

Eideticom NoLoad™ FPGA Accelerator

Product Brief

Overview

- Eideticom NoLoad™ (NVMe Offload) FPGA accelerator platform
- NVMe 1.3 compatible interface provides seamless integration for all CPU platforms
- Support for acceleration of storage workloads i.e. Erasure Coding, Deduplication & Compression
- Hardware Eval Kits (HEKs) available now targeting COTS PCIe cards from Nallatech® & AlphaData®
- Compatible/validated with Broadcom®, Mellanox® and Q-Logic® RDMA NIC's

Capacities

- 1.5 - 32 GB RAM Drive
- 0.5 – 16 GB Controller Memory Buffer (CMB)
- Up to 4 TB Flash memory supported on some HEK SKU's

Capabilities

- ISA-L compliant RS Erasure Coding engine
- Deduplication - support for SHA-1, SHA-2 & SHA-3 (with hashing)
- LZ77 + Huffman encoding GZIP compliant compression
- Library of dev tools including 256 TB drive simulator

NVMe Feature Support

- NVMe 1.3 compatible interface
 - o Admin queue
 - o 16 I/O queues
 - o Supports NVMe Scatter Gather Lists (SGLs)
 - o CMB support (all modes)

Performance

- Available under NDA

Eval Kit Target Platforms

NoLoad-HEK-NT250S

- Nallatech® 250S FlashGT+ Card target
- Add In Card (AIC) Half Height Half Length (HHHL) Form Factor
 - o Single slot x8 connector
- PCIe Gen3x8 or Gen4x8 (also supports x4 mode)
- Xilinx® KU15P Kintex® Ultrascale Plus™ FPGA
- Up to 4 TB of Flash Memory

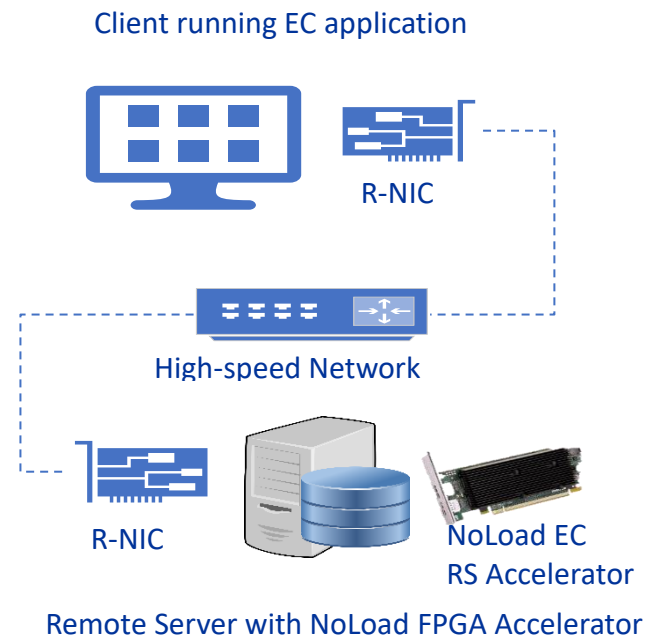
NoLoad-HEK-AD9V3

- AlphaData® ADM-PCIE-9V3 Card target
- AIC HHHL Form Factor
 - o Single slot x16 connector
- PCIe Gen3x16 or PCIe Gen4x 8 (also supports x4 mode)
- Xilinx® VU3P Virtex® Ultrascale Plus™ FPGA

The Disaggregation of FPGA Accelerators using NoLoad™ and NVMe over Fabrics

Get your FPGA's "out of the box" and shared across the datacenter

- Emerging NVMe over Fabrics ecosystem allows NoLoad NVMe accelerator namespaces to be accessed/shared across network fabrics such as Ethernet
- NoLoad™ FPGA acceleration sharing across the network fabric enables FPGA disaggregation
- Eideticom demo showcases a local client accessing a NoLoad™ Reed-Solomon (RS) Erasure Code (EC) accelerator on a remote server via RDMA NVMe over Fabrics



NoLoad™ and Peer-2-Peer processing

The case for Peer-2-Peer processing

- PCIe End-Points (EPs) are getting faster and faster e.g. GPGPUs, RDMA NICs, NVMe SSDs
- Bounce buffering all IO data in system memory is a waste of system resources and reduces QoS for CPU memory (the noisy neighbor problem)

The solution:

- NoLoad™ plus p2pmem Linux kernel framework for allowing PCIe EPs to DMA to each other whilst under host CPU control
- CPU/OS still responsible for security, error handling etc.
- 99.99% of DMA traffic now goes direct between EPs

