
Excelero
Excelero NVMesh
Lowest-Latency Distributed Block Storage for IBM Spectrum Scale

Excelero

Founded in 2014

Experienced Executive Team

Well-funded

Strong Patent Portfolio

Web-Scale Infrastructures

technicolor

NASA

CMA

GE

SciNet

teuto.net

Lenovo

SUPERMICRO

Mellanox TECHNOLOGIES

Western Digital

BROADCOM

Micron

intel

SAMSUNG

NVMesh®

Software-defined Block Storage

Lowest Overhead
Local Flash Latency across the Network

100% Software-defined
Use any Hardware

Block Storage
Use any File System

Excelero NVMesh Logical Architecture

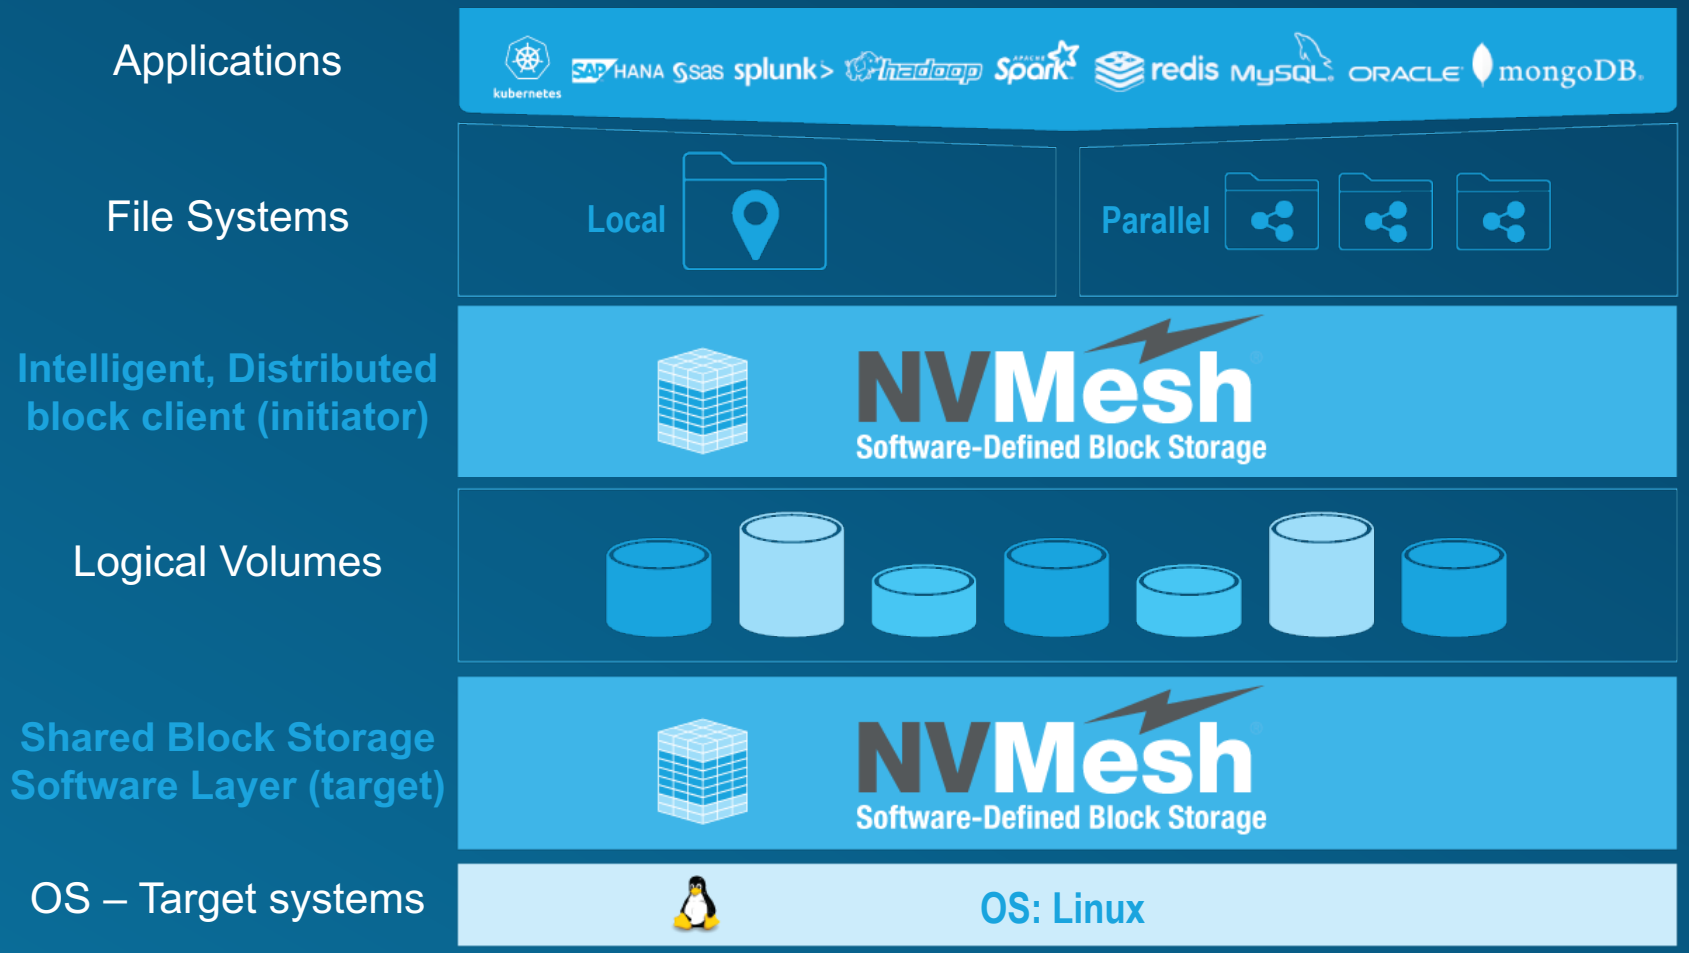

Flexible, Distributed Deployment Models

Converged - Local NVMe drives in Application Servers

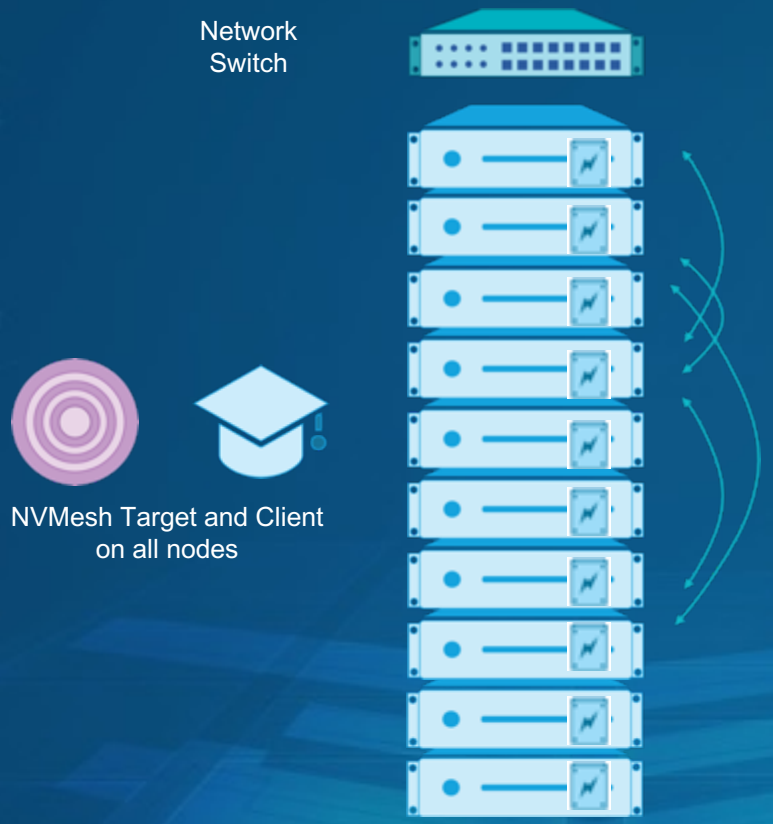

- Single, unified storage pool
- NVMesh client and target running on all nodes
- NVMesh bypasses server CPU
- Various protection levels
- No dedicated storage servers needed
- **Linearly scalable**
- **Highest aggregate bandwidth**

Flexible, Distributed Deployment Models

Top-of-Rack Flash

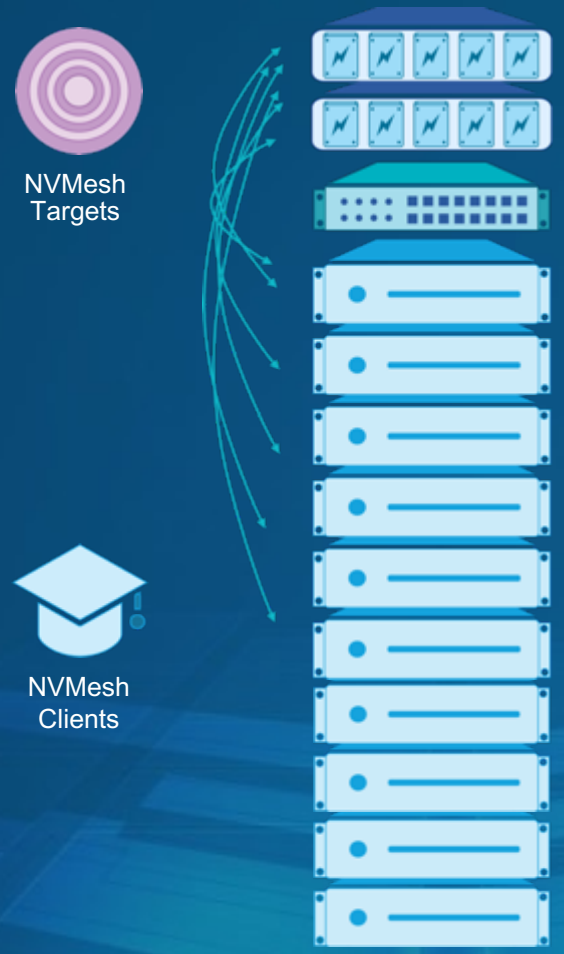

- Single, unified storage pool
- NVMesh Target runs on dedicated storage nodes
- NVMesh Client runs on application servers
- Applications get performance of local NVMe storage
- Various Protection Levels
- **Linearly scalable**

MeshProtect - Accelerated & Simplified Data Protection

No Redundancy

Concatenated

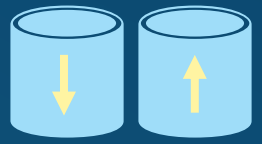

MeshProtect 0
(up to 128 wide)

Mirroring

MeshProtect 1

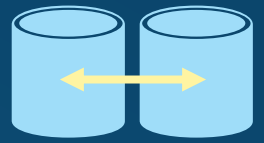

MeshProtect 10
(Striped and mirrored)

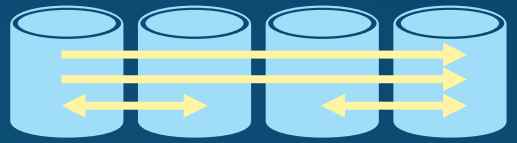

Parity-based

MeshProtect 60 (8+2)

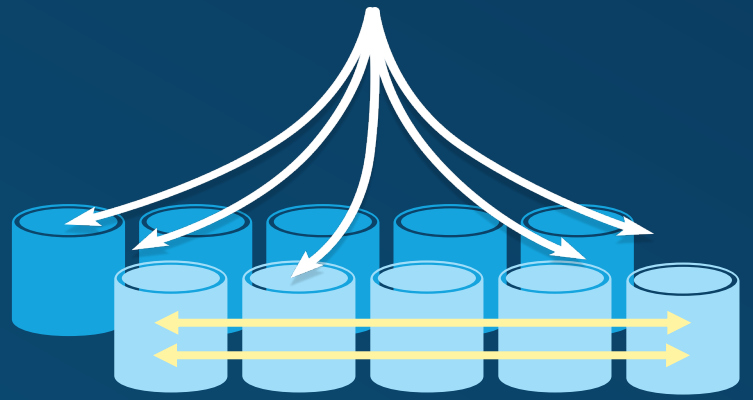

Use-Cases

NVMesh Accelerates Spectrum Scale GPFS

The World's Fastest AI GPU Farms Require an Extremely Capable Storage system

A single NVIDIA DGX-2 is capable of ingesting data at rates of 100GB/s and Millions of IOPS!

It's Hard enough to feed one.

How about at data center scale?

Data Center-Scale AI Computation

NVMesh + Spectrum Scale – Scalable Storage for AI/ML/DL

- Access remote NVMe at local speed
- Distributed and scalable storage infrastructure
- Share storage resources across multiple GPU servers
- Eliminates need to copy data locally
- Datasets can be larger than what can fit inside the DGX
- Zero-CPU usage on both GPU and Storage access

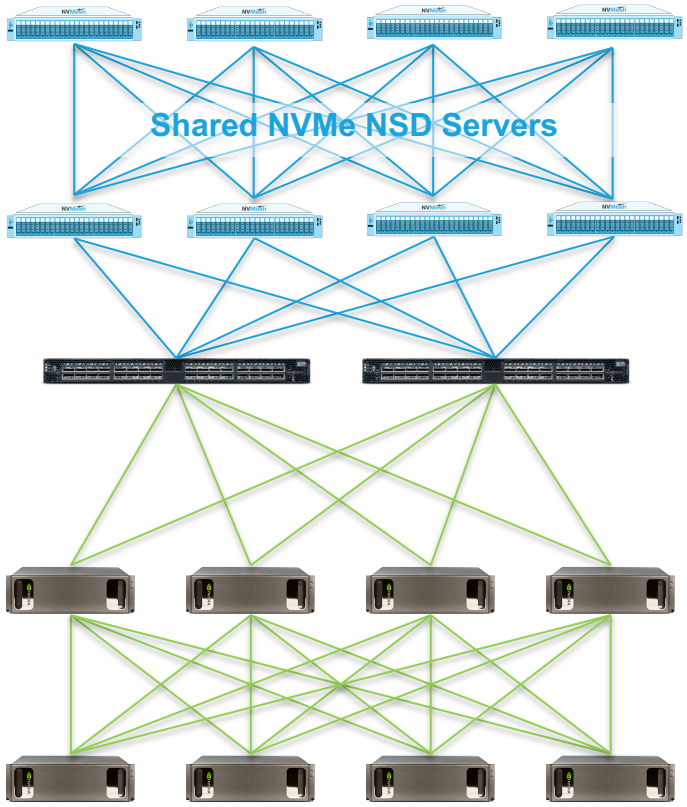

NVMesh®

NVIDIA

Excelero

Machine & Deep Learning data input pipeline optimization

Customer Problem:

- Certain analytic loads require low latency and/or high IOPs.
  Unrestricted pipelining and storage IO, in general, is required to reduce the execution time of learning tasks.

NVMesh Solution Benefits:

- Allows ML/DL workflows to extract and acquire the data required for the learning task at the very lowest latency.
- **Enables highly efficient pipelining**
- Eliminates CPU, GPU, and TPU data starvation
- Allows to invest into compute resources rather than storage appliances
- Truly SDS – HW platform independency

Without pipelining, the CPU and the GPU/TPU sit idle much of the time:

With NVMesh accelerated pipelining, idle time diminishes significantly:

Boston Talyn Server + NVMesh + Spectrum Scale - RESNET-50 Results

Tensorflow DGX-1 Talyn Benchmarks

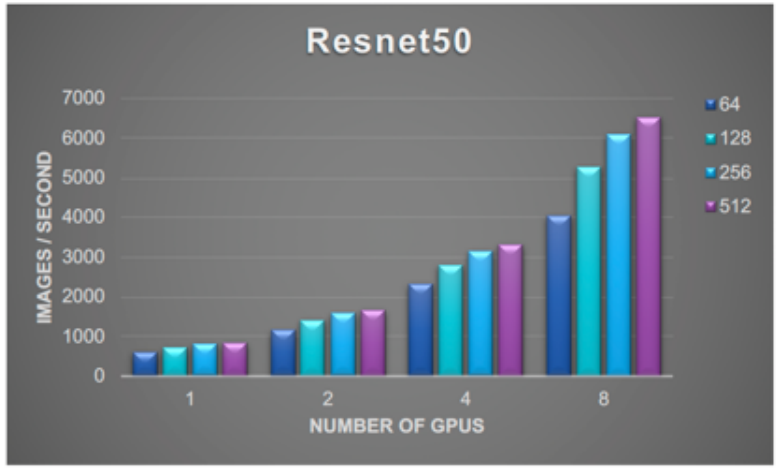

	Resnet50			
	1	2	4	8
64	603	1147	2316	4046
128	730	1408	2808	5281
256	807	1575	3149	6097
512	845	1669	3321	6530

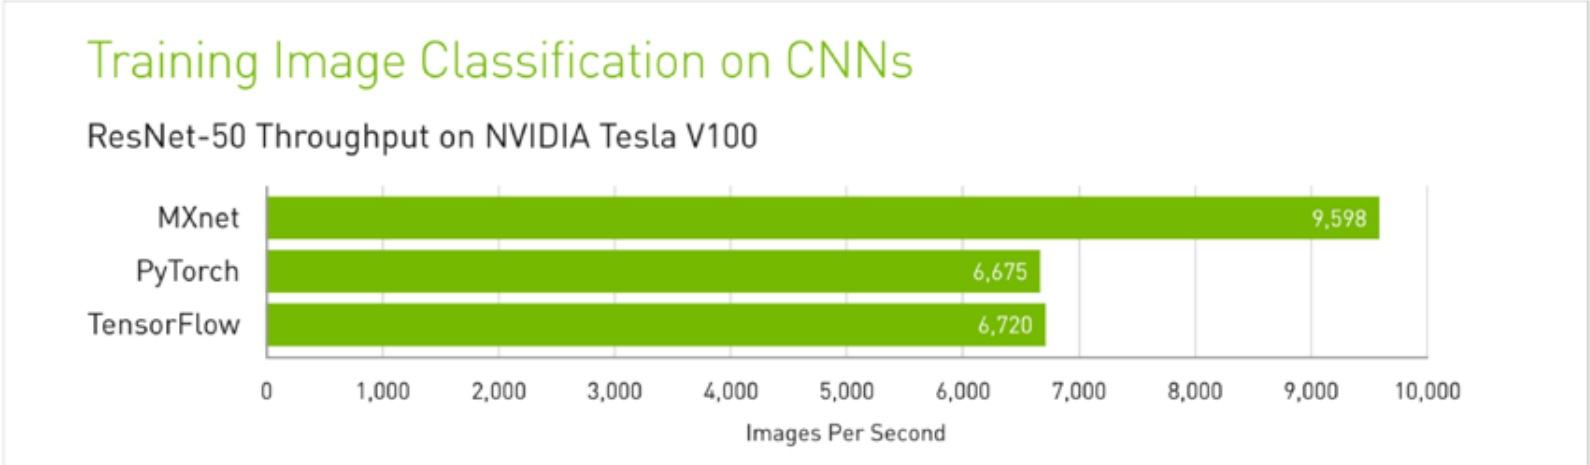

https://developer.nvidia.com/deep-learning-performance

NVMesh for Distributed LROC and HAWC

Local Read-Only Cache

- No local flash required
- Flexibility & Granularity
- Lowest Read IO latency
- Highest Read IOPs

Highly-Available Write Cache

- Cost efficient highly available flash
- Flexibility & Granularity
- Lowest Write IO latency
- Higher Write IOPs

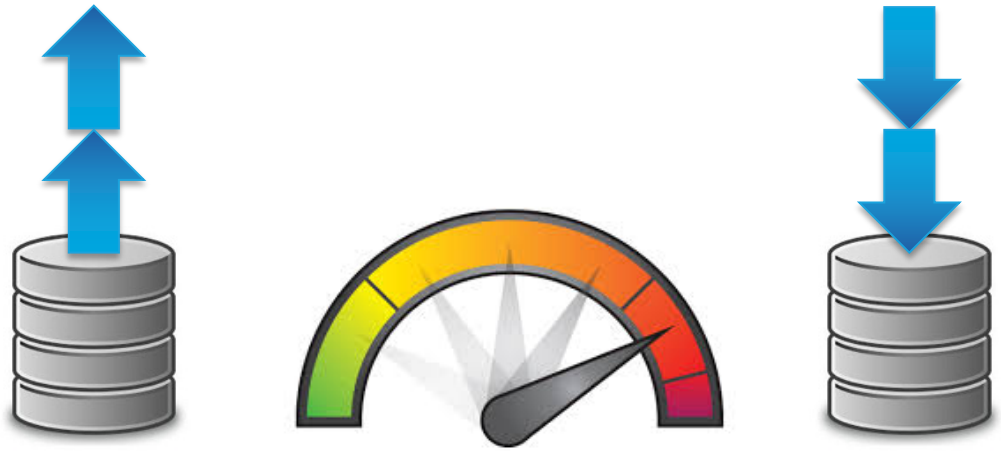

Extremely High Performance Burst-Buffer Cache

Buffer Burst Cache

- Reliable and predictable performance
- Highest sustained write throughput
- Lowest sustained write latency
- Leverage the Buffer Burst Cache resource ***when*** and ***where*** you need it
- Ensure job SLA's and accelerate checkpoints

Customer Success

80 pooled NVMe drives delivers **230GB/s** of throughput and **20M** random 4k IOPS

SciNet | compute canada | calcul canada

Pooling NVMe Within GPFS NSDs enables new Science use-cases

Use Case

- Large-scale modeling, simulation, analysis and visualization
- Visualizes supercomputer simulation data on 100s of compute nodes

Problem

- Finish check pointing faster and start running the job
- Achieve performance of 230GB/Sec at the lowest price achievable

Lenovo | Mellanox TECHNOLOGIES | Excelero

Solution

NVMesh by Excelero enables SciNet to create a petabyte-scale unified pool of high-performance distributed flash retaining the speeds and latencies of directly-attached media

Questions?


Excelero
Thank you!