

OpenStack Virtual Machines Without I/O Virtualization Overhead

Tackle your most demanding shared storage performance requirements with NVMesh Reference Architecture utilizing Supermicro NVMe optimized server solutions and Mellanox SR-IOV and RDMA technologies.



What do you get when you take the lowest latency, highest performance software defined block storage and pair it with a hardware platform designed from the ground up to serve NVMe flash?

The Answer - NVMesh Reference Platform

The Dilemma: Bare-metal Performance or Virtualization?

When deploying virtual machines in OpenStack environments customers are forced to choose between high performance storage or application migration ability. High performance requires you to run your virtual machines on local flash devices such as NVMe SSD. While delivering local flash performance, it traps VMs on a physical platform. In order to get the full benefits and freedoms of virtualization, customers may choose to utilize all-flash-arrays (likely via iSCSI) as remote shared storage, but this introduces high latency and dramatically reduces performance.

Benefits for OpenStack:

- Eliminate virtualization overhead
- Enable high performance, transactional applications
- Cinder integration
- Highest Performance Block Storage
 - Lowest Latency (10's μ s)
 - High Throughput (10's GB/s)
- Massive Scalability
 - 10's Millions of IOPs
 - 100's of TB of NVMe flash

The Solution: Excelero NVMesh

Excelero, Supermicro and Mellanox have partnered to deliver the world's fastest software defined storage solution for OpenStack environments. It delivers local flash performance with data protection while eliminating virtualization overhead. Excelero NVMesh software delivers direct-attached storage speeds and latencies yet is networked and distributed. Combined with the Supermicro 2U, 24x 2.5" NVMe SSDs storage platform and Mellanox SR-IOV and RDMA technologies, this complete solution delivers disaggregated, high speed, low-latency flash storage. The shared storage solution can be read/written by OpenStack VMs over an Ethernet network in the 10's of micro-seconds. This is 20-30 times faster than leading all-flash-arrays. It's as fast as local, in-server flash, allowing the freedom and flexibility of VMs without the traditional virtualization overhead. Unlike general purpose storage solutions, NVMesh is ideal for highly transactional databases, real time analytics, multi-terabyte indexes and other low latency demanding applications.



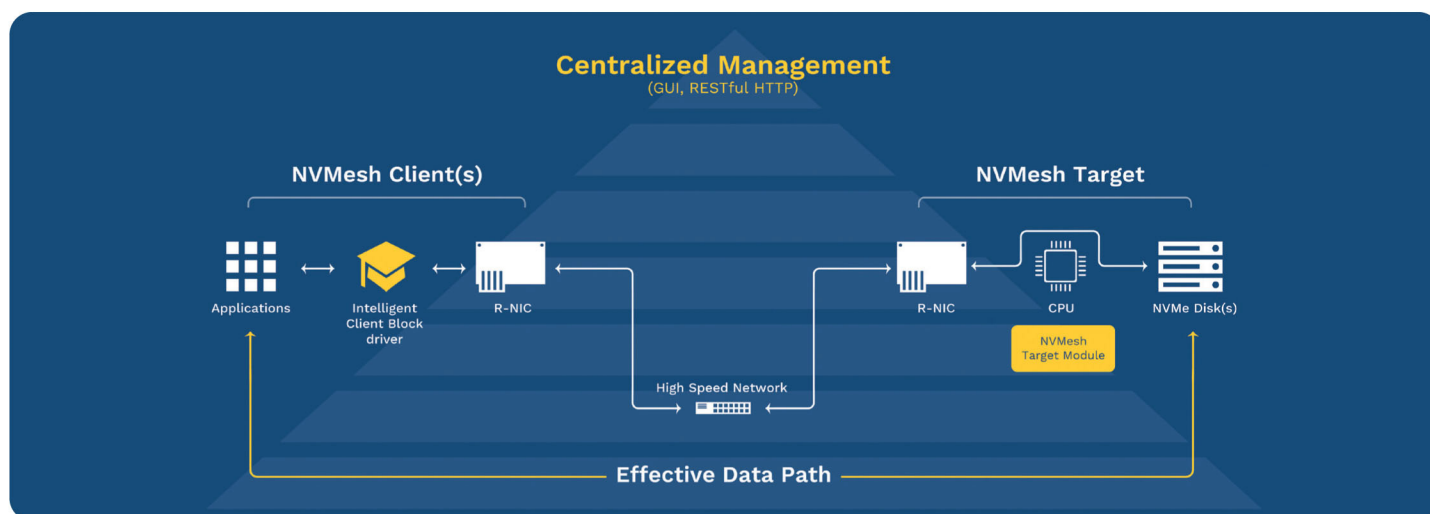
Supermicro 2028U-TN24R4T+

NVMe allows remote NVMe devices to perform at local speeds and latencies in OpenStack virtual environments in under 100 microseconds

Software Architecture and Components

NVMe is comprised of three software modules:

- The Storage Management Module functions as a centralized web based GUI, RESTful API and controls overall system configuration. It also interfaces with Docker Persistent Volumes plugins and OpenStack Cinder drivers.
- The Target Module identifies and reports NVMe target drives to the Management Module. It validates initial connections to NVMe drives (from clients) but is not in the data path. It is installed on the reference platform.
- The Client Block Driver (a Cinder plugin in OpenStack environments) runs in the virtual machine image and delivers local performance from remote NVMe devices and is installed on any client image making use of NVMe logical block volumes.



Specifications

Storage Software Features	
Volumes and Data Protection	Striping (RAID 0), Mirroring (RAID 1/10 – requires second server) Volumes may be striped over multiple drives and/or servers
Availability	Mirrored volumes, Multi-path (Active/Active)
Management and Provisioning	Web GUI, RESTful API, OpenStack Cinder Driver, Docker Persistent Volume Plugin
Maximum Clients/Targets in one NVMe	128
Performance per Reference Platform	
Random Read/Write 4K IOPs	Up to 4 Million/2.8 Million
Read/Write Throughput	Over 16GB/s
Read/Write Latency from Client to Server	As low as 90µs/25µs
Reference Platform Features	
Server	2U 24 NVMe SYS-2028U-TN24R4T+
Storage Network Ports	Up to 4 x 100GbE ports (2 NICs, PCIe Gen3 x16)
Availability/Redundancy	1+1 1600W Titanium Level Power Supplies, 4 Hot Swap Fans
Remote Accessibility	IPMI 2.0 with virtual media over LAN and KVM-over-LAN support
Form Factor	2U Rackmount, 17.2"W, 3.5"H, 27.76"D, RoHS Compliant

Excelero, the Excelero logo and NVMe, are trademarks of Excelero, Inc. in the United States and other countries. All other trademarks used herein are the property of their respective owners. © Copyright 2016 Excelero, Inc. All rights reserved.