Excelero

SCALABLE AND PERSISTENT STORAGE FOR CONTAINERS

SOLUTION BRIEF

INTRODUCTION

While container technology has only been available for a few years, we are seeing significant growth of container applications in enterprise production environments. The container ecosystem is booming: with plenty of solutions for container management, orchestration, monitoring and security. The top pain point for container environments is 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, following the examples set by tech giants like Google and Amazon.

FAST-GROWING CONTAINER MARKET

The promise of container technologies is application portability, scalable application design and efficient hardware utilization. Per 451 Research's latest Cloud-Enabling Technologies Market Monitor report, the application container market is expected to grow to $2.7bn by 2020. With an estimated CAGR of 40% through 2020, this growth is much faster than, for example, virtualization PaaS or Devops. While the technology is much newer, containers may even have a broader impact on the market than OpenStack. This is also confirmed by the breadth and diversity of emerging offerings based on or supporting containers: orchestration platforms, monitoring and management tools, and security layers.

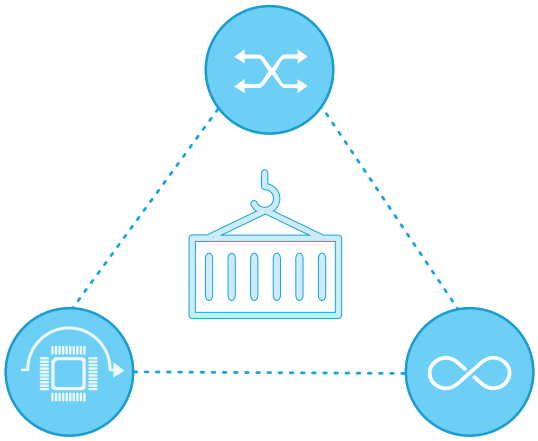

In spite of this fast-growing ecosystem, many customers are still holding back on deploying container technology at massive scale due to the lack of tools for managing persistent container storage. Most companies have started moving the stateless parts of their applications to containers first and are 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 server SAN supports persistent, low latency container storage for hyperscale 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. 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 be striped, mirrored or both while enjoying centralized management, monitoring and administration.

In short, applications/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.

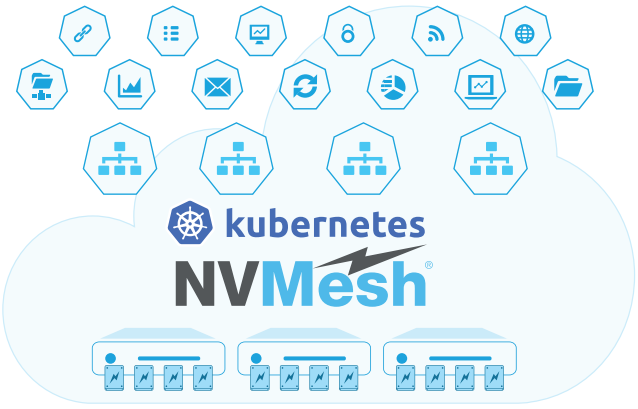

The solution was designed for infrastructure architects and developers who want to embrace microservices style deployments but desire more performant approaches than NFS or persistent volumes on traditional all-flash-arrays (AFAs).

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. It is also assumed that not more than 20% of the available bandwidth will be used for recovery of redundancy, so that a reasonable level of service is maintained during the rebuild process. This number should be tunable to match the administrator's preference between risking data loss versus reducing service levels during rebuild.

Excelero's NVMesh 100% Server SAN 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.

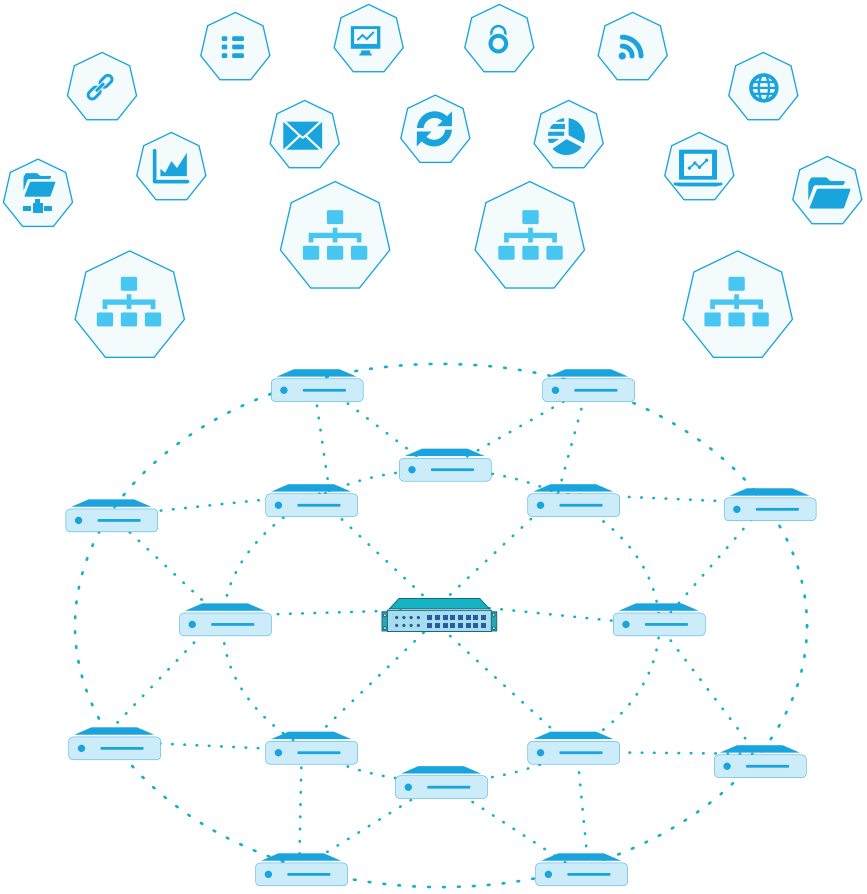

NVMesh BENEFITS FOR MICROSERVICES

SCALE & PERFORMANCE

- 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
- Leverage high IOPS, high bandwidth or mixed

EFFICIENCY

- Maximize the utilization of your NVMe flash devices
- Choose hardware from any server, storage and network vendor
- Easy to manage & monitor, reduces the maintenance TCO
- Balance CPU and storage resources

FLEXIBILITY

- 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
- Mix different storage media types to optimize for cost, scale or performance
- Scale storage and compute separately, as needed

ABOUT EXCELERO NVMesh

NVMesh features Elastic NVMe, a distributed block layer that allows any application 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. In short, applications 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 is deployed as a virtual, distributed non-volatile array and supports both converged and disaggregated architectures, giving customers full freedom in their architectural design.