Product Brief

NoLoad™ U.2 Computational Storage Platform

Overview

Eideticom's NoLoad™ Computational Storage FPGA accelerator platform in a 2.5" U.2 NVM Express® (NVMe) drive form factor

NoLoad's NVMe compatible interface provides seamless integration for all CPU platforms and has been validated on Intel, AMD, ARM and IBM Power8/9 CPUs

NoLoad supports high performance acceleration of storage and compute workloads on FPGA i.e. Erasure Coding, Deduplication, Compression, Analytics, etc

Compatible/validated with Broadcom®, Mellanox® and Q-Logic® RDMA NIC's

Validated at PCIe Gen4 rates with Gen4 compliant hosts

Capacities

- 1.5 - 32 GB RAM Drive
- 0.5 – 16 GB NVMe Controller Memory Buffer (CMB)

Capabilities

- GZIP/ZLIB/Deflate compliant compression core: >3 GB/s
- GUNZIP/ZLIB/Inflate decompression core: >3 GB/s
- ISA-L compliant RS Erasure Coding engine: >3 GB/s
- Deduplication - support for SHA-1, SHA-2 & SHA-3 (with hashing)
- Supports easy integration of user specified and developed acceleration functions

NVMe Feature Support

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

Performance

- Available under NDA

NoLoad™ U.2 and Peer-2-Peer processing

The case for Peer-2-Peer (P2P) processing

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

The solution:

- NoLoad™ U.2 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
- Application: P2P Compression offload

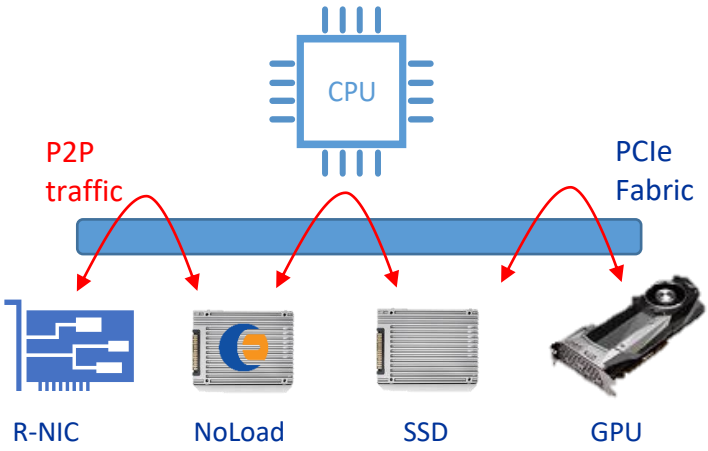

The Disaggregation of FPGA Accelerators using NoLoad™ and NVMe over Fabrics

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 client accessing NoLoad™ accelerators on a remote server via RDMA or TCP/IP NVMe over Fabrics
- Eideticom's API supports PCIe and/or Ethernet/RDMA/TCP with no changes in application code