
Accelerate Everything.

**Computational Storage Is the Answer for Huge Data and Deep Problems (Apparently)**

Andrew Maier, Director of Solutions, Flash Memory Summit 2019 (with some help from Stephen Bates ;-))

# B.C.S
# (Before Computational Storage)

# Let There Be Light!

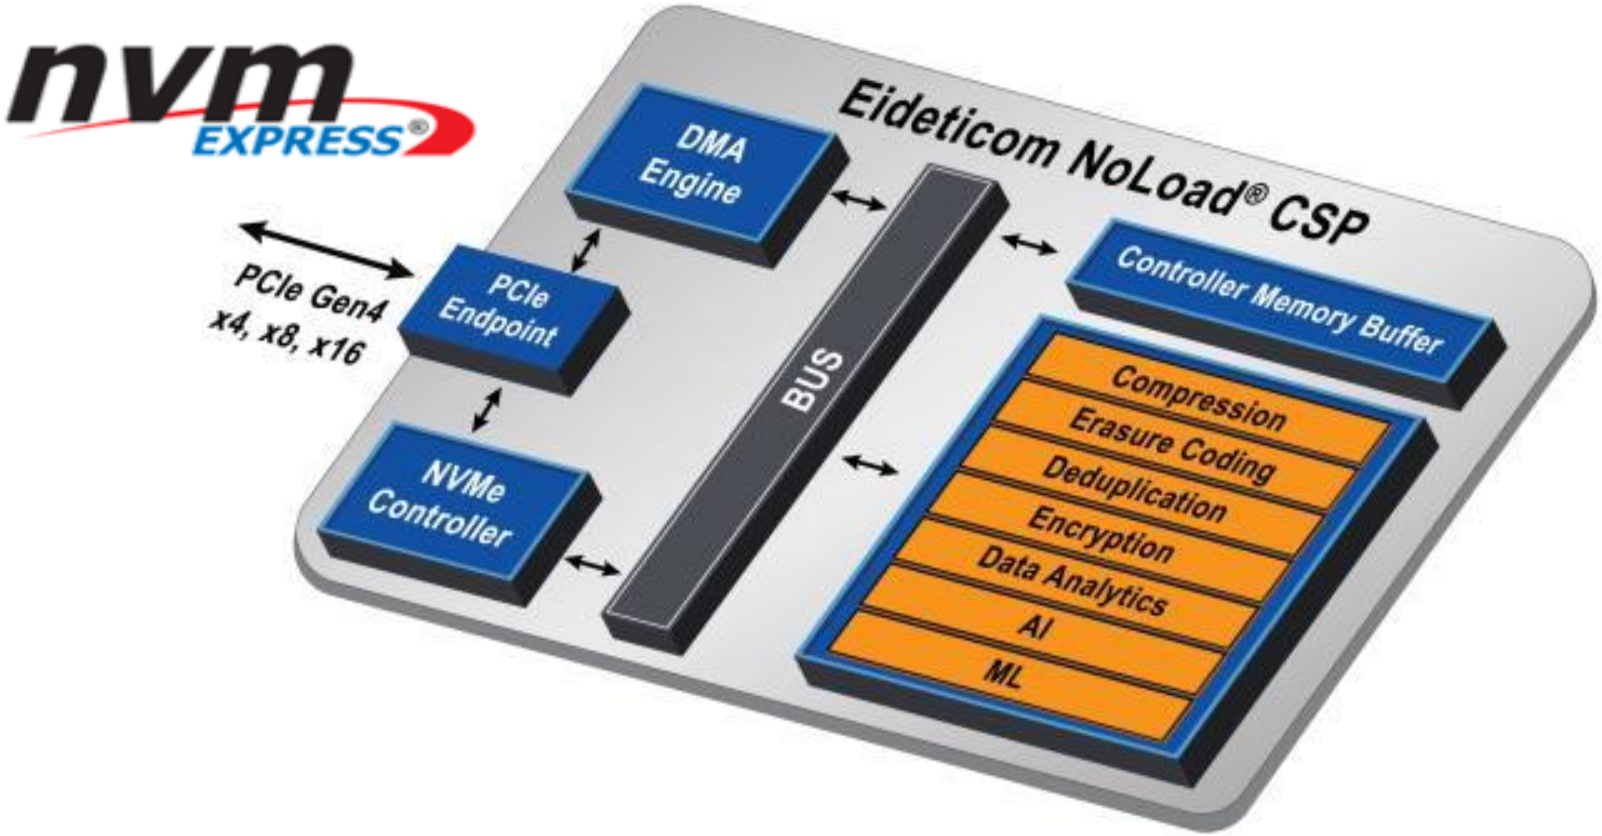


**Best-In-Class Storage and Analytic Acceleration delivered via an NVMe-based Computational Storage Processor.**

# NoLoad® CSP – Hardware Platforms

## NoLoad® CSP U.2

- Standard U.2 NVMe form-factor: Utilizing SFF-8639 connector
- BittWare 250-U2

## NoLoad® CSP Alveo

- Standard GPU form-factor: x16 PCIe
- Deployed on Xilinx Alveo U200, 250 or U280


## NoLoad® CSP E1.S EDSFF

- Standard E1.S NVMe form-factor
- BittWare 250-E1.S Hardware

New!

# NoLoad® CSP - Software

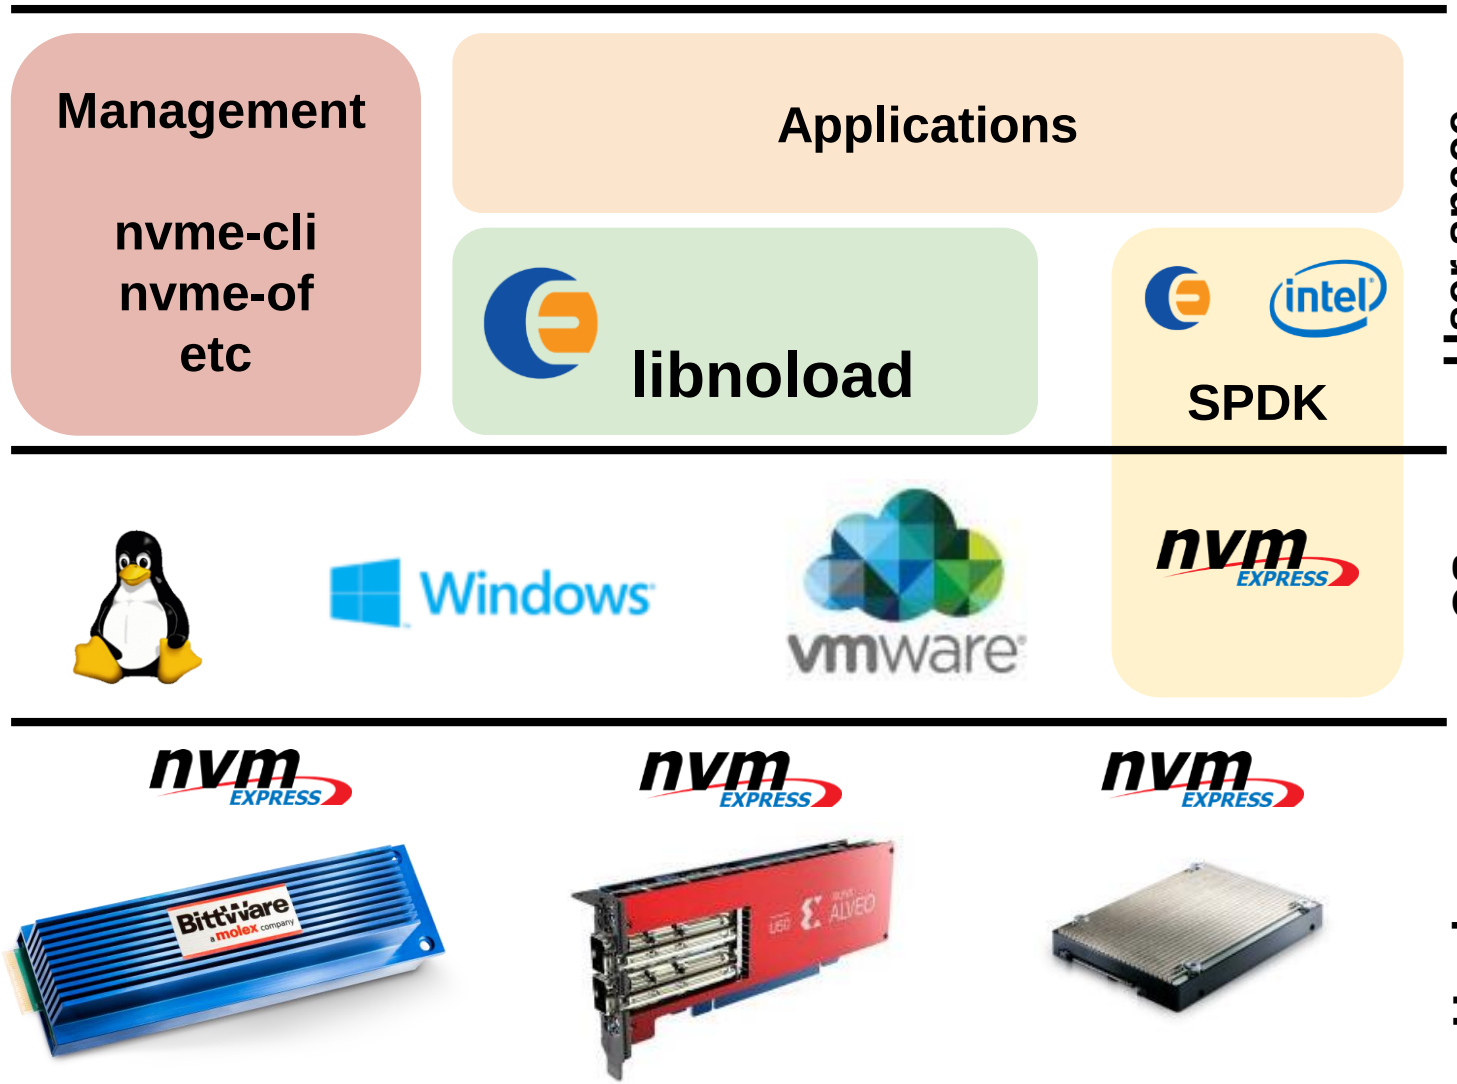


- Both kernel and User space frameworks supported
- No changes to OS
- Use In-box NVMe drivers
- NoLoad® CSP Hardware

# A.C.S
# (After Computational Storage)

# End Solutions – Application Acceleration: RocksDB

EIDETICOM

## Bottom Line

- 6x more transactions per sec
- 2.5x more efficient
- 4x reduced NAND costs
- Improved QoS

## Details

- Eideticom's NoLoad CSP
- Xilinx Alveo U280 (HBM)
- Dell R7425 PowerEdge server
- RocksDB
- Linux Operating System
- 2 NoLoad instances with compression offload

RocksDB

OCP SUMMIT


nvm EXPRESS

# End Solutions – Application Acceleration: Hadoop

## Bottom Line

- Significant improvement in Job Throughput and Cost per Job Reduction
- Lower TCO/TCA
- Improved QoS

## Details

- Eideticom's NoLoad CSP
- Xilinx Alveo
- Dell R7425 PowerEdge server
- Hadoop MapReduce
- Linux Operating System
- NoLoad CSP iwith compression and EC offload


# End Solutions – Application Acceleration: ZFS

## Bottom Line

- 2+ GB/s compression per NoLoad U2
- 26X improvement in CPU loading
- ZFS performance scales linearly with number of NoLoads

## Details

- Eideticom's NoLoad CSP in U.2
- ZFS updated to integrate directly with Eideticom's NoLoad CSP
- Benchmarking on Dell AMD EPYC server
- Supports Burst Buffer architectures by providing a fast storage layer
- Test set with 30% compressible data

**Dell R7425 AMD EPYC Server**

**Scalable number of NoLoad U2s**

**Dell NVMe SSDs**

# End Solutions – NVMe-oF

## Bottom Line

- NoLoad Accelerators located in a remote server can be accessed by any client with a RDMA or TCP/IP connection

## Details

- Disaggregation of NoLoad Accelerators using NVMe-oF
- NoLoad Accelerators identify as NVMe Namespaces, which can be accessed/shared using NVMe-oF
- NoLoad Accelerators (compression, EC) shared across RDMA & TCP/IP
- Broadcom BCM58800 SmartNIC

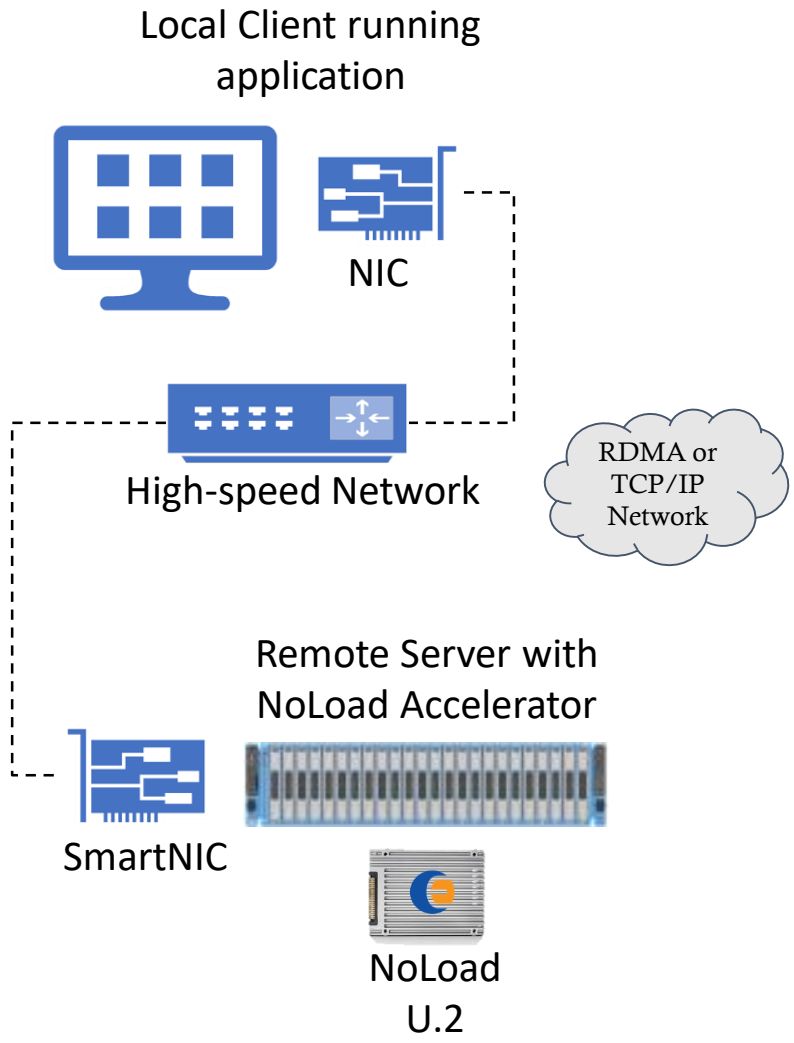

# Conclusions

- The large consumers of accelerators (hyperscalers) want vendor-agnostic, consumable interfaces, software and management stacks. NVMe gives them all that.
- Eideticom's NoLoad CSP is the world's first UNH certified NVMe-based CSP.
- Using NoLoad we can accelerator customers applications and filesystems.

**This is all getting standardized so cusomters can enjoy a scalable, vendor agnostic framework for acceleration and computational storage!**


Eideticom HQ
3553 31st NW,
Calgary, AB,
Canada T2L 2K7

Eideticom (Bay Area)
168 South Park,
San Francisco, CA 94107
USA

www.eideticom.com

Contact: sales@eideticom.com