

Stand alone NVMe-oF Acceleration Solution



U50 based 100Gb ethernet NVMe Storage JBOF with Application Acceleration

INTRODUCTION

The Xilinx NVMe-over-Fabric (NVMe-oF™) reference design and U50 solution was created with the idea of adding computational storage into next generation networked storage solutions. Using Remote Direct Memory Access (RDMA) this design provides a low latency, high performance, industry standard interconnect for up to 24 NVMe SSDs. This platform provides the flexibility to define custom acceleration functions within a NVMe-oF compliant environment and eliminates the need for an external processor or Network-Interface-Card (NIC) in a Just-a-Bunch-of-Flash (JBOF) enclosure enabling a highly integrated and cost-effective storage solution.

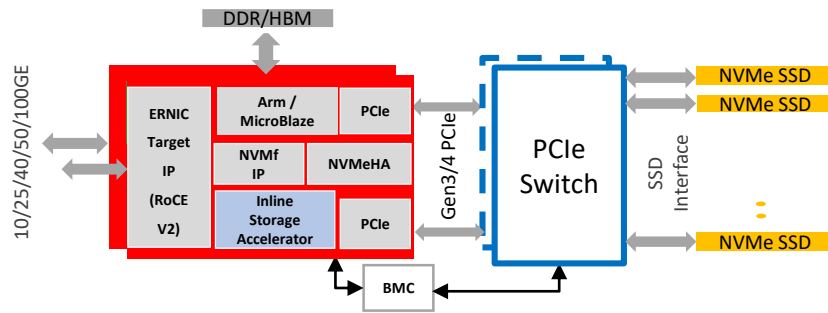


Figure 1

SOLUTION OVERVIEW

The design leverages a standard NVMe-oF protocol and standard Ethernet ports to present storage to a server that looks like a local SSD, but is in fact a remote storage namespace. This enables the efficient pooling and sharing of storage resources across datacenter servers. These remote namespaces can dramatically reduce the storage cost, footprint and power within datacenters. With additional capability for hardware acceleration service, the Alveo U50 solution becomes a true disaggregated computational storage accelerator.

This Xilinx solution provides reliable transport of NVMe frames with low latency, high throughput and massive scalability to remote hosts. A block diagram of a typical system solution is depicted above in Figure 1. The Xilinx NVMe-oF reference design implements the NVM express over fabric protocol and the RDMA NIC protocol in the single highly integrated Xilinx FPGA contained in an Alveo U50 add in card with a significant amount of programable logic remaining for use as computational storage accelerators.

The key data transfer commands in the NVMe-oF protocol are offloaded entirely to hardware, while the software running on an embedded CPU in the Xilinx device processes the control commands. This gives this Xilinx solution significant performance advantages over processor-only implementations. The Xilinx ERNIC IP integrated into this reference design provides reliable transport, flexibility in network interconnect and the performance to support line speed bandwidth. An implementation supporting 24 drives can be accomplished using just over 30% of the U50 resources, leaving the rest of for customizable acceleration engines or other types of differentiation.

SOLUTION BRIEF



- 100Gb/s NVMe-oF Alveo U50 based solution
- Support for up to 24 NVMe SSDs
- Six U50s in a single chassis
- 100GbE line rate performance
- Ability to add Hardware Accelerators

Stand alone NVMe-oF Acceleration Solution



U50 based 100Gb Ethernet NVMe Storage JBOF with
Application Acceleration

SOLUTION DETAILS

Feature Overview	Description
Supported RDMA Protocol	RoCEv2
Network Interface	Up to 100Gb Ethernet 100GbE, 50GbE, 40GbE, 25GbE and 10GbE
SSD Interface	Up to one PCIe Gen3 x16 or two Gen4 x8 interfaces (up to 24 SSDs with the use of a PCIe switch)
Number of Hosts	A maximum of 128 hosts
Send and Receive Queue-Pairs (QPs)	Up to 255, which is QP1- QP255
Completion Queues (CQ)	Up to 255
Completion Queues (CQ) Queue Depth	64 entries per queue
Latency	1 us additional latency per NVMe command traversing the IP. Networking latency, SSD latency, and number of commands in a transfer are additional.
Performance (Single U50, 4KB, 8 SSDs, OIO=64)	2.5M IOPs, 10 GB/s
Management Interfaces	SMBus, NVMe-MI
Future support	NVMe 1.4+ NVMe 1.1+ NVMe/TCP
Inline Accelerator Examples	Storage services: <ul style="list-style-type: none">• (De)Compression• (De)Encryption• Data protection Database Acceleration: <ul style="list-style-type: none">• Scan• Filter• Aggregate
Resource Utilization	The full solution resources include NVMF-IP, ERNIC-IP, CMAC, NVMeHA, AXI-DMA and DDR-MIG IP. The resource utilization depends on selected configuration and additional accelerators

TAKE THE NEXT STEP <https://www.xilinx.com/products/intellectual-property/nvmeof.html>

Learn more about [Alveo accelerators](https://www.xilinx.com/products/boards-and-kits/alveo.html) <https://www.xilinx.com/products/boards-and-kits/alveo.html>



Adaptable. Intelligent.