Excelero

Leveraging NVMe SSDs at Scale for NASTRAN Analytics

USE CASE

INTRODUCTION

Nastran is an application for multidisciplinary structural analysis that was initially developed to help design more efficient space vehicles such as the Space Shuttle. NASTRAN – NASA STRucture ANalysis – was released to the public in 1971 by NASA. Several commercial releases have been issued since. Today, Nastran is used to analyze the behavior of elastic structures when designing automotive parts, aircraft, railroad tracks and bridges, power plants, skyscrapers, etc. Nastran workloads are very compute and data IO intensive, a perfect use case for NVMe SSDs with Excelero NVMesh®.

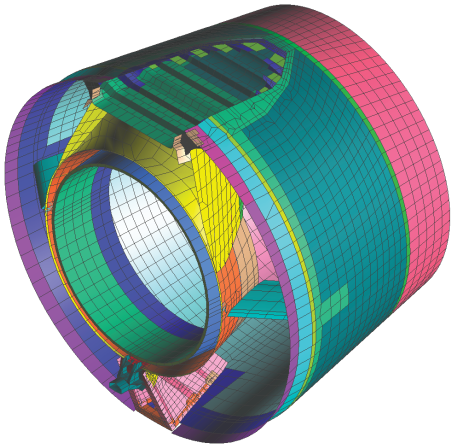

Figure 1[2]

NASTRAN ANALYTICS

Conceptualized by NASA, developed by CSC and now commercialized by several software providers, such as MSC Software, Siemens and NEi Software, Nastran is "an application used by engineers to perform static, dynamic, and thermal analysis across the linear and nonlinear domains, complemented with automated structural optimization and award winning embedded fatigue analysis technologies, all enabled by high performance computing."[1]

Nastran Analysis models tend to be very data-intensive, requiring lots of compute cycles and IO operations. To reduce the processing time and get to results faster, customers are building large, high-performance computing clusters, leveraging parallelization techniques, GPU technologies and NVMe SSDs. As Nastran jobs require processing many small inputs, the high IOPs/throughput and low-latency characteristics of NVMe are ideal. However, the inability to share NVMe across the network, to the entire cluster, is an important limitation to overcome. Deploying NVMe SSDs cluster-wide with Excelero's NVMesh enables customers to leverage the full performance capabilities of their NVMe devices across the entire Nastran cluster.

[1]Source: http://www.mscsoftware.com
[2]Source: http://www.mscsoftware.com/product/msc-nastran

NVMesh FOR NASTRAN

Excelero delivers low-latency (25µs) distributed block storage for high-performance data analytics workloads. Excelero's NVMesh enables shared NVMe across any network and supports any parallel file system. Distributed workloads can leverage the full performance of NVMe SSDs with the convenience of centralized storage while avoiding proprietary hardware lock-in and reducing the overall storage TCO.

NVMesh is a Software-defined Block Storage solution that features Elastic NVMe, a distributed block layer that allows unmodified applications to utilize pooled NVMe storage devices across a network at local speeds and latencies. Distributed NVMe storage resources are pooled with the ability to create arbitrary, dynamic block volumes that can be utilized by any host running the NVMesh block client. These virtual volumes can be striped, mirrored or both while enjoying centralized management, monitoring and administration.

Particularly interesting for large clusters, is the fact that even cluster node servers without local NVMe SSDs have access the NVMe pool at full speed. This can result in a significant savings. NVMesh is the only solution that lets local NVMe storage resources to be pooled and then logically distributed over a compute cluster – allowing maximum utilization of capacity and storage bandwidth. There's no modification required to the HPC applications.

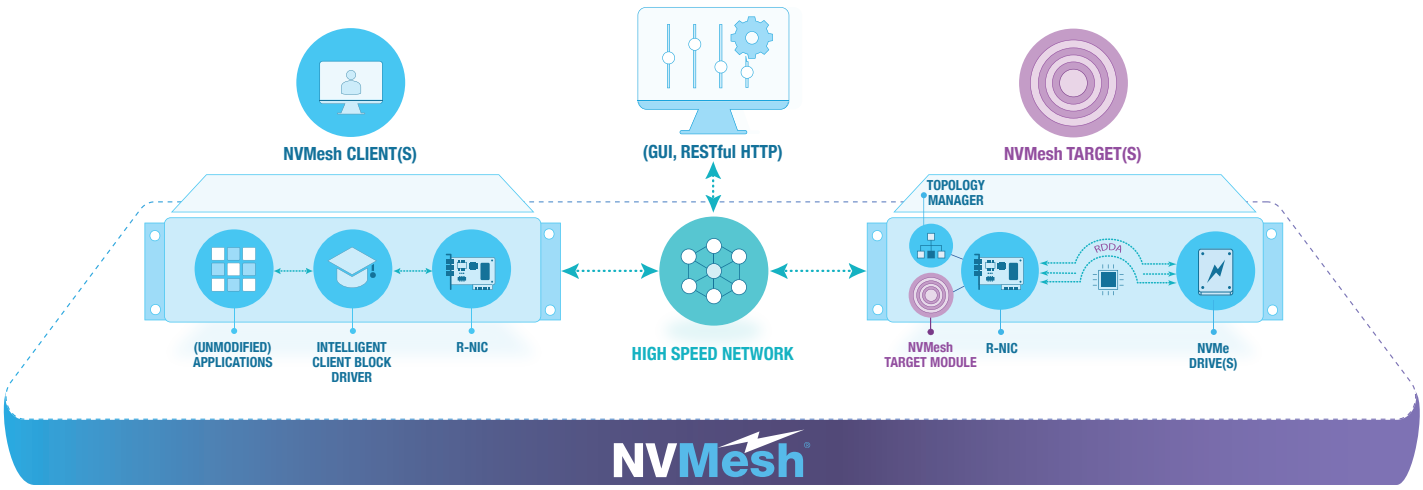

RUNNING MORE JOBS ON THE SAME CLUSTER FASTER

A performance benchmark of MSC Nastran Lenovo servers with pooled and shared NVMe (NVMesh) demonstrated significant performance improvements against an environment where the flash resources were being managed by the parallel file system. The NVMesh environment also showed drastic resource utilization improvements. The benchmark was based on a common industry test workload using an MSC Nastran SOL 400, which was designed to solve complex problems in materials science, specifically for determining how stiff structures, such as skyscrapers, will deform under a heavy load. The benchmark test used the SOL 400 applications with a standard, publicly accessible dataset.

A single SOL 400 job running on the cluster took 20% less time to complete using NVMe pooled with NVMesh. While that alone would be an impressive result, the NVMesh environment delivered the same performance benefits even when running multiple jobs in parallel, without requiring any additional flash resources. Just like FC SAN with a shared storage controller replaced local HDDs some decades ago, Excelero NVMesh now replaces local SSDs with a shared pool of NVMe storage.

MSC NASTRAN SOL 400 BENCHMARK

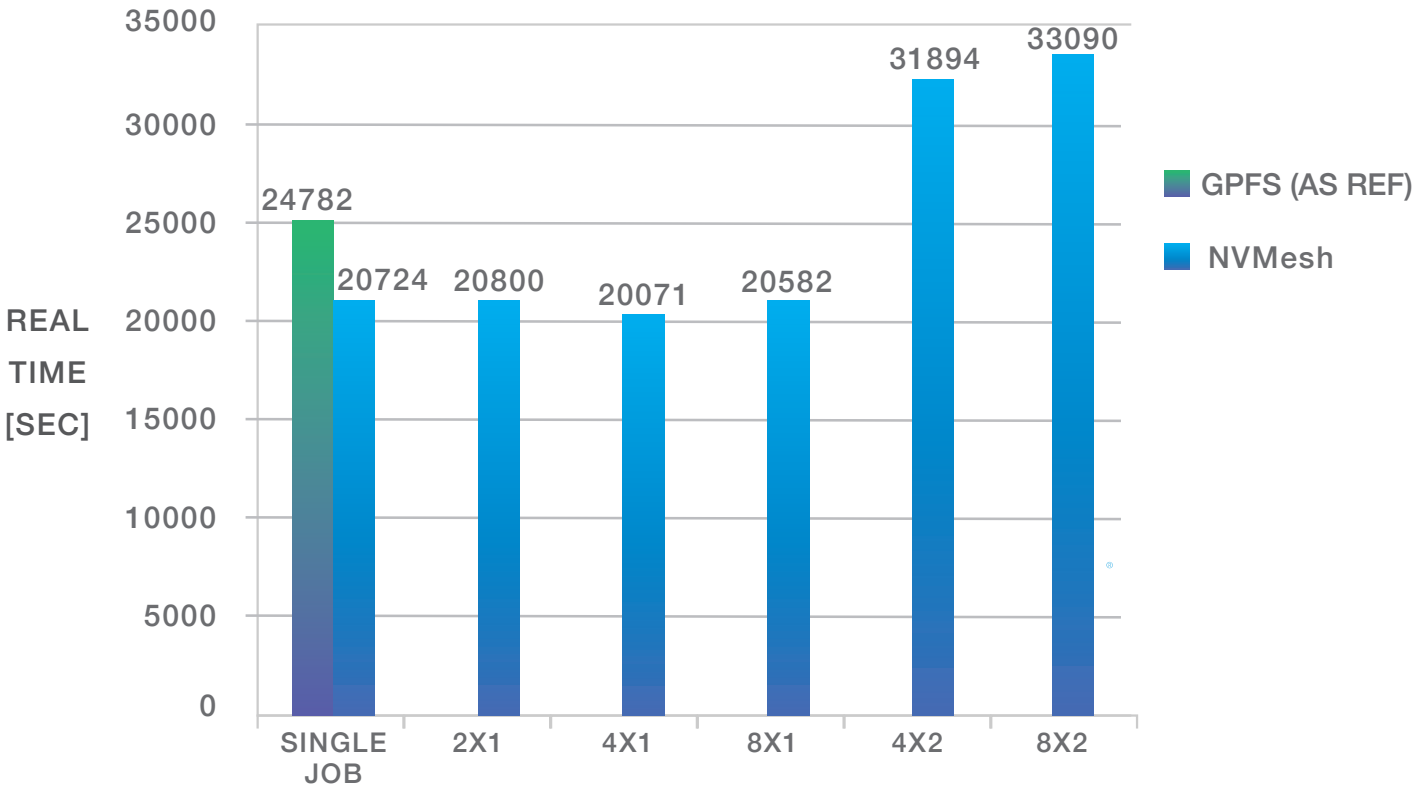

APPLICATION: MSC NASTRAN 2016

SMP=4, MEM=24GB
Max disk usage: 62GB
I/O transferred: 8.6TB
1 or 2 jobs per node

NVMesh

3 Lenovo file servers
2x 800GB NVMe / server
RAID0 across 6 NVMe devices
Up to 10GBps write (single node)
Up to 8.5GBps write (8 clients)
4KB blocks (32 blocks/stripe)

NVMesh BENEFITS

MAXIMIZE RESOURCE UTILIZATION

NVMesh was designed to enable customers to maximize NVMe utilization (capacity, performance & endurance) across their infrastructure. Leveraging the new parity-based data protection, NVMesh delivers over 80% greater storage efficiency, which helps further drive down the NVMe TCO.

STANDARD HARDWARE & PROTOCOLS

NVMesh was designed from the ground up to support any hardware. With the added support for traditional network fabrics and protocols, enabling NVMe over Ethernet, Fibre Channel and Infini-band, customers don't have to invest new networking technologies to deploy NVMesh.

MORE DIAGNOSTIC CAPABILITIES

NVMesh enables users to analyze cluster-wide and per-object performance and utilization, and build a customized dashboard from a selection of data visualization widgets.

NVMesh SPECS

DATA MANAGEMENT & PROTECTION

Multiple Transports	NVMeoF, Patented RDDA
MeshProtect	Virtual volumes in flexible choice of redundancy - Concatenated, RAID 0, RAID 1, RAID 10, Parity-based N(3-11)+M(1-2)
Multiple Drive Types	NVMe, NVMf, SATA, SAS, 3D-XPoint
Failure Domains	Customizable Host, Rack & Row aware

MANAGEMENT & MONITORING

Interactive Interfaces	Web GUI & CLI commands
Automated Provisioning	RESTful API, OpenStack Cinder and CSI

NEXT-GEN DATA CENTER

Flexible Topologies	Physically converged, disaggregated or mixed
Scalability	1000's of nodes up to 3TB/s and 640 Million IOPs, over 65 PB in one cluster
High Performance	Only +5µs additional latency, 100% linear performance scalability
Scale-Out Architecture	Logical volumes can span across drives and hosts for virtually unlimited single volume scalability
Connectivity	Ethernet (TCP/IP and RoCEv2), InfiniBand