Excelero

NVMesh®

SCALABLE AND PERSISTENT STORAGE FOR CONTAINERS

DATA SHEET

INTRODUCTION

The use of containers is rising sharply. After early adoption for cloud applications and later for enterprise production environments, we are now seeing massive interest in containers for AI workloads. As such, the container ecosystem is booming, with plenty of solutions for container management, orchestration, monitoring and security.

The top pain point for container environments has been scalable and persistent storage. Traditional storage solutions were not designed for container application design and introduce performance, scalability and flexibility constraints. Container environments requiring the performance of local flash but the flexibility and data protection of centralized storage have been left without a viable solution. Instead of using proprietary, hardware-based storage solutions, infrastructure architects and developers who want to embrace microservices style deployments are seeking to leverage intelligent software and standard servers with state-of-the-art flash.

FAST-GROWING CONTAINER MARKET

The promise of container technologies is application portability, scalable application design and efficient hardware utilization. Per MarketsandMarkets, the global application container market is expected to grow from $1.2 billion in 2018 to $4.98 billion by 2023, at a CAGR of 32.9% during the forecast period. Allied Market Research predicts even faster growth, towards $8 billion by 2025. This growth is confirmed by the breadth and diversity of emerging offerings based on or supporting containers: orchestration platforms, monitoring and management tools, and security layers. In spite of this fast-growing ecosystem, many customers are still holding back on deploying container technology at massive scale due to the lack of tools for managing persistent container storage. Most companies have started moving the stateless parts of their applications to containers first and are now looking for solutions for stateful/persistent services as they increase their container adoption.

Customers are not only concerned about storage management, but also about data loss, which calls for persistent storage that effectively scales with the number of containers. Robust, low-latency container storage is long overdue as more IT teams look to microservices requiring persistence to improve the speed and reliability of cloud environments. Therefore, Excelero integrates with leading-edge innovations including Kubernetes pods so that our customers can build IT architectures at data center scale backed by more efficient Flash storage.

NVMesh® PERSISTENT CONTAINER STORAGE

NVMesh Elastic NVMe supports persistent, low latency container storage for scale-out architectures utilizing Kubernetes. This unique offering makes use of pooled, redundant NVMe storage for container applications requiring persistent volumes, so enterprises can obtain both local flash performance and container mobility at data center scale.

Container users need persistent storage that is scalable enough for stateful applications yet also offers mobility to help protect against drive or host failure. By leveraging Kubernetes with NVMesh, Excelero is enabling IT teams to have exactly that: containers that have high performance storage with both persistence and mobility.

Usage of NVMesh-backed containers ranges from persistent local storage specific to a single container to providing a SAN-like backend for a container based deployment of a distributed file system. Containers can consume a private volume using a local file system such as XFS or EXT4 or even directly as block storage. Using the PVC ReadWriteOnce access mode, NVMesh can ensure that only one container is accessing the storage to avoid corruption with Kubernetes driven failover scenarios. Alternatively, large data sets can be shared as a form of reference data for a large number of containers using the PVC ReadOnlyMany access mode. To deploy a distributed file system using containers, utilize NVMesh-backed volumes, often in a pseudo SAN-mode with the PVC ReadWriteMany access mode.

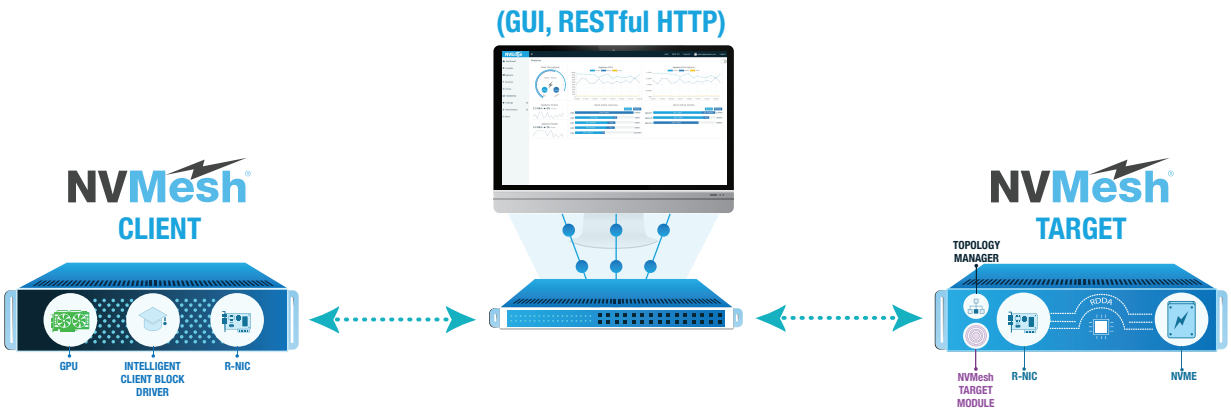

NVMesh is a Software-Defined Block Storage solution that features Elastic NVMe, a distributed block layer that allows applications to utilize pooled NVMe storage devices across a network at local speeds and latencies. Distributed NVMe storage resources are pooled with the ability to create arbitrary, dynamic block volumes that can be utilized by any host running the NVMesh block client. These virtual volumes can optionally be striped and mirrored or erasure coding can be employed for data protection. NVMesh provides centralized provisioning, monitoring and administration. In short, applications and containers can enjoy the latency, throughput and IOPS of a local NVMe device while at the same time getting the benefits of centralized, redundant storage. NVMesh provides the ability to attach volumes ubiquitously, enabling containers to access storage on any server at any time. NVMesh is deployed as a virtual, distributed non-volatile array and supports both converged and disaggregated architectures, giving customers full freedom in their architectural design.

NVMesh BENEFITS for MicroServices

- Persistent storage that scales for stateful applications
- Leverage the full performance of your NVMe flash at any scale, over the network
- Predictable application performance – ensure that storage is not a bottleneck
- Scale your performance and capacity linearly
- Containers in a pod can access persistent storage presented to that pod, but with the freedom to restart the pod on an alternate physical node
- Mobility to help protect against drive or host failure
- Choice of architecture: converged, disaggregated or mixed
- Choice of Kubernetes PVC access mode to match the storage to the application and file system requirements

The solution was designed for infrastructure architects and developers who want to embrace microservices style deployments but desire more performant approaches than NFS or persistent volumes on traditional all-flash-arrays (AFAs). NVMesh provides containers and applications the performance of local flash with the flexibility and data protection of centralized storage. With NVMesh, the Kubernetes' advanced orchestration layer can deliver pooled NVMe with local latency and performance. In this approach, containers in a pod can access persistent storage presented to that pod, but with the freedom to restart the pod on an alternate physical node.

Excelero's NVMesh 100% Software-defined Storage platform further benefits container deployments with its approach by shifting data services from centralized CPU to complete client-side distribution. It virtualizes the NVMe devices and unifies the capacity into a single pool of high-performance storage in an approach that makes data locality irrelevant; a breakthrough in enabling local latency and speeds on the network using standard hardware. Because NVMesh does not impose a "CPU tax" on targets sharing NVMe drives, it allows for complete converged deployments without the normal SDS penalty. This allows NVMesh to scale performance linearly at near 100% with a virtual, distributed non-volatile array without requiring additional dedicated storage servers or appliances.

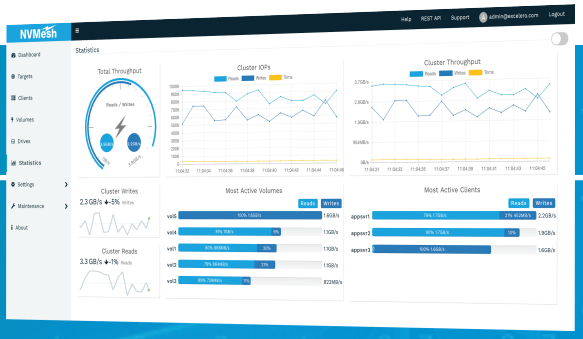

NVMesh FEATURES

- Distributed block storage
  - Logical block volumes
  - Any local or distributed file system
- Elastic NVMe
  - Share NVMe between Servers or add external NVMe
  - Local NVMe performance across the network
- Scale-out architecture
  - No controller bottlenecks
  - Scale performance and capacity as needed (any scale)
- Software-defined Storage
  - Use standard servers
  - Leverage state-of-the art storage and networking