Extending the Power of Virtualization to Storage

WHITE PAPER

vmware®

Server virtualization has changed the way IT runs data centers across the world. According to Gartner, as of 2014 more than 70 percent of x86 server workloads have already been virtualized[1]. Most enterprises now operate with a "virtual first" policy where applications are deployed on virtual platforms by default.

Despite the broad adoption of server virtualization, the storage architecture that supports it has not kept pace. Server virtualization imposes new requirements on storage, increasing the demand for storage systems capable of responding to the highly random, spiky I/O profiles of virtualized applications and new workloads like VDI. The emergence of cloud-based services and rapid development and deployment of new applications increase end-user expectations for capacity, performance, and agility.

From a management standpoint, virtualization brings a new foundational storage construct through which VMs consume storage services – the virtual disk – and increases the expectations for rapid storage provisioning and ease of ongoing management in a dynamic virtual environment with consolidation of multiple VMs and mixed workloads.

Regardless of these new storage demands, today's data center storage architectures are largely driven and defined by hardware. Traditional methods for provisioning and maintaining storage resources are no longer sufficient to achieving business goals. While traditional storage system capacity may increase or new, higher performance systems are introduced, the fundamental challenges of a hardware-centric approach remain – it's complex, rigid, and inefficient.

Traditional Storage and Virtualization is a Mismatch

Data storage has become one of the most significant pain points in enterprise IT[2]. Over the next two years, storage capacity is growing at 41 percent per year[3] and IT admins are having trouble storing, cataloguing and making this data available at all times. Despite traditional external network storage systems being the proven method to store data, there are pain points associated with high capital expenditures, challenges with scaling, complexity of operations and difficulty in meeting SLAs, the last one being the most pressing according to IDC data[4].

Traditional Storage is Silo'd and Complex

While modern storage arrays offer powerful capabilities to store, manage and protect data, they have poor knowledge of dynamic application requirements, and are often unable to react quickly to changes.

Typical storage systems have their own constructs for delivering storage (LUNs, volumes), handled through separate workflows by different teams. Critical data services (snapshots, replication, caching, dedupe, etc.) are often tied to specific storage arrays, inhibiting standardized approaches across multiple storage types. These traditional storage environments demand specialized, (often expensive) vendor-specific skillsets, isolated storage teams, and fragmented management processes.

Meeting Storage SLAs and Aligning to Business Needs is Cumbersome

Provisioning and ongoing management of traditional storage is a rigid process, making it difficult to support dynamic application service levels. It requires up-front planning to create pre-configured static pools of capacity with pre-defined performance settings and a particular set of data services – independent from the application. This hardware-centric, bottoms-up approach leads to inefficient provisioning. Applications are mapped to these fixed pools of capacity in an effort to identify the "best fit" to their requirements. As applications change, storage requirements change and the VMs need to be migrated to a new LUN or volume. Performance scaling is not granular, requiring laborious planning and large investment outlays. With the lack of end-to-end visibility from the virtual layer down into the storage layer it is hard to troubleshoot and monitor performance issues for each VM. Overall, a lack of application awareness from the storage system makes today's operational model slow to adapt to changes.

---

1. Gartner x86 Server Virtualization Magic Quadrant, 2014
2. ComputerWeekly.com, June 2013
3. IDC, Worldwide Enterprise Storage Systems Forecast, November 2013
4. IDC, Storage Predictions 2014, January 2014

CapEx is High and Scaling is Expensive

The rigidity and complexity of existing storage environments leads to inefficient storage consumption and ROI. To avoid storage provisioning complexity and delays, organizations often resort to requesting more storage than they truly need. This overprovisioning leads to storage resources being allocated, but not necessarily consumed. And since storage resources aren't virtualized, the idle storage can't be easily reallocated or shared, forcing potentially more frequent purchases. Consequently, not only do storage systems end up representing a large portion of the IT budget, but scaling storage resources also happens in chunks by either paying for expensive upgrades or forklifting existing hardware to buy costly new storage infrastructure on a recurring basis.

A New Approach is Needed

To meet today's challenges, IT needs a fundamentally more efficient operational model– one that's aligned to the business and can deliver agility and efficiency. VMware's goal is to bring to storage the same efficient operational model that server virtualization brought to compute. This new model changes how storage is provisioned, and managed – it's defined by software, not the hardware. And it starts with the application, not the storage array. Software-Defined Storage, and the Software-Defined Data Center, enable a more agile and efficient environment, delivering speed, simplicity and cost-effective solutions.

Software-Defined Data Center

The Software-Defined Data Center is an open architectural approach that extends the principles of virtualization — abstraction, pooling, and automation — to all data center resources and services. The Software-Defined Data Center is the ideal infrastructure that enables the continuous development of today's mobile applications, dramatically speeding up and simplifying the initial provisioning and ongoing management of the virtualized infrastructure of compute, networking and storage. VMware's Software-Defined Data Center approach leverages core data center virtualization technologies to transform data center economics and business agility through automation and non-disruptive deployment that embraces and extends existing compute, network and storage infrastructure investments.

How Does it Work?

Software-Defined Storage relies on the hypervisor and two new architectural elements - the virtual data plane and the policy-driven control plane:

- **Virtual data plane** – virtualizes physical resources, abstracting them into logical, VM-centric pools of capacity that can be flexibly consumed and managed. VMware's implementation is vSphere Virtual Volumes™ on SAN/NAS devices, and VMware Virtual SAN™ on x86 server storage.
- **Policy-driven control plane** – acts as the bridge between applications and infrastructure, enabling common management across heterogeneous storage tiers and dynamic control of storage service levels through VM-centric, policy-driven automation. VMware's implementation is Storage-Policy-Based Management.

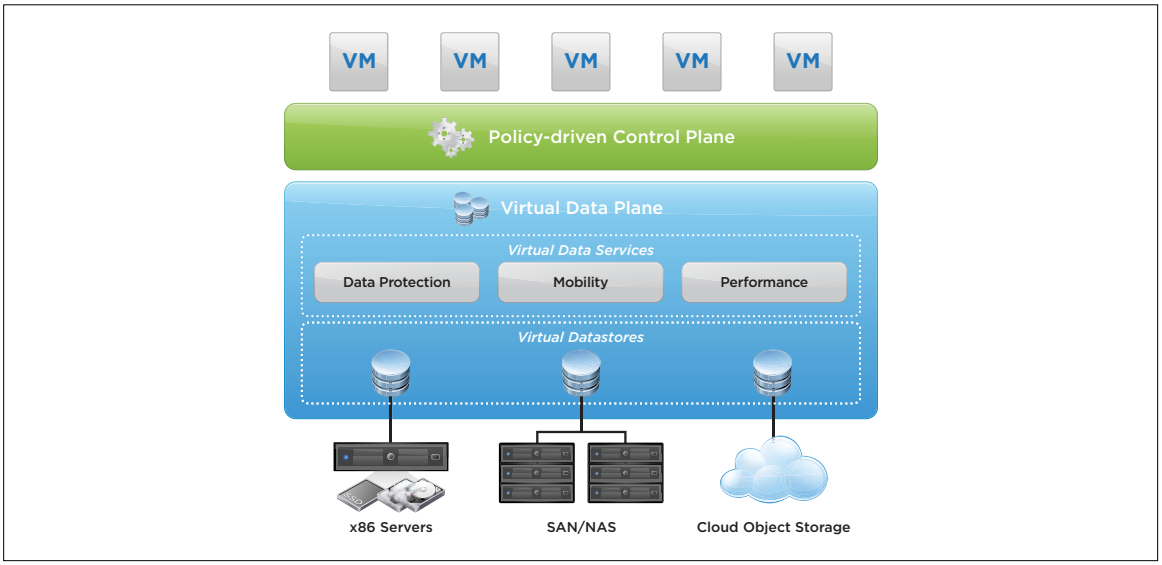

Figure 1: Software-Defined Storage Architecture

Virtual Data Plane

This is a new abstraction layer that pools storage into flexible VM-centric virtual datastores where virtual disks are natively represented on underlying physical infrastructure. Using Virtual Volumes and Virtual SAN, the storage infrastructure expresses available capabilities (performance and data services) to the virtual data plane to enable automated provisioning and dynamic control of storage services levels through programmatic APIs.

Compare this to traditional storage environments that are predicated on rigid infrastructure-centric LUNs (Logical unit numbers) or storage volumes that essentially are static allocations of storage service levels (capacity, performance and data services). In this situation, storage provisioning is dictated by the hardware available, and applications are rigidly assigned resources in a "best fit" scenario with little room for flexibility or change.

Policy-Driven Control Plane

This new management layer provides common orchestration and automation across all storage tiers whether on external arrays, x86 servers or cloud storage. Unlike traditional control planes, this policy-driven control plane uses policies to define different classes of storage and automates the composition of these storage services at scale, while enabling adjustments at the VM level. As application requirements change, policies are updated and the hypervisor leverages the virtual data plane and virtual datastores to automatically and non-disruptively provision the new storage resources. The control plane is programmable via public APIs that can be used to consume and control policies via scripting and cloud automation tools for self-service consumption of storage.

In a typical storage environment, the control plane is tied to each storage device – each array with its own its own management tools, generating uniquely defined "classes of service" or static pre-allocations of resources, difficult to change and largely disassociated from specific application requirements.

The Hypervisor's Role

As the abstraction layer sitting at the intersection between applications and infrastructure, the hypervisor occupies a privileged position within the IT stack that allows it to balance all IT resources – compute, memory, storage and networking – needed by an application. The hypervisor is inherently application-aware and has direct visibility into each application (i.e. virtual machines and virtual disks) and its requirements. With a virtual data plane and policy-based management in place, the hypervisor now has the ability to provision storage resources in a similar fashion to compute resources. Now the hypervisor can:

- Automate the delivery of storage levels (capacity, performance) and data services to applications through a new control plane.
- Abstract the storage infrastructure into flexible pools of capacity where these storage levels and data services can be delivered to individual virtual machines.
- Manage and dynamically provision storage more precisely and efficiently through storage policy-based management templates.

Making Software-Defined Storage a Reality

VMware's Software-Defined Storage model shifts the operational model of storage from the bottoms-up array-centric approach of today's storage, to a tops-down VM-centric model. As a result, storage services are automatically and precisely aligned to application requirements, saving you time and money.

VMware Virtual SAN

Implementing Software-Defined Storage starts with VMware Virtual SAN, which provides radically simple, hypervisor-converged storage for virtual environments. Virtual SAN delivers enterprise-class, high performance storage for any virtualized application, including business critical applications. The hardware-independent, scale-out architecture leveraging x86 servers can drastically lower overall storage TCO while providing the flexibility to scale in a linear, predictable manner. By pooling server-attached flash-based storage (in a flash-only or hybrid SSD/HDD architecture) Virtual SAN can deliver high performance and predictable response times. Its VM-centric management approach, along with its seamless integration with VMware vSphere® and the entire VMware stack, makes it the simplest storage platform for VMs.

Virtual SAN can be configured as hybrid or all-flash storage. In a hybrid storage architecture, Virtual SAN pools server-attached HDDs and SSDs to create a distributed shared datastore that abstracts the storage hardware and provides a software-defined storage tier for virtual machines. Flash is used as a read cache/writer buffer to accelerate performance and magnetic disks are used for data persistence. Alternately, Virtual SAN can be deployed as an all-flash storage architecture in which flash is intelligently used only as a write cache while SSDs provide high endurance data persistence.

VMware vSphere Virtual Volumes

Extending the Software-Defined Storage model to existing storage investments allows IT managers to experience the benefits of a flexible storage infrastructure throughout the data center. vSphere Virtual Volumes is VMware's integration framework to abstract and pool existing SAN and NAS systems into the virtual data plane. Virtual Volumes defines a new virtual disk container (Virtual Volume) independent of the underlying physical storage representation (LUN, file system, object, etc.). With Virtual Volumes, the virtual disk becomes the primary unit of data management at the array level. This turns the virtual datastore into a VM-centric pool of capacity allowing application-specific requirements to drive storage provisioning decisions while leveraging the capabilities provided by existing storage arrays. Virtual Volumes allows IT managers to continue to extract the benefits from current storage systems, while embracing the new operational model of Software-Defined Storage.

Storage Policy-based Management

Both Virtual Volumes and Virtual SAN rely on VMware's Storage Policy-Based Management (SPBM) operational model. SPBM is a common storage policy framework that allows control and automation of VM-centric storage policies across Virtual SAN and through Virtual Volumes, external storage.

SPBM captures storage service levels requirements (capacity, performance, availability, etc.) in the form of logical templates (policies) to which VMs are associated. SPBM automates VM placement by identifying available datastores that meet policy requirements and enables the hypervisor to dynamically instantiate necessary data services. Through policy enforcement, SPBM also automates service level monitoring and compliance throughout the lifecycle of the VM. When an application's storage requirements change, the policy is updated and changes are automatically and instantly provisioned.

Software-Defined Storage: The Evolution of Storage

In order to adequately keep up with the rising stress on the storage infrastructure in our now virtualized world, a new approach is necessary. The answer is Software-Defined Storage, a new approach to storage that enables a fundamentally more efficient operational model. With Software-Defined Storage, IT admins can deliver the agility, speed and cost-effective storage solutions needed for today's virtual environments.

Read more about VMware Software-Defined Storage at: http://www.vmware.com/software-defined-datacenter/storage.html

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com