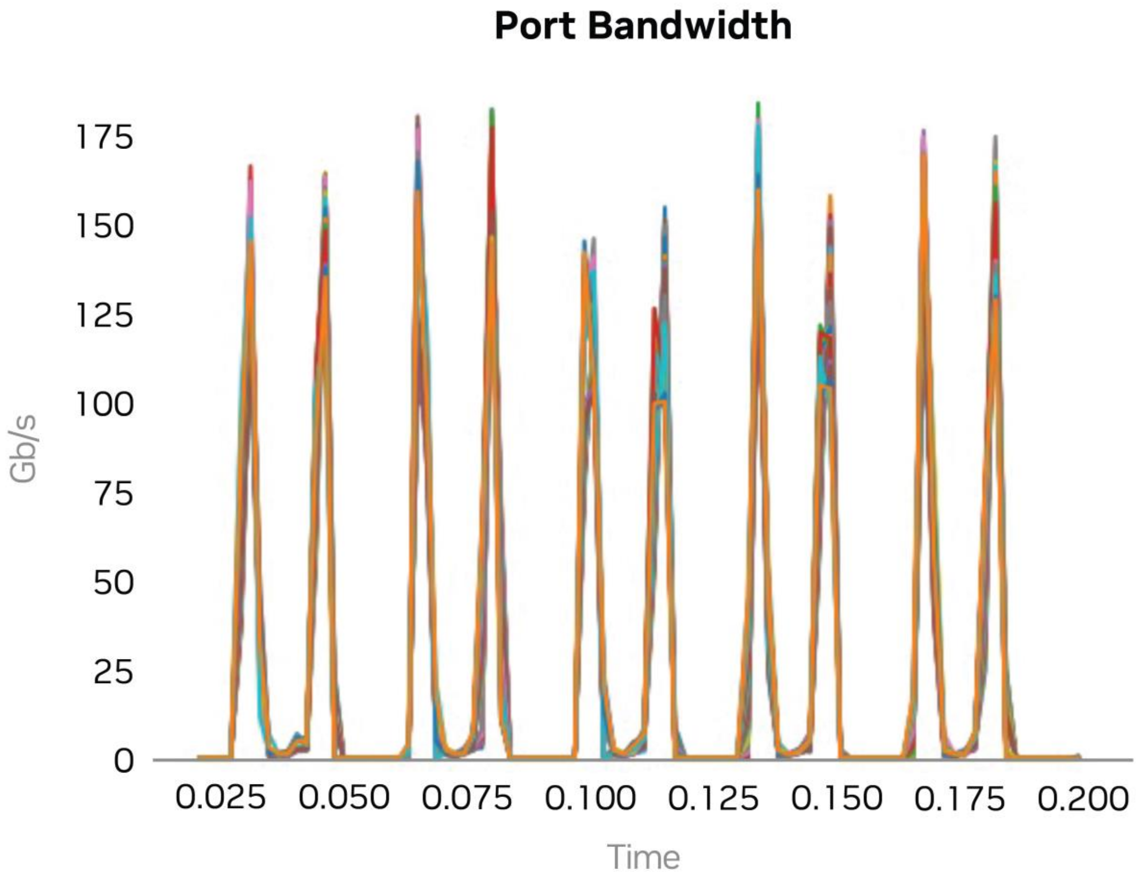

Considering these challenges, many cloud environments have turned to methods like Data Center Quantized Congestion Notification (DCQCN) to proactively identify and react to network congestion. DCQCN employs Explicit Congestion Notification (ECN) to notify sending devices of forthcoming network congestion before packet loss occurs. While DCQCN is proven to be a practical method to alleviate network congestion, it falls short in meeting the demands of generative AI clouds, where traffic patterns are so bursty that notifications might arrive after congestion has occurred.

Facing the challenges of network congestion in AI cloud infrastructures, Spectrum-X introduces a powerful solution driven by the tight integration of the BlueField-3 SuperNIC and Spectrum-4 switch. This solution significantly enhances network efficiency and accelerates AI workload performance. At its core, the solution primarily relies on an innovative network-aware congestion algorithm that utilizes real-time telemetry data streamed from network switches to manage and prevent network congestion. It taps into the Spectrum-4 switch's in-band telemetry capabilities to inform the sender's BlueField-3 SuperNIC about the current network utilization status. In the event of congestion buildup, it sends timely alerts. Based on this information, the SuperNIC adjusts transmission rates as needed, ensuring congestion doesn't spread further. Importantly, BlueField-3 SuperNICs execute the congestion control algorithm, handling millions of congestion control events per second with microsecond reaction latency and making precise rate decisions.

By integrating this network-aware congestion control approach, NVIDIA optimizes network bandwidth and minimizes latency, thereby accelerating AI training and inference workloads. Additionally, it ensures consistent and predictable performance for every tenant job by streamlining communication transmissions at data center scale.

Figure 9 Spectrum-4 Detects Congestion Spots in Real-Time; BlueField-3 Adjusts the Transmission Rate

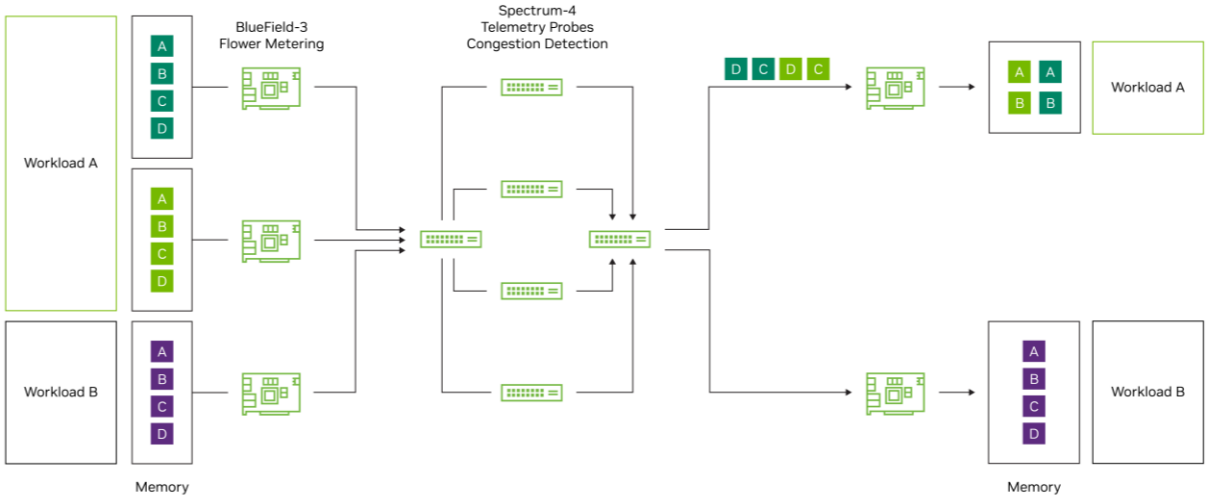

In response to the burstiness of AI collective operations, Spectrum-X has introduced an innovative telemetry technology to discover network congestion at a very granular level. Unlike traditional telemetry tools, which struggle to keep pace in dynamic AI networks, Spectrum-X's telemetry gathers comprehensive, high-frequency data that is leveraged to enhance data transmission and optimize network efficiency. The following figure illustrates Spectrum-X's approach to telemetry collection. In this example, AI communication bursts occur approximately every 5 milliseconds. Each blue dot represents a network transmission rate at a specific time. This high-frequency sampling is essential for revealing the bursty nature of AI networks and effectively managing congestion at the data center level.

Figure 10 Spectrum-X's Telemetry Samples Data in High-Frequency to Uncover Bursty AI Network Flows

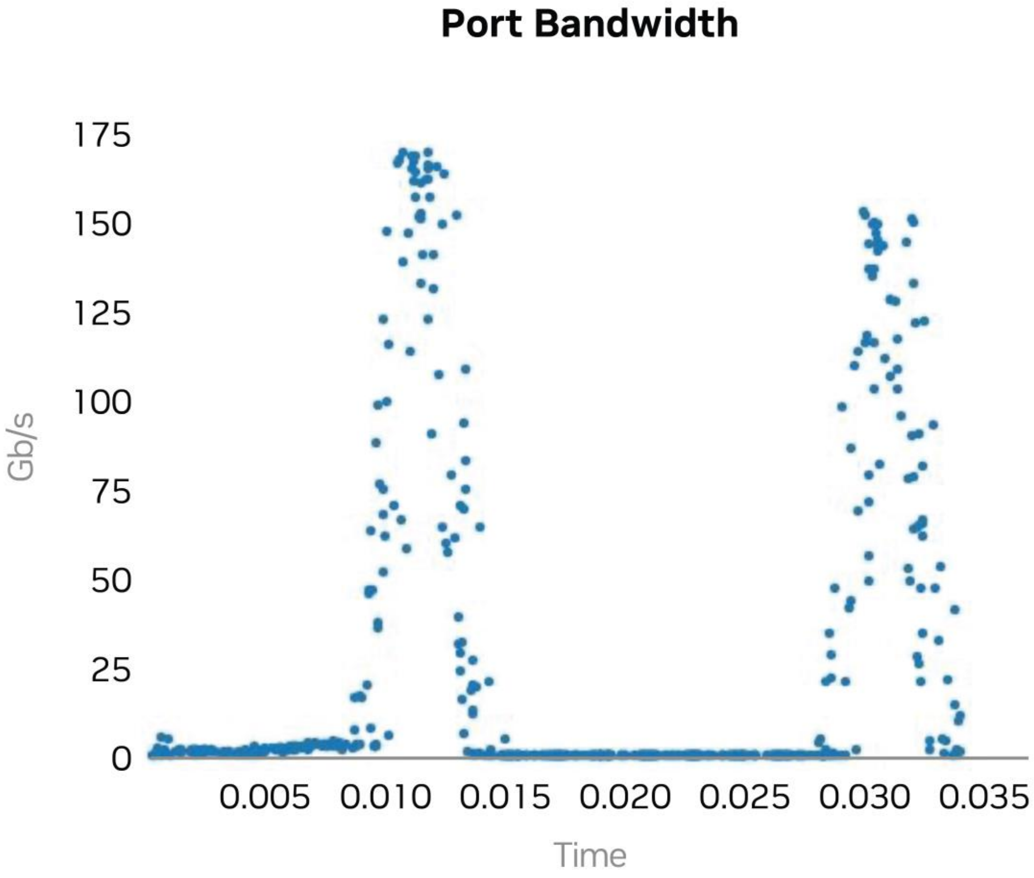

Taking a closer look at the BlueField-3 SuperNIC reveals its advanced features. This advanced network accelerator prioritizes full programmability to meet the demands of users seeking to implement customized congestion control algorithms. The need for programmability arises from the distinctive traffic patterns typical of AI workloads and data center network topologies, which require further optimization to achieve peak performance.

By integrating an advanced Datapath Accelerator (DPA), the BlueField-3 SuperNIC provides a dedicated compute engine, optimized for I/O-intensive and low-code packet processing. It allows users to run home-grown congestion control algorithms tailored to their network environments. Powered by the BlueField-3's DPA processor, the programmable model is built upon three key components:

1. DOCA Programmable Congestion Control (PCC) software library: This library equips developers with the necessary tools to formulate AI-optimized congestion control algorithms on the DPA processor.
2. DOCA DPA: This compute engine, integral to BlueField-3 is optimized for I/O-intensive and low-code packet processing. It allows users, whether they are customers or partners, to run home-grown congestion control algorithms tailored to their network environments.
3. DOCA FlexIO driver—As a low-level SDK, this driver is essential for loading DPA programs into the DPA, managing the DPA's memory, and creating the execution handlers, as well as the needed hardware rings and contexts.

The combination of the DPA processor and the multi-layer DOCA programmable model elevates the potential of the BlueField-3 SuperNIC platform, positioning it as a pivotal tool to advance networking for the next wave of AI applications.

Secure and Accelerated VPC Networking for AI Computing

Modern cloud data centers are designed from the ground up to be secure, multi-tenant computing environments. General-purpose clouds use various network technologies to establish virtual private clouds (VPCs), ensuring tenant traffic is properly isolated. However, the advent of AI clouds, with their dedicated AI compute networks, adds complexity to the creation of VPCs and the isolation of tenant workloads.

First and foremost, AI compute networks must provide high-throughput and low-latency connectivity for GPU servers. Software-defined network solutions running solely on CPUs are inadequate for the high-performance connectivity essential in AI compute networks. In addition, as many AI cloud environments offer bare-metal as-a-service (BMaaS), it becomes impractical to deploy tenant networking software on the compute nodes. Recognizing this limitation, many bare-metal cloud environments deploy EVPN VXLAN on network switches as an effective way to establish tenant isolation. While this provides a solution for AI compute networks, it lacks advanced functionality like access-lists and security groups, and it does not scale effectively, particularly when expanding to tens-of-thousands of GPUs.

The BlueField-3 SuperNIC equips cloud architects with the tools to overcome these challenges by implementing secure, zero-trust VPC networking tailored for the AI compute plane. With accelerated switching and packet processing (ASAP^2) technology, the BlueField-3 SuperNIC enables the best of both worlds, software-defined and hardware-accelerated network connectivity. The NVIDIA ASAP^2 technology stack provides a range of network acceleration capabilities and full programmability through the DOCA FLOW SDK. The SuperNIC delivers several orders of magnitude faster performance compared to non-accelerated network environments. Out-of-the-box, the BlueField-3 SuperNIC offers two paths to create secure, multi-tenant, and high-performance AI compute network environments:

- OVS (Open vSwitch) / OVN (Open Virtual Network)-based SDN acceleration solution
- VXLAN (Virtual Extensible LAN) - EVPN (Ethernet VPN)-based network solution

While both SDN and EVPN VXLAN are designed to create multi-tenant networks, they each do it differently. SDN centralizes control and abstracts network resources, whereas EVPN VXLAN distributes control using a BGP-based control plane coupled with MAC learning.

The BlueField-3 SuperNIC seamlessly offloads and accelerates both SDN and EVPN-based solutions. In both cases, the software stack runs exclusively on the SuperNIC and is powered by the following DOCA capabilities:

Figure 11. BlueField-3 SuperNICs Enable Secure and Accelerated Networks for Multi-Tenant AI Clouds

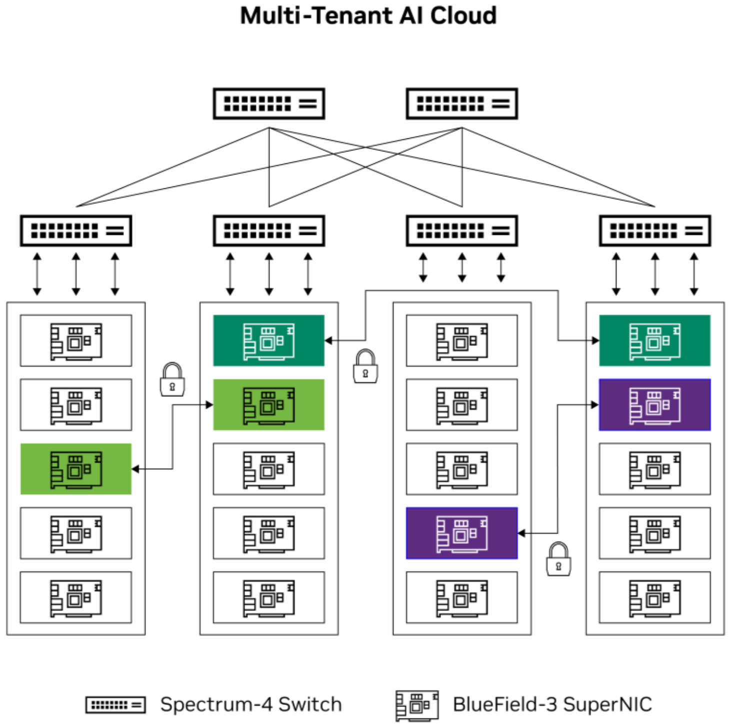

In addition, the BlueField-3 SuperNIC provides inline IPsec encryption acceleration at speeds of up to 400Gb/s. What's notable about this acceleration engine is its compatibility with other inline accelerations, allowing AI cloud builders to encrypt all East-West communications on the AI compute network. This not only adds an extra layer of cyber protection but also increases the security posture of the AI platform. The DOCA IPsec software library provides an API for developers to enable BlueField-accelerated flow encryption and decryption.

The BlueField-3 SuperNIC is ideally suited for securing and accelerating VPC networking in any cloud environment, especially in bare-metal, multi-tenant AI clouds. The integrated compute subsystem within the network I/O path lays a secure foundation for deploying tenant networking solutions and enforcing fine-grained network policy, further bolstering the overall security of the AI cloud platform.

Power-Efficient, Low-Profile Design

Modern data centers are increasingly constrained caused by the power demands of cooling systems, rising energy costs, and limited energy supply. By combining GPU-accelerated AI computing with BlueField-3 SuperNICs, data centers can efficiently accommodate more AI workloads while operating within a tight power budget.

Featuring a sub-75-watt, half-height, half-length (HHHL) PCIe form factor, the BlueField-3 SuperNIC (model B3140H) is designed with energy efficiency in mind. Compatible with most enterprise-class servers, its distinct low-profile design facilitates effective scaling to match the number of GPUs in a system. Large-scale GPU systems can host up to eight SuperNICs in a single server chassis, accelerating AI performance by providing full 400Gb/s bandwidth for each GPU in the system. To ensure high-availability of mission-critical systems, the BlueField-3 SuperNIC (model B3220L) is also available in a full-height, half-length (FHHL) PCIe form factor, providing two ports of 200Gb/s connecting to two network switches.

In addition, the SuperNIC doesn’t require external power, relying solely on the PCIe bus. This simplifies the server hardware design and allows for more efficient SuperNIC scaling in power-constrained environments.

BlueField-3 SuperNIC Powers NVIDIA-Accelerated Systems

NVIDIA's accelerated computing systems are designed to meet the stringent performance requirements of a growing range of complex and diverse AI applications, all while ensuring cloud manageability and security. Embodying these principles, NVIDIA is integrating BlueField-3 SuperNICs throughout its data center platforms, including HGX™ H100, OVX™ L40S, etc., designed to deliver accelerated AI performance in cloud and enterprise environments. By integrating the BlueField-3 networking platforms, NVIDIA capitalizes on its expertise in accelerated computing and AI, streamlining the network traffic flow, accordingly:

- BlueField-3 SuperNIC for East-West (E-W) traffic: Handles traffic between AI systems inside the cluster, typically used for distributed AI training, collective operations, and other AI computational tasks.
- BlueField-3 data processing units (DPUs) for North-South (N-S) traffic: Manages user traffic to and from the AI cluster, including connections to external resources like cloud management systems, remote data storage nodes, and other data center environments, or Internet connections.

The following diagrams illustrate the optimized networking model for AI clouds, and the NVIDIA HGX H100 system configurations, integrating BlueField-3 SuperNICs and DPUs.

Figure 12. NVIDIA-Accelerated Computing Systems Provide an Optimized Networking Model for AI Clouds

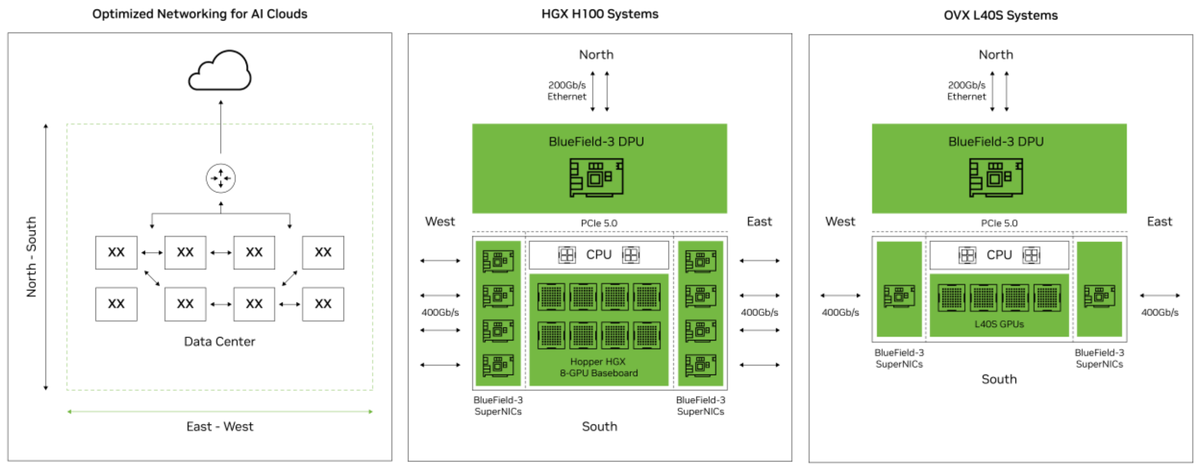

Most NVIDIA HGX H100-class systems host eight BlueField-3 SuperNICs (B3140H), one for every GPU in the system. This 1:1 configuration ensures extreme AI performance with every GPU benefiting from 400Gb/s RoCE bandwidth with all the capabilities previously discussed. L40S-based systems follow a 2:1 configuration where every two GPUs share one BlueField-3 SuperNIC, for a total bandwidth of 200Gb/s per GPU.

Conclusion

In the age of generative AI, the necessity for GPU-accelerated computing is undisputed. GPU-accelerated computing stands as a cornerstone. Yet, traditional Ethernet often falls short in delivering the requisite network performance and efficiency demands of AI computing. To bridge this gap, NVIDIA offers Spectrum-X, an optimized Ethernet networking platform for enhancing AI clouds. At its core, the BlueField-3 SuperNIC emerges as a novel network accelerator, purpose-built to supercharge hyperscale AI workloads. It ensures lightning-fast and efficient communication between GPU servers, for peak AI workload efficiency. Furthermore, the SuperNIC establishes a secure multi-tenant data center environment, while guaranteeing deterministic and isolated performance for tenant jobs. With its power-efficient HHHL PCIe design, the BlueField-3 SuperNIC seamlessly integrates into flagship NVIDIA- accelerated systems, significantly boosting performance for AI traffic on the east-west network inside the cluster.

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 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 NONINFRINGEMENT, 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, NVIDIA Spectrum, BlueField, are trademarks and/or registered trademarks of NVIDIA Corporation and/or its affiliates 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

© 2023 NVIDIA Corporation. All rights reserved. NOV

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com