

Mango StorageBoost™ NVMe over TCP Initiator

Enhance Your Infrastructure with
Ultra-scale Disaggregated Storage System

Why Disaggregated Storage?

High-performance and High-capacity

With the rapid increase in data production, driven by advancements in AI and the growing complexity of large language models (LLMs), efficient storage management has become a pressing challenge. LLMs demand significant amounts of data and computational resources, placing additional strain on traditional storage infrastructures. In response, storage disaggregation technologies are emerging as a key solution, enabling the separation of compute and storage to enhance system efficiency.

Lower CapEX and OpEX

Disaggregated storage effectively lowers capital expenditures (CapEx) by mitigating over-provisioning and facilitating more efficient utilization of storage resources. Rather than maintaining excess capacity for each individual server, centralized storage pools can be dynamically allocated, enhancing overall resource utilization. Furthermore, disaggregated storage contributes to a reduction in operational expenditures (OpEx) by enabling independent management of storage resources separate from CPUs. This separation streamlines management processes and minimizes operational inefficiencies, allowing for the independent updating or expansion of storage components without necessitating a complete overhaul of the entire server system.

Why NVMe/TCP?

NVMe over Fabrics (NVMe-oF) has emerged as a leading technology for storage disaggregation in modern data centers. Among the various transports available, NVMe/TCP is widely recognized for its scalability and ease of deployment, particularly in comparison to RDMA-based protocols like RoCEv2, which require specialized network hardware such as RDMA network cards and switches [1, 2, 3, 4]. This characteristic positions NVMe/TCP as a more accessible and flexible solution for many organizations. However, despite these advantages, NVMe/TCP has historically encountered performance challenges due to higher overheads, which can hinder its efficiency for high-performance workloads such as artificial intelligence and machine learning.

Mango StorageBoost™ - NVMe/TCP Initiator (NTI) addresses these challenges by offering an efficient and easy-to-use disaggregation solution for NVMe/TCP initiators.

Hardware-accelerated Storage Disaggregation Solution

Enhance your Datacenter without
Compromising Performance

Key Highlights

Unparalleled Performance

4KB (IOPS)	Random Read	5.7M
	Random Write	5.7M
	Random RW Mix	10.3M
128KB (GB/s)	Sequential Read	23.6
	Sequential Write	23.5
	Sequential RW Mix	41.6

Easy-to-Deploy Solution: No Modification on Infrastructure and Application

- Expose DPU as a standard NVMe-PCIe device to user
 - > Compatible with NVMe 1.3/2.0
- Expose DPU as a standard NVMe/TCP Initiator to NVMe/TCP target servers
 - > Compatible with NVMe-oF 1.0c
 - > Compatible with standard TCP/IP network

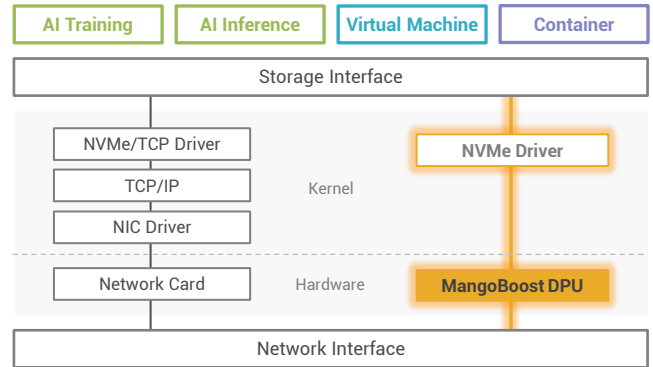
Configurable and Adaptable Solution

- Specification can be tailored per customer's request

Seamless Integration

Mango StorageBoost™ - NVMe/TCP Initiator is seamlessly integrated into existing storage systems. By presenting the DPU as a standard NVMe PCIe device, the storage systems are able to utilize NTI with no SW modification. On the target side, NTI can connect to any NVMe/TCP target servers through a standard TCP/IP network and Ethernet switches.

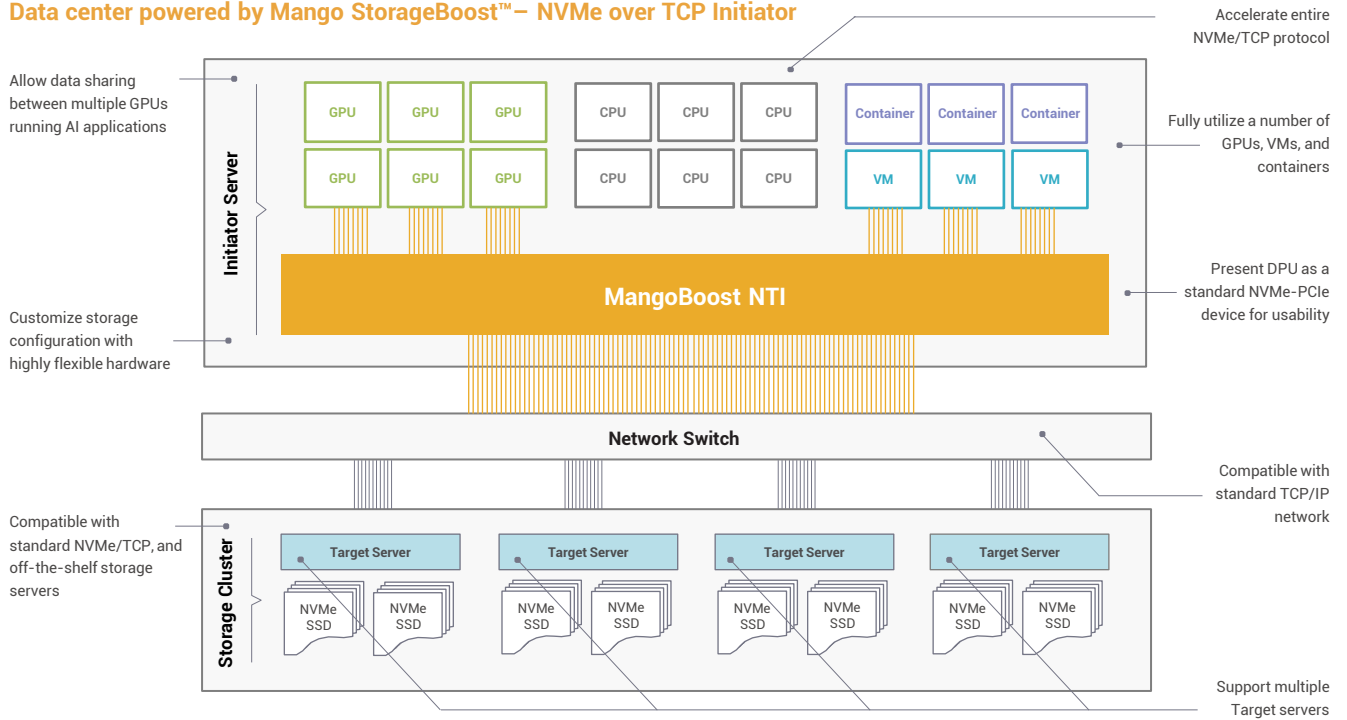
Effortless integration of NTI as a replacement for existing storage stacks, leveraging its compatibility with both storage and network interfaces.



Fast and Flexible Storage Solution with Disaggregated Infrastructure

Mango StorageBoost™ - NVMe/TCP Initiator is designed to meet the rigorous demands of diverse workloads, delivering exceptional performance and usability. With its high bandwidth and low latency, NTI fully leverages accelerators such as GPUs and supports cloud services, including virtual machines and containers, as well as targeted storage systems. NTI integrates seamlessly into existing storage networks without necessitating modifications to the current infrastructure. This capability empowers modern data centers to fully harness their potential, optimizing both energy efficiency and cost-effectiveness.

Data center powered by Mango StorageBoost™ – NVMe over TCP Initiator



READY TO GET STARTED?

To learn more about accelerating your AI infrastructure, visit mangoboost.io

DISCLAIMERS

The performance claims in this document are based on the internal cluster environment. Actual performance may vary depending on the server configuration. Software and workloads used in performance tests may have been optimized for performance only on MangoBoost products. Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. Results that are based on pre-production systems and components as well as results that have been estimated or simulated using MangoBoost reference platform for informational purposes only. Results may vary based on future changes to any systems, components, specifications, or configurations. Statements in this document that refer to future plans or expectations are forward-looking statements. These statements are based on current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in such statements. MangoBoost does not guarantee any specific outcome. Nothing contained herein is, or shall be relied upon as, a promise or representation or warranty as to future performance of MangoBoost or any MangoBoost product. The information contained herein shall not be deemed to expand in any way the scope or effect of any representations or warranties contained in the definitive agreement for MangoBoost products.

The information contained herein may not be reproduced in whole or in part without prior written consent of MangoBoost. The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. MangoBoost assumes no obligation to update or otherwise correct or revise this information and MangoBoost reserves the right to make changes to the content hereof from time to time without any notice. Nothing contained herein is intended by MangoBoost, nor should it be relied upon, as a promise or a representation as to the future.

MANGOBOOST MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

— Industry-standard Benchmark and Real-World Use Cases

Discover how **Mango StorageBoost™ - NVMe/TCP Initiator** plays a crucial role in enhancing AI infrastructure and optimizing AI systems.

1. MLPerf™ Storage benchmark

MLPerf™ Storage benchmark measures how fast storage systems can supply training data when a model is being trained. MangoBoost submitted the **industry-first** DPU-accelerated NVMe/TCP Storage system on MLPerf Storage v1.0 [5], demonstrating the **best Ethernet-based system result**.

2. Mango StorageBoost™ - GPU Storage Boost (GSB)

Mango StorageBoost™ - GPU Storage Boost (GSB) is a software solution for efficient data movement between GPUs and storage. When combined with **Mango StorageBoost™ - NVMe/TCP Initiator**, GSB enables highly efficient GPU data movement, even over remote storage systems based on the TCP/IP transport layer without any CPU involvement. MangoBoost co-presented with AMD at SDC 2024 [6], showcasing large language model (LLM) training utilizing AMD MI300X GPUs, enhanced by the **Mango StorageBoost™** solution.

[1] Answering Your Questions: NVMe™/TCP: What You Need to Know About the Specification, <https://nvmeexpress.org/answering-your-questions-nvme-tcp-what-you-need-to-know-about-the-specification-webcast-qa/>

[2] Pavilion compares RoCE and TCP NVMe over Fabrics performance, <https://blocksandfiles.com/2018/08/16/pavilion-compares-roce-and-tcp-nvme-over-fabrics-performance/>

[3] SRNIC: A Scalable Architecture for RDMA NICs, NSDI 23, <https://www.usenix.org/conference/nsdi23/presentation/wang-zilong>

[4] RDMA over Ethernet for Distributed Training at Meta Scale, SIGCOMM 24, <https://dl.acm.org/doi/abs/10.1145/3651890.3672233>

[5] <https://mlcommons.org/2024/09/mlperf-storage-v1-0-benchmark-results/>

[6] Accelerating GPU Server Access to Network-Attached Disaggregated Storage using Data Processing Unit (DPU), <https://www.sniadeveloper.org/events/agenda/session/666>

READY TO GET STARTED?

To learn more about accelerating your AI infrastructure, visit mangoboost.io

DISCLAIMERS

The performance claims in this document are based on the internal cluster environment. Actual performance may vary depending on the server configuration. Software and workloads used in performance tests may have been optimized for performance only on MangoBoost products. Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. Results that are based on pre-production systems and components as well as results that have been estimated or simulated using MangoBoost reference platform for informational purposes only. Results may vary based on future changes to any systems, components, specifications, or configurations. Statements in this document that refer to future plans or expectations are forward-looking statements. These statements are based on current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in such statements. MangoBoost does not guarantee any specific outcome. Nothing contained herein is, or shall be relied upon as, a promise or representation or warranty as to future performance of MangoBoost or any MangoBoost product. The information contained herein shall not be deemed to expand in any way the scope or effect of any representations or warranties contained in the definitive agreement for MangoBoost products.

The information contained herein may not be reproduced in whole or in part without prior written consent of MangoBoost. The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. MangoBoost assumes no obligation to update or otherwise correct or revise this information and MangoBoost reserves the right to make changes to the content hereof from time to time without any notice. Nothing contained herein is intended by MangoBoost, nor should it be relied upon, as a promise or a representation as to the future.

MANGOBOOST MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.