Flash Memory Summit

NVMe over Fabrics (NVMe-oF) For Containers

Scaling Cloud-Native Applications With Elastic And High-Performance Storage

TOSHIBA

Sudhakar Mungamoori
Director Storage Solutions, Toshiba Memory America Inc.

Venkat Ramakrishnan
VP Engineering, Portworx

portworx

Flash Memory Summit 2018
Santa Clara, CA

What is NVMe over Fabrics (NVMe-oF™)?

A new storage networking protocol, optimized for flash, *exclusively for NVMe commands*

Based on *Industry-standard* Networks

Networked flash, with *the performance and latency of native NVMe*

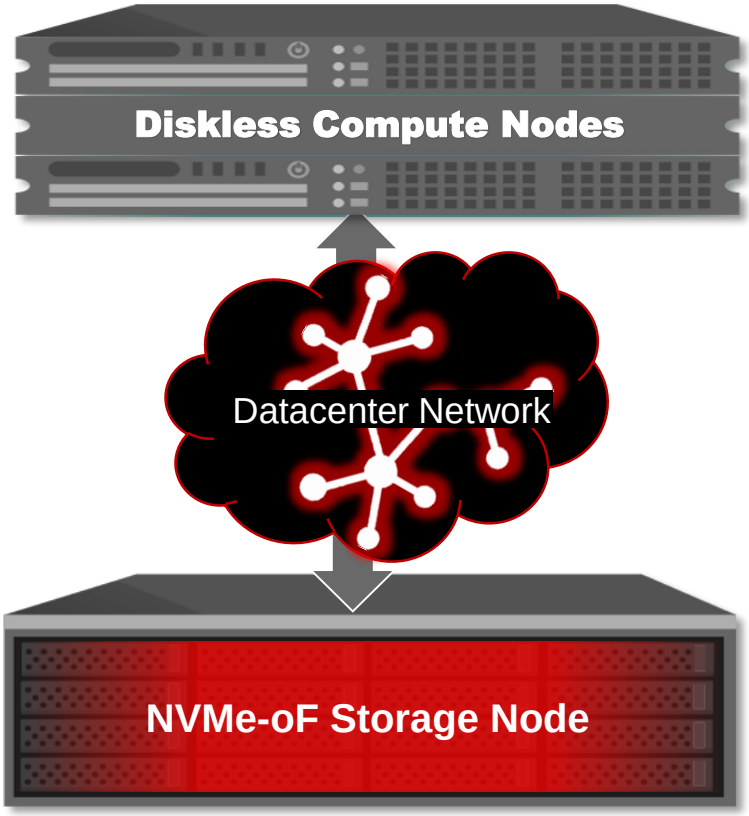

NVMe-oF Enables "Shared Accelerated Storage"

Direct Attached SSDs

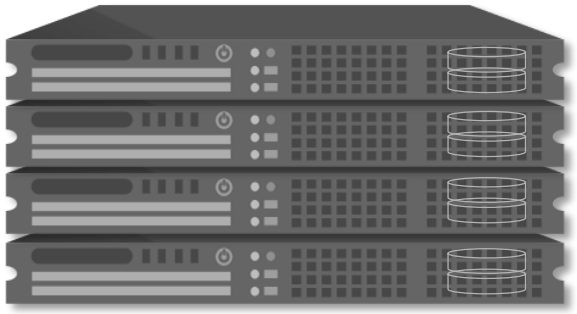

No matter what size SSD you choose, it may miss the mark; tough to optimize configs.

Pooled Storage with NVMe-oF

Pool the physical drives, dynamically create any size "logical drive" for each node

With NVMe-oF - Any service can run on any server
Each service gets "just the right amount" of high performance, low latency storage

Toshiba KumoScale™ Shared Accelerated Storage Software

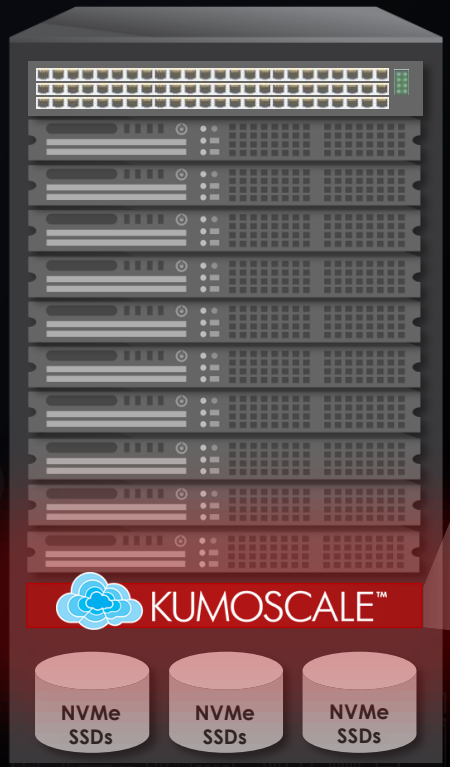

Kubernetes®
OpenStack®
Mesosphere®
Intel® RSD

RESTful API

Management Engine

Virtualization / Abstraction

Fabric Module
(front end)

NVMe Module
(back end)

Network Driver

NVMe Driver

NVMe-oF Network

NVMe SSDs

Toshiba KumoScale Software

Container Orchestration Transforming Datacenter Cloud Architectures

- Modernize on Commodity Hardware
- Enable Operational Agility
- Build Elastic Private and Hybrid Clouds
- Efficient Utilization of Infrastructure

Challenges With Stateful Containers

- Modern Applications are Data intensive and Require Stable, Persistent and Performance Storage
- Expensive to Migrate Stateful Applications Across Nodes
- DAS Single-Node Solution for Persistent Storage is Inelastic
- Dynamic Storage Scaling is Impossible in DAS Architectures

NVMe-oF Solves Container Storage Problems

TOSHIBA + portworx

- Dynamic provisioning of NVMe storage namespaces per container.
- Fast, isolated, highly available storage volumes for containers – Scheduler Integrated.
- Storage orchestration that optimizes for performance and cost
- Application aware live migration for performance and scalability.

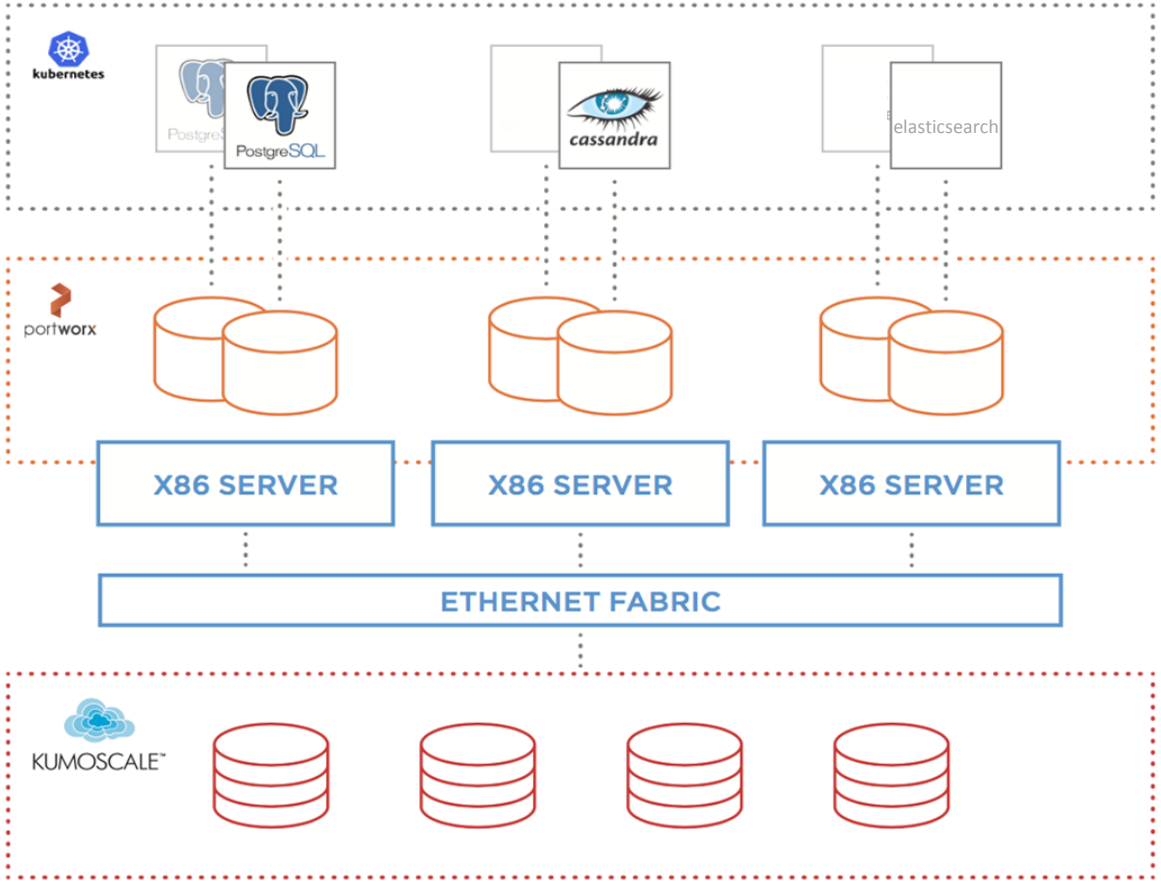

Portworx Introduction

Storage for DevOps

Portworx takes existing infrastructure

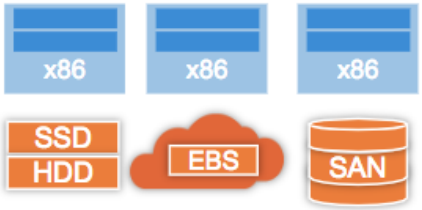

Fingerprints, aggregates, tiers, and monitors storage media

(elastic container storage)

integrates with schedulers

Integrates provisioning, control, and cluster scale (up to 1000 servers)

and enables self-service IT for stateful containers

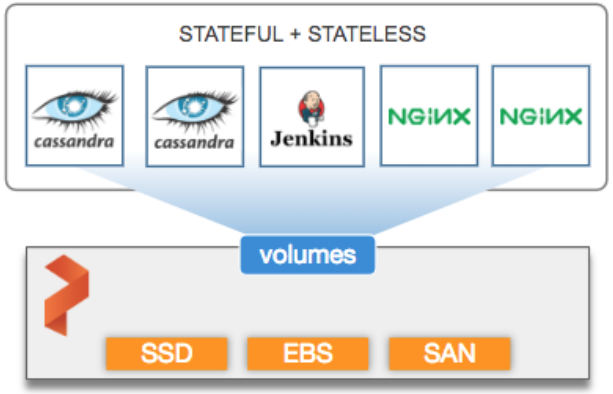

Dynamically creates volumes with schedulers, resizes, encrypts, and moves across clouds

KumoScale NVMe-oF : Portworx API Integration

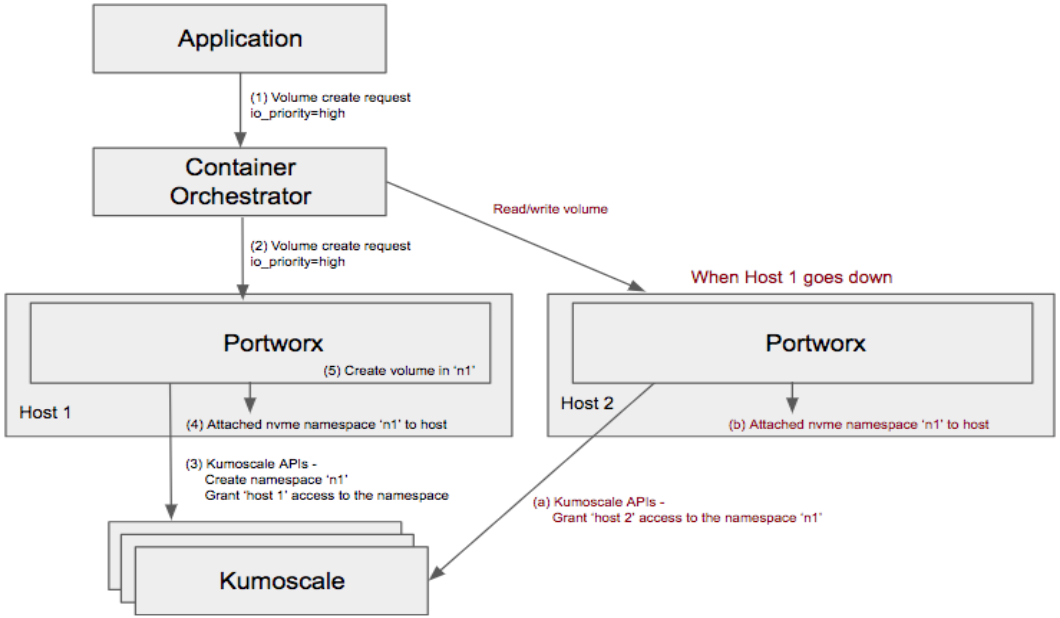

Reference Topology – Kubernetes Use-case

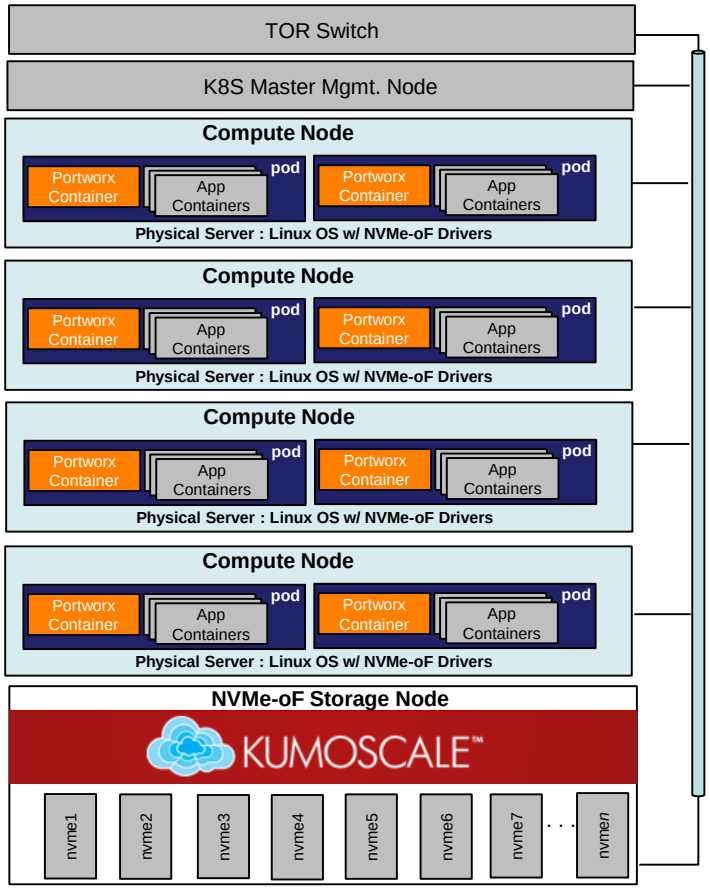

Compute Node

- One or more Pods w/ containers per physical server.

Portworx Container

- PX dynamically manages pool of storage (Kumoscale NVMe-oF volumes) to hold data for stateful containers
- PX provides replication of data across volumes.

K8S Master

- Manages scheduling and orchestration Container replication and restart to defined policy.

KumoScale Node

- One or more Kumoscale NVMe-oF Storage Nodes.

KumoScale NVMe-oF : Portworx
Demo Video (1m 15s)

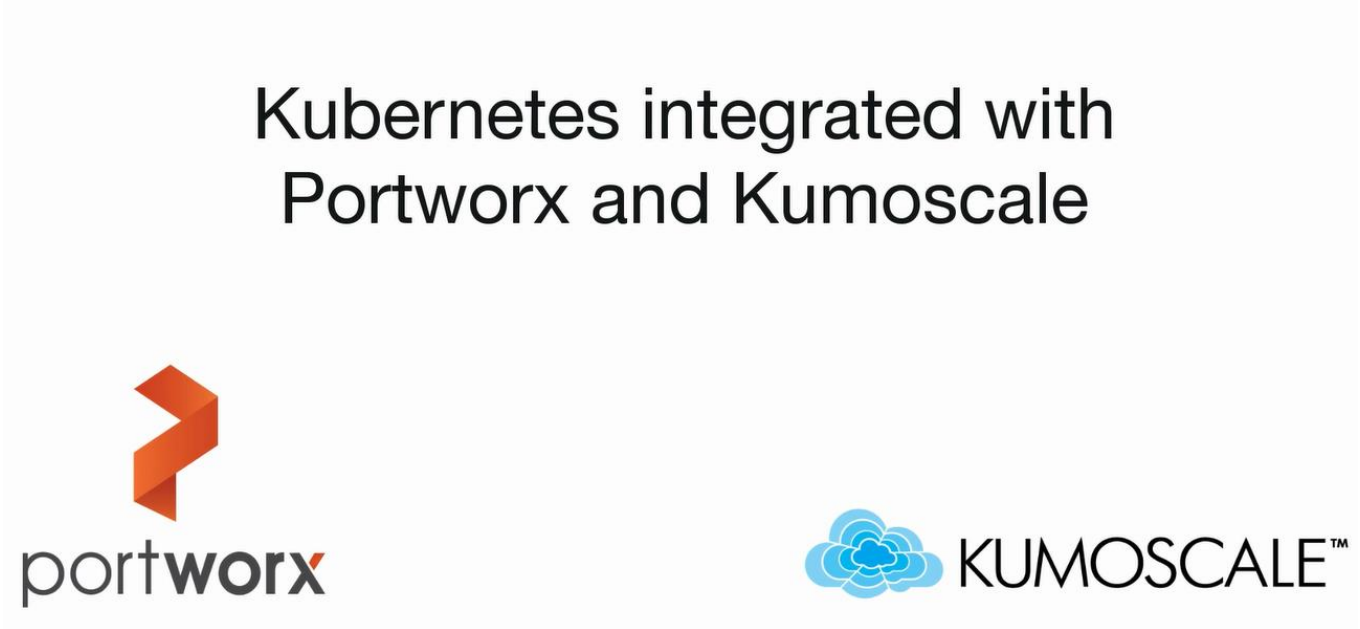

Notes:
- KumoScale is a trademark of Toshiba Memory Corporation.
- NVMe and NVMe-oF are trademarks of NVM Express, Inc.
- Kubernetes is a registered trademark of Linus Torvalds in the U.S. and other countries.
- The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.
- Mesosphere is a registered trademark of Mesophere, Inc.
- Intel is a trademark of Intel Corporation or its subsidiaries in the U.S. and/or other countries.
- Docker and the Docker logo are trademarks or registered trademarks of Docker, Inc. in the United States and/or other countries.
- PostgreSQL is a trademark or registered trademark of The PostgreSQL Global Development Group or its subsidiaries in Canada, the United States and/or other countries.
- Cassandra is a trademark of The Apache Software Foundation.
- Elastisearch is a registered trademark of Elasticsearch BV.
- The Jenkins logo is a registered trademark of Software in the Public Interest, Inc. in the United States. https://jenkins.io/
- NGINX is a trademark of Nginx Software Inc.
- All other company names, product names, and service names mentioned herein may be trademarks of their respective companies.

Information in this presentation, including product pricing and specifications, content of services, and contact information is current and believed to be accurate on the date of the publication, but is subject to change without prior notice. Technical and application information contained here is subject to the most recent applicable Toshiba product specifications.

Thank You

Visit Toshiba Booth #307 For Live Demo of NVMe over Fabrics

Sudhakar.Mungamoori@taec.toshiba.com
Venkat@portworx.com