

POOLING NVMe WITHIN GPFS NSDs ENABLES EFFICIENT BURST BUFFER

CASE STUDY

SciNet



Lenovo

USE CASE

Large-scale modeling, simulation, analysis and visualization

CHALLENGE

Complete checkpoints within 15 minutes to meet availability SLA's

SOLUTION

NVMesh enables a petabyte-scale unified pool of distributed high-performance NVMe flash as burst buffer for checkpointing

RESULTS

- 80 pooled NVMe devices
- 148 GB/s of write burst (device limited)
- 230GB/s read throughput (network limited)
- Well over 20M random 4k IOPs

BENEFITS

- Meets the 15 minutes checkpoint window
- Extremely cost effective
- Unheard of burst buffer bandwidth

SciNet is Canada's largest supercomputer center, providing Canadian researchers with computational resources and expertise necessary to perform their research at massive scale. The center helps power work from the biomedical sciences and aerospace engineering to astrophysics and climate science. SciNet is currently building what will be the fastest supercomputer in Canada. As the project is funded by government branches, such as the University of Toronto and Compute Canada, this new supercomputer needs to meet high levels of availability to ensure high ROI for the supercomputer. One way to increase availability is by using a burst buffer for checkpointing. This case study lays out how Excelero's NVMesh® enables SciNet to build a petabyte-scale unified pool of distributed high-performance NVMe as a burst buffer for checkpointing. The NVMe pool delivers 230GB/s of throughput and well over 20M random 4k IOPs and enables SciNet to meet its availability SLA's.

High-performance computing applications consist of complex sets of processes that sometimes run for weeks. When any of these processes is interrupted, this could destroy the results of the entire compute job. This problem becomes worse as supercomputers become more powerful – imagine the challenge for Canada's soon to be largest supercomputer. Therefore, parallel computing applications use the concept of checkpoint-restart. This technique allows compute jobs to be restarted from the most recently saved checkpoint in case of an interruption.

Checkpoints are typically saved in a shared, parallel file system; SciNet has chosen GPFS. But as clusters become larger and the amount of memory per node increases, each individual checkpoint becomes larger and either takes more time to complete or requires a higher-performance file system. When a system is checkpointing it's not computing, which reduces the availability score of the system. To shorten those moments of unavailability, SciNet decided to implement a burst buffer leveraging Excelero's NVMesh.

WHAT IS A BURST BUFFER?

A burst buffer is a fast and intermediate storage layer between the non-persistent memory of the compute nodes and persistent storage – the parallel file system. This layer is configured to take a burst of write IO at a very high rate. Once the burst (checkpoint) is complete, the written data is “drained” to the parallel file system, using the GPFS policy engine. This allows checkpoints to finish rapidly so that systems meet availability SLA’s.

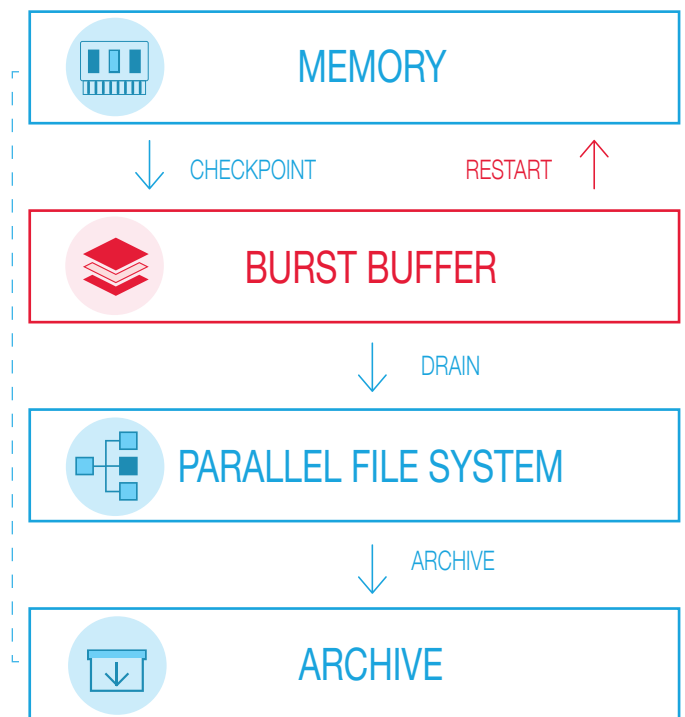
When flash storage is used as the burst buffer pool it has the added advantage of facilitating a faster restart (when needed) as checkpoint restarts often impose a very large random read load on the underlying storage. With those benefits in mind, the SciNet burst buffer was sized to hold

two checkpoints so the most recent completed checkpoint would be available for restart. To maximize performance, and reduce the checkpoint window, SciNet decided to leverage higher performance NVMe SSDs.

THE NVMesh® BURST BUFFER

Excelero enables customers to build high-performance burst buffers without needing additional proprietary arrays. NVMesh customers can use standard NVMe drives in their application servers to build a local burst buffer or build converged file system / block server appliances. The main advantage of the latter method is that it adds redundancy with centralized management while at the same time preserving all compute resources for the applications themselves.

HPC STORAGE ARCHITECTURE WITH BURST BUFFER



NVMesh and its patented Remote Direct Drive Access (RDDA) technology allows customers to logically disaggregate NVMe drives in the compute nodes away from CPU resources. This way, local NVMe drives may be used by remote compute nodes without consuming local CPU.

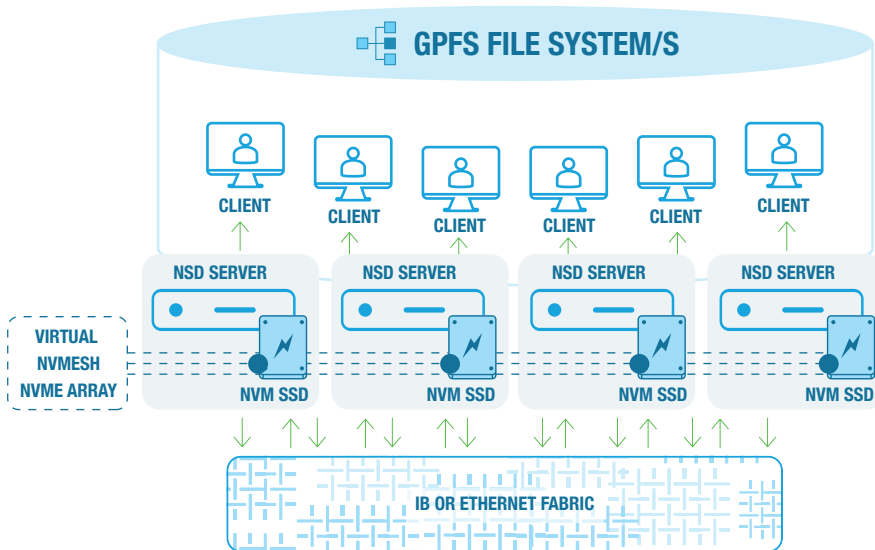
As a result, every compute node can have a local NVMe SSD (or multiple drives) and all the drives are pooled for use by the cluster. In the most simplistic form, half of each drive can be used as a local burst buffer while the other half is reserved for the redundant copy of

a peer. Thus, when a node fails, its scratch is preserved and accessible by an alternate node - any node on the fabric.

SCINET'S NVMeSH BURST BUFFER IMPLEMENTATION

Excelero's NVMeSH enables SciNet to create a petabyte-scale unified pool of distributed high-performance NVMe flash retaining the speeds and latencies of directly-attached media. The NVMe pool, consisting of 80 NVMe devices in just 10 NSD servers gives about 148 GB/s of write burst (device limited) and 230GB/s (network limited) of read throughput and well over 20M random 4k IOPs. This configuration is more than sufficient to meet the 15 minutes checkpoint window that is required to meet availability SLA's that were defined for the new super-computer.

CONVERGED NSD SERVER TOPOLOGY



For SciNet, NVMeSH is an extremely cost effective method to achieve unheard of burst buffer bandwidth by adding commodity flash drives and NVMeSH software to compute nodes and low latency network fabric that was already provided for the supercomputer itself. NVMeSH provides redundancy without impacting target CPUs: this enabled standard servers not to act just as block targets but also as file servers. Integration with SciNet's parallel file system was very straightforward as it looks like a simple block device.

NVMeSH BENEFITS FOR BURST BUFFER

- Petabyte-scale unified pool of high-performance flash retaining the speeds and latencies of directly-attached media.
- Supports large-scale modeling, simulation, analysis and visualization.
- Visualizes supercomputer simulation data on 100s of compute nodes.
- Finish checkpointing faster and start running the job.
- Achieve highest performance at the lowest price.
- Leverage the full performance of your NVMe SSDs at scale, over the network.
- Scale your performance and capacity linearly.
- Easy to manage & monitor, reduces the maintenance TCO.
- Utilize hardware from any server, storage and network vendor. No vendor lock-in.

ABOUT SCINET

SciNet is Canada's largest supercomputer centre, providing Canadian researchers with computational resources and expertise necessary to perform their research on scales not previously possible in Canada. SciNet powers work from the biomedical sciences and aerospace engineering to astrophysics and climate science. SciNet is part of Compute Canada, a national infrastructure for supercomputing-powered innovation, and is funded by CFI, NSERC, the Ontario Government, Fed Dev Ontario, and the University of Toronto.

ABOUT EXCELERO

Excelero delivers low-latency distributed block storage for HPC workloads. Founded in 2014 by a team of storage veterans and inspired by the Tech Giants' shared-nothing architectures for web-scale applications, the company has designed a software-defined block storage solution that meets the low-latency performance and scalability requirements of the most demanding AI and machine learning applications.

ABOUT NVmesh

Excelero's NVMesh enables shared NVMe across any network and supports any local or distributed file system. Customers benefit from the performance of local flash with the convenience of centralized storage while avoiding proprietary hardware lock-in and reducing the overall storage TCO. The solution has been deployed for hyper-scale Industrial IoT services, machine learning applications and massive-scale simulation visualization.

