Solution White Paper

Managing Dynamic Workloads in the Cloud Data Center

Toshiba and Portworx enable scalable storage for disaggregated NVMe-oF™ clouds

TOSHIBA

portworx

Executive Summary

In a world of dynamic workflows, pressure is high to deliver optimal performance and ensure cost-efficiencies in the cloud data center. Infrastructure professionals are exploring new strategies and technologies to ensure the private cloud can support evolving business needs. With the wider adoption of NVMe™ SSDs and the development of NVMe™ over Fabrics (NVMe-oF™), disaggregated storage offers a promising approach for increasing performance and reducing data center expense.

As the market leader in cloud native storage for containers, Portworx has been at the forefront of improving efficiencies in the multi-cloud world. Together with Toshiba, Portworx has introduced the world's first Kubernetes storage solution integrated with NVMe-oF by using Toshiba KumoScale™ shared accelerated storage software. The resulting solution helps companies with unpredictable demand environments automate the microservices data lifecycle, improve the efficiencies of scale out, and simplify ongoing management of the private cloud.

Today's business dynamics put cloud infrastructure professionals under ever-increasing pressure. Whether managing varying workloads, charged to reduce costs, or tasked with optimizing capacity while delivering continuous performance improvements, infrastructure teams face a constant compute and storage juggle. As data intensive technologies like artificial intelligence and machine learning become more pervasive, the need to process exponential amounts of data fast is only growing.

At the same time workloads are expanding, internal business customer demands have become less forgiving. The presumption of immediate IT changes with low latency services is pervasive. This requires those responsible for the cloud infrastructure to fully leverage datacenter resources while squeezing out every possible ounce of performance improvement. New technologies play an essential role in this. From cloud scale infrastructure to SSDs to interconnected technologies such as 100 Gigabit Ethernet (GbE), organizations are looking for advantages in all areas.

Shifting Strategies

In the not-so-distant past, public cloud seemed to be the holy grail for more scalable performance improvements. As the approach has matured, we're seeing companies recognize its strengths for certain applications while shifting other

workloads to a private cloud for cost savings, performance and greater control. Many organizations have found that a hybrid approach is best for business, but across both public and private cloud environments we've seen that shared CPU and storage bottlenecks can hamper performance and scalability. As automated orchestration frameworks have matured, you now have more options for how to design your infrastructure for optimal performance.

Assessing Your Options

Understanding the price-performance trade offs and benefits of new technology is key to delivering the performance you need for the long-term. Currently there are different ways to minimize the impact of and solve for bottlenecks.

Ramp up compute and storage to handle peak usage and performance. For example, to manage the spike in shoppers on Black Friday/Cyber Monday, you would add servers to accommodate peak demand. As this is the most simplistic solution, it is also the most problematic. First and foremost, it's expensive. Most companies don't have an unlimited amount to spend on infrastructure. Plus, it's incredibly difficult to effectively manage growing SKUs and inefficient, as you will find yourself with unused compute and storage capacity at most times.

Overprovision SSDs to provide better write performance and longer drive life. Unfortunately, this approach to reducing bottlenecks also limits available storage capacity because some usable capacity is reserved for overprovisioning. In essence, you've solved for one problem only to create a different one.

Convert to all-flash arrays for performance consistency and predictability. Scalability of this approach is challenging due to expense, which is particularly frustrating given the likelihood of underutilization at most times and bottlenecks at peak usage. Additionally, this strategy is not feasible with cloud orchestration and other technologies that facilitate a hybrid environment.

Implement NVM Express (NVMe) protocol to optimize flash SSDs by placing them directly into the compute node. This reduces latency and boosts performance by delivering a vast improvement in drive connectivity directly connected to the server CPU. The downside is that NVMe SSDs often offer far more performance and capacity than a single compute node can consume, which leaves storage stranded and unavailable for other applications. Additionally, if compute fails, the storage capacity becomes unavailable. Generally speaking, this has been considered overkill for most applications.

Disaggregate storage by building unique storage nodes for different applications. This enables faster scaling but can run up against CapEx limitations, lead to higher energy use, and exponentially increase logistical complexity with multiple SKUs.

Without a wholly effective alternative using existing resources, it is clear that organizations need a new deployment model for cloud storage that offers economic efficiency, dynamic optimization, and lower latency. NVMe over fabrics (NVMe-oF) combines some of the core advantages of the alternatives outlined above in a breakthrough protocol that optimizes performance, while eliminating the deficiencies.

Deriving More Value From Both Storage and Compute Infrastructure

NVMe-oF is a storage architecture that runs NVMe commands over a fabric transport to enable NVMe-based storage to be shared across the network and scaled to thousands of devices with similar performance to direct attached storage. By allowing diverse applications to dynamically share high performance, high capacity NVMe SSD storage, this disaggregated approach makes NVMe flash much more efficient and affordable.

When disaggregated NVMe-oF is implemented, storage capacity is no longer constrained by the number of drives

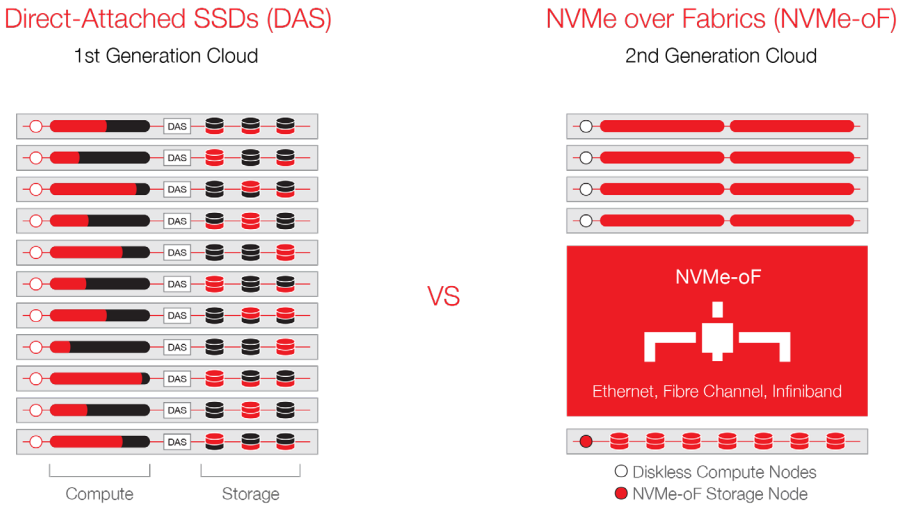

A complete NVMe-oF software stack includes a compatible Linux kernel, NIC and NVMe drivers, and shared accelerated storage software.

each server can support. NVMe-oF offers high bandwidth and throughput and can handle large numbers of queues and commands. This makes both rapid scaling and stateful application live-migration of running containers feasible.

With Ethernet able to offer 100Gb/second per link with microsecond latencies, NVME-oF completes the end-to-end optimization of the stack to deliver the maximum benefits of low-latency NVMe SSDs.

NVMe-oF also makes it possible to reduce server SKU count by eliminating the need for higher-powered servers for specific workloads or peak periods. Instead, general-purpose servers deliver broadly sustainable support for most uses. And since NVMe-oF uses existing network protocols such as Fibre Channel, Ethernet, iWARP, RoCE, or Infiniband, it can be deployed today. As a result, NVMe-oF enables a more flexible, scalable infrastructure.

KumoScale Shared Accelerated Storage Software

Taking your datacenter to a new cost/performance level with NVMe-oF requires software to manage NVMe flash storage, as well as interface to automated cloud applications/tools. Toshiba KumoScale shared accelerated storage software seamlessly integrates with widely-available orchestration and provisioning frameworks, including Kubernetes®, OpenStack®, and Intel RSD to facilitate NVMe-oF implementations. KumoScale manages NVMe SSD storage functionality, enabling the creation of networked storage nodes that can be immediately deployed at scale. It also self-provisions networking and management interfaces to optimally support diverse application demands. This means you can live migrate stateful applications across compute nodes, abstract namespaces, manage NVMe SSDs to maximize performance and endurance, and mask physical drive complexity.

NVMe-oF in the Real World

The combined impact of shared accelerated storage software and data services automation were the key to the world's first Kubernetes storage target integrated with NVMe-oF. Using Toshiba flash storage, KumoScale shared accelerated storage software, and Portworx data resiliency and scheduler integration for stateful containers, this solution optimizes dynamic creation of highly available container-granular storage.

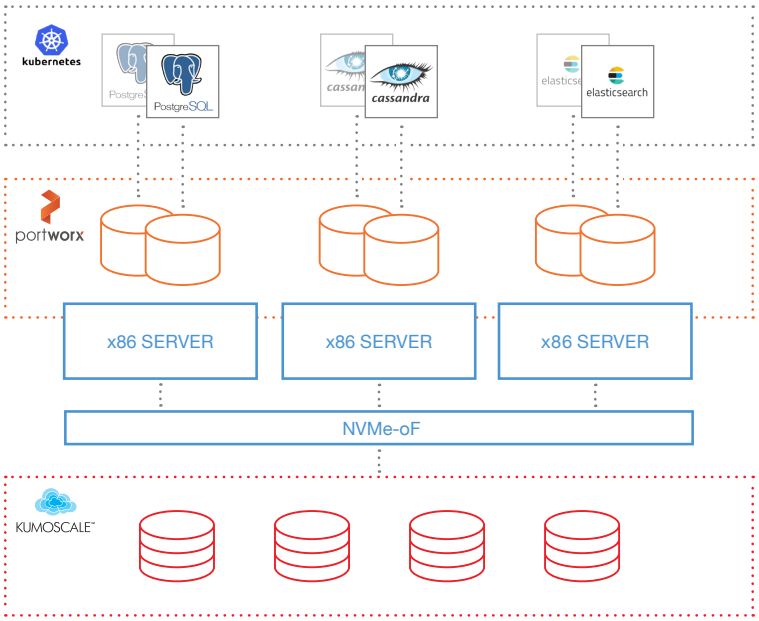

The Portworx and Toshiba KumoScale solution automatically connects and provisions out of the box enabling easy installation and direct management of mission critical databases and other stateful services via Kubernetes. Starting with a user request to deploy a stateful container via the Kubernetes command line kubectl, Portworx communicates with KumoScale to provision an NVMe-oF storage namespace from the underlying Toshiba NVMe SSDs. This namespace is formatted by Portworx and mounted into the docker container created by Kubernetes. These volumes are provisioned within milliseconds as applications are launching, which eliminates manual provisioning processes that slow down common operations like deployment and failover. All applications get their own NVMe storage created as it is needed.

In addition to dynamic volume provisioning, Portworx protects data against hardware-level failures by replicating data across NVMe namespaces, enabling high availability of databases running on Kubernetes. These volumes also have enhanced security via encryption using customer-controlled keys.

Through Kubernetes integration, Portworx enables automation of the entire data and application lifecycle, for example snapshotting a group of volumes and moving them to a secondary disaster recovery site, cloning a data set to use in a machine learning workflow, or resizing a volume without any downtime. The NVMe-oF model can also help companies maximize performance in bursty or unpredictable demand environments through automation of the microservices data lifecycle. Hence, companies can run significantly more dense microservice patterns and scale out effectively across thousands of services, multiple sites, and multiple clouds, if needed. The disaggregation and reduced complexity of NVMe-oF offers greater flexibility to run apps on the right infrastructure and simplifies ongoing management by reducing the "noisy neighbor" problem and facilitating real-time ops and optimization through a common data services layer in the private cloud. This disaggregated approach allows you to manage and scale workloads independently for greater efficiency, giving companies cloud-like elasticity and the price/performance benefits of owning the hardware.

Transforming the Data Center

NVMe-based SSDs offer unparalleled performance for the demanding workloads of today and beyond. With NVMe-oF, organizations can fully exploit this leading-edge network technolgy to better utilize infrastructure and operations investments. And as the inventor of flash memory and a pioneer in SSDs, Toshiba continues to lead the way in standards development collaboration for NVMe and NVMe-oF. Advancing new approaches and protocols for disaggregation and accessing the full potential of flash storage is key to delivering performance with real-world and long-term utility. As workloads fluctuate and the demands on data expand exponentially, having an end-to-end approach for scaling and provisioning in a dynamic world are key to continued success. With disaggregated NVMe-oF, you have the efficiency and agility you need to adapt to today's shifting market realities.

TOSHIBA

Toshiba Memory America, Inc. is the U.S.-based subsidiary of Toshiba Memory Corporation, a leading worldwide supplier of flash memory and solid state drives (SSDs). From the invention of flash memory to today's breakthrough 96-layer BiCS FLASH™ 3D technology, Toshiba continues to lead innovation and move the industry forward.

To learn more about Toshiba Memory America, visit: www.toshiba.com/tma

portworx

Portworx is the cloud-native storage company modern enterprises trust to manage data in containers. Portworx dramatically reduces storage, compute and infrastructure costs for running mission critical multi-cloud applications with zero downtime or data loss for customers such as GE Digital, Verizon, NIO and Lufthansa Systems. The company is headquartered in Los Altos, Calif., and investors include Mayfield, Sapphire Ventures and GE Ventures.

For more information about Portworx, visit portworx.com or follow @portwx.