What's New: VMware Virtual SAN 6.1

VMware Storage Business Unit Documentation
v 1.0/August 2015

vmware®

Contents

INTRODUCTION **3**
1.1 OVERVIEW OF SOFTWARE-DEFINED STORAGE (SDS) 3
1.1.1 Data Plane *5*
1.1.2 Policy-Driven Control Plane *6*
1.2 OVERVIEW OF HYPER-CONVERGED INFRASTRUCTURES (HCI) 7
1.3 VMWARE VIRTUAL SAN (VSAN) 9
1.4 CONFIGURATION AND MANAGEMENT 11
1.5 HARDWARE CHARACTERISTICS 12
1.6 WHAT'S NEW VMWARE VIRTUAL SAN 6.1 14
1.6.1 Enterprise Availability and Data Protection *15*
1.6.2 Maximum Performance and Low Latencies *18*
1.6.3 Enhanced Usability and Management *19*
1.6.4 Advanced Management & Troubleshooting *20*
CONCLUSION **22**
ACKNOWLEDGEMENTS **22**
ABOUT THE AUTHOR **23**

Introduction

1.1 Overview of Software-Defined Storage (SDS)

VMware's Software-Defined Storage vision and strategy is to drive transformation through the hypervisor, bringing to storage the same operational efficiency that server virtualization brought to compute.

As the abstraction between applications and available resources, the hypervisor can balance all IT resources – compute, memory, storage and networking – needed by an application. With server virtualization as the de-facto platform to run enterprise applications, VMware is uniquely positioned to deliver Software-Defined Storage utilizing the pervasiveness of this software tier.

By transitioning from the legacy storage model to Software-Defined Storage with Virtual SAN, customers will gain the following benefits:

- **Automation of storage class of service at scale:** Provision virtual machines quickly across data center using a common control plane (SPBM) for automation.
- **Self-Service capabilities:** Empower application administrators with cloud automation tool integration (vRealize Automation, PowerCLI, OpenStack).
- **Simple change management using policies:** Eliminate change management overhead and use policies to drive infrastructure changes.
- **Finer control of storage class of service:** Match virtual machine storage requirements exactly as needed with class of service delivered per virtual machine.
- **Effective monitoring/troubleshooting with per virtual machine visibility:** Gain visibility on individual virtual machine performance and storage consumption.
- **Safeguard existing investment:** Use existing resources more efficiently with an operational model that eliminates inefficient static and rigid storage constructs.

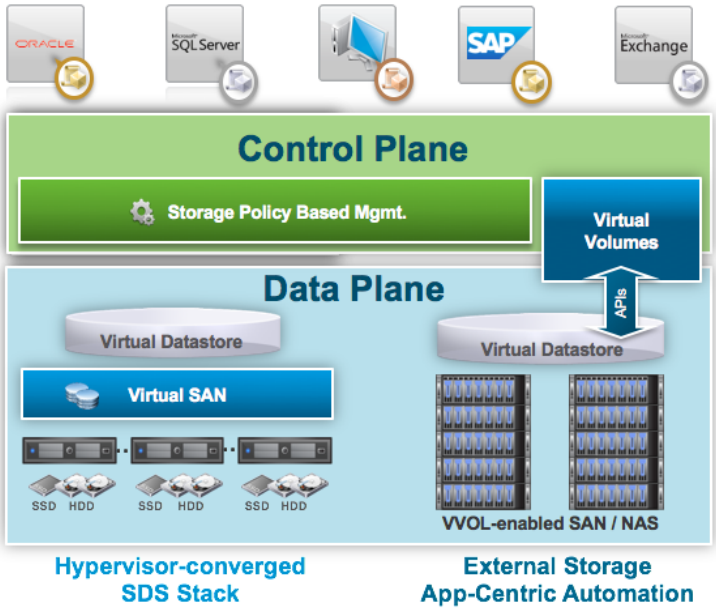

Figure 1: Software-Defined Storage Conceptual Diagram

The goal of Software-Defined Storage is to introduce a new approach that enables a more efficient and flexible operational model for storage in virtual environments. This is accomplished in two ways:

- Abstracting the underlying hardware into logical, pools of capacity that can be more flexibly consumed and managed in a virtual machine-centric fashion. This is analogous to the functions served by server and network virtualization, and VMware refers to this storage virtualization as the Data Plane.

- The abstraction of the Data Plane enables additional functions that an array may provide to be offered as data services for consumption on a per-virtual machine basis. Data Services can provide functionality such as compression, replication, snapshots, de-duplication, availability, migration and data mobility, performance capabilities, disaster recovery, and other capabilities. While the data services may be instantiated at any level of the infrastructure, the virtualized data plane allows for these services to be offered via policy on a per-VM basis.

- Implementing an automation layer that enables dynamic control and monitoring of storage services levels to individual virtual machines across heterogeneous devices – VMware refers to this as the Policy-Driven Control Plane

1.1.1 Data Plane

The Data Plane is responsible for both, storing data and applying data services (snapshots, replication, caching, etc). While data services may be provided by a physical array or implemented in software, the virtual data plan abstracts the services and will present them to the policy-driven control plane for consumption and applies the resultant policy to the objects in the virtual datastore.

In today's model, the data plane operates on rigid infrastructure-centric constructs (LUNs or storage volumes) that typically are static allocations of storage service levels (capacity, performance and data services), independently defined from applications.

In the VMware Software-Defined Storage model, the data plane is virtualized by abstracting physical hardware resources and aggregating them into logical pools of capacity (virtual datastores) that can be more flexibly consumed and managed. Additionally, to simplify the delivery of storage service levels for individual applications, the Data Plane makes the virtual disk the fundamental unit of management around which all storage operations are controlled. As a result, exact combinations of data services can be instantiated and controlled independently for each virtual machines.

For each virtual machine that is deployed, the data services offered can be applied individually: Each application can have its own unique storage service level and capabilities assigned to it at its time of creation. This allows for per-application storage policies, ensuring both simpler yet individualized management of applications without the requirement of mapping applications to broad infrastructure concepts like a physical datastore.

In the Software-Defined Storage environment, the storage infrastructure expresses the available capabilities (performance and data services) to the control plane to enable automated provisioning and dynamic control of storage services levels through programmatic APIs.

These storage services may come from many different locations: Directly from a storage array, from a software solution within vSphere itself, or from a third party location via API. These capabilities are given to the control plane for consumption and expression by policies. The ability to pull in multiple sources of data services and abstract them to a policy engine gives the administrator the ability to create unique policies for each virtual machine in accordance with their business requirements, consuming data services from different providers in each

VMware's implementation of the Data Plane is delivered through VMware's own Virtual SAN in the case of x-86 hyperconverged storage and vSphere Virtual Volumes in the case of external SAN/NAS arrays.

1.1.2 Policy-Driven Control Plane

In the VMware Software-Defined Storage model, the control plane acts as the bridge between applications and storage infrastructure. The control plane provides a standardized management framework for provisioning and consuming storage across all tiers, whether on external arrays, x86 server storage or cloud storage.

The policy-driven control plane is the management layer responsible for controlling and monitoring storage operations. In today's model, the control plane is typically, tied to each storage device – each array is operated in a different way - and implements a "bottom-up" array-centric approach in which storage service levels are aggregated into physical tiers or "classes of services", which are static pre-allocations of resources and data services tied to the infrastructure.

Upon provisioning, an application is rigidly mapped to these pre-configured storage containers. These storage containers are rarely aligned to precise application boundaries, and their capabilities need to be defined broadly to encompass the requirements of a broad set of applications.

This restricts the ability of a storage container to be focused specifically on the business requirements of an individual application. To circumvent this problem, storage administrators may be asked to create numerous purpose-built datastores, increasing the management overhead and complexity associated with storage.

Through Software-Defined Storage, the storage classes of service become logical entities controlled entirely by software and interpreted through policies. Defining and making adjustments to these policies enables automating the provisioning process at scale, while dynamically controlling individual service levels over individual virtual machines at any point in time.

This makes the Software-Defined Storage model able to more flexibly adapt to ongoing changes on specific application requirements. Policies become the control mechanism to automate the monitoring process and to ensure compliance of storage service levels throughout the lifecycle of the application.

The control plane is programmable via public APIs, used to consume and control policies via scripting and cloud automation tools, which in turn enable self-service consumption of storage to application tenants as well as a variety of external management frameworks.

VMware's implementation of the policy-driven control plane is delivered through Storage Policy-Based Management (SPBM).

1.2 Overview of Hyper-converged Infrastructures (HCI)

Traditional external storage systems still tend to be deployed in silos creating the need for specialized, vendor-specific skillsets and fragmented management processes. Additionally, storage CAPEX is high because scaling, upgrades or replacements tend to be very expensive.

Today's businesses demand less complex technology services that don't require legions of IT staff and specialized experts to operate. Instead, businesses want their IT functions to operate with the cloudlike efficiency of companies like Google and Amazon.

The ubiquity of datacenter virtualization coupled with technology advances like high density flash, non-volatile memory technologies and increasing CPU densities of multicore processors has led to the emergence of a new architecture for SDDC – Hyper-Converged Infrastructure (HCI).

What is a Hyper-Converged Infrastructure (HCI)?

A new approach to the Software-Defined Data Center, called Hyper-convergence, is emerging. Hyper-Converged Infrastructure (HCI) leverages the hypervisor to deliver compute, networking and shared storage from a single x86 server platform.

The software-driven architecture allows the convergence of physical storage onto commodity x86 servers, enabling a building block approach with scale-out capabilities. The use of commodity x86 server and storage hardware allows datacenters to operate with agility on a highly scalable, cost-effective infrastructure.

The foundational components of VMware's Hyper-Converged Infrastructure solution are VMware vSphere and VMware Virtual SAN. Together they allow the convergence of compute, storage and networking onto a single, integrated layer of software that, can run on any commodity x86 infrastructure.

VMware vSphere abstracts and aggregates compute and memory resources into logical pools of compute capacity while Virtual SAN, embedded in vSphere, pools server-attached storage to create a high performance, shared datastore for virtual machine storage.

Customers can further extend vSphere to achieve the convergence of networking using VMware NSX which decouples network resources from underlying hardware and brings the operational model of a virtual machine to the datacenter network.

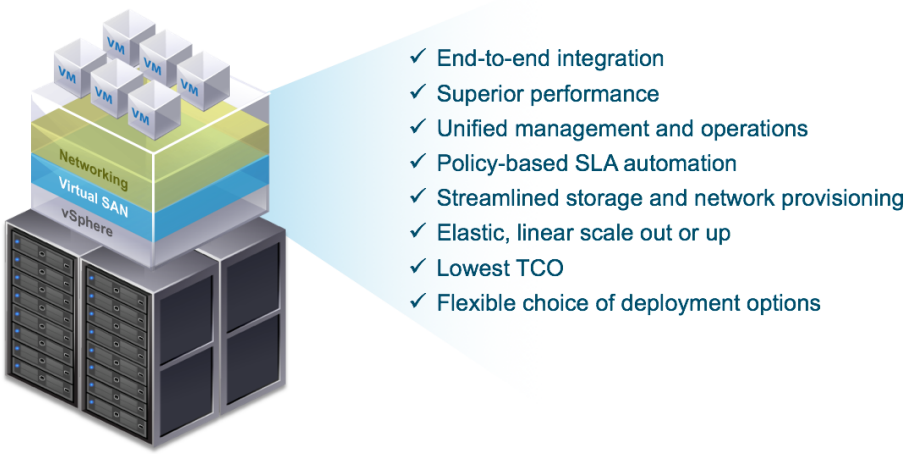

Figure 2: VMware's Hyper-Converge Infrastructure

By transitioning to this new Hyper-Converged infrastructure model with VMware vSphere and VMware Virtual SAN, customers will gain the following benefits:

- **End-to-end integration and unified management:** VMware HCI solution provides a single layer of compute, networking and storage software. Management is unified through common tools and interfaces. Features like HA, vMotion, DRS etc. work seamlessly across the VMware stack.

- **Streamlined Provisioning and Automation:** SLAs can be controlled through 'virtual machine-level' policies that can be set and changed on-the-fly. Virtual SAN dynamically self-tunes and load balances, adapting to changes in workload conditions to ensure that each virtual machine has the storage resources it needs, as defined by its policy. The end-to-end vSphere integration and policy-driven approach automates manual tasks and makes management of compute, storage and networking extremely easy.

- **Elastic, linear scale out or up:** Virtual SAN architecture allows for elastic, linear and non-disruptive scaling. Capacity and performance can be scaled simultaneously by adding new hosts (scale-out); or capacity and performance can be scaled independently by adding new drives to existing hosts (scale-up).

- **Lowest TCO:** The "grow-as-you-go" scaling approach as well as the ability to use commodity hardware, implies that the overall TCO of the HCI solution is dramatically lower than using standalone hardware or converged systems.

- **Choice of deployment options:** VMware provides the broadest choice of deployment models with multiple options across vendors. Certified hardware through Virtual SAN Ready Nodes guarantee flexibility while the EVO family of Integrated Systems is designed for simplification of procurement, deployment and management.

1.3 VMware Virtual SAN (VSAN)

Virtual SAN is VMware's enterprise-class storage software-defined solution for Hyper-Converged Infrastructure. It is seamlessly embedded in the hypervisor, designed to deliver enterprise class, elastically scalable, high performance shared storage for vSphere Virtual Machines that is radically simple to consume and manage.

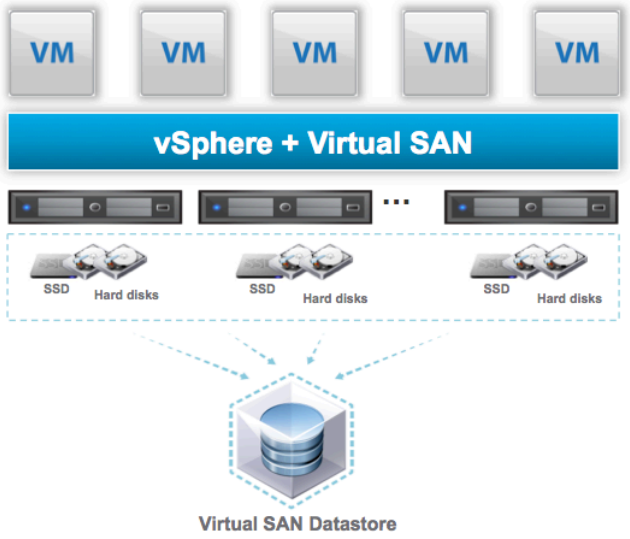

Figure 3: VMware Virtual SAN

Virtual SAN pools server-attached magnetic and flash devices to create a distributed shared datastore that abstracts the storage hardware and provides a hyper-converged compute, storage, and network optimized solution for virtual machines. VMware customers can benefit from the many unique values delivered by this transformational technology:

- **Radically Simpler** – Deploy easily through the vSphere web client and automate management using storage policies

- **Elastic Scalability** – Scale out or up performance and capacity by adding a new host to the cluster or new drives to existing hosts. Start small with 3 hosts and scale all the way to 64 hosts per cluster.

- **High Performance** – being embedded in the hypervisor Virtual SAN can deliver up to 7M IOPS with all flash or 2.5M IOPS with hybrid configurations

- **Lower TCO** – Leveraging standard x86 hardware components and reducing operational overhead can lower storage TCO by up to 50%

Virtual SAN's tight integration with the VMware stack of features, and products work seamlessly with no additional overhead.

Virtual SAN is designed to deliver to core storage services for virtualized production environments, with greater performance, lower latency, greater scalability and flexibility and at a fraction of the price. Just like vSphere can be deployed on a wider range of hardware options, and it's applicable to a variety of use cases. For example:

- **Virtual Desktop Infrastructure (VDI)** – Traditional storage solutions can account for 38% of the total cost per desktop, Virtual SAN, included with Horizon Advanced and Enterprise, can reduce the cost of virtual desktops to as low as $36/desktop by leveraging server-side economics providing significant cost savings and a grow as you go model that can scale out or up – while providing better performance and better scalability

- **IT Ops (management cluster, DMZ)** – IT Ops applications are consumed directly by IT people and are critical to manage the IT environment. To enable full control of these applications, an easy way to upgrade, troubleshoot, and provide security isolation, a separate set of resources is required. Virtual SAN provides a separate datastore with an efficient way to manage all the storage operations, critical for these type of applications.

- **Private Cloud** – Requires rapid provisioning of virtual machines and high level of automation on both compute and storage. Virtual SAN storage policy based management enables a self-service environment that can provision storage to VMs with complete SLA automation.

- **Remote IT** – usually characterized with no to a small number of IT people with limited amount of time and domain expertise. Virtual SAN simple deployment and management model removes the need for specialized IT staff and reduces the time needed to manage the storage in these locations.

- **Test/Dev (DevOps)** – Test/Dev environments require both easy and fast ways to provision VMs for an efficient application development cycle. Virtual SAN's ability to quickly spin up and down VM storage with the required SLA provides the Test/Dev team the much necessary agility they require.

1.4 Configuration and Management

From setup to on-going management Virtual SAN is radically simple; it is embedded directly in the hypervisor and does not require any additional software to be installed – it can simply be enabled in a few clicks.

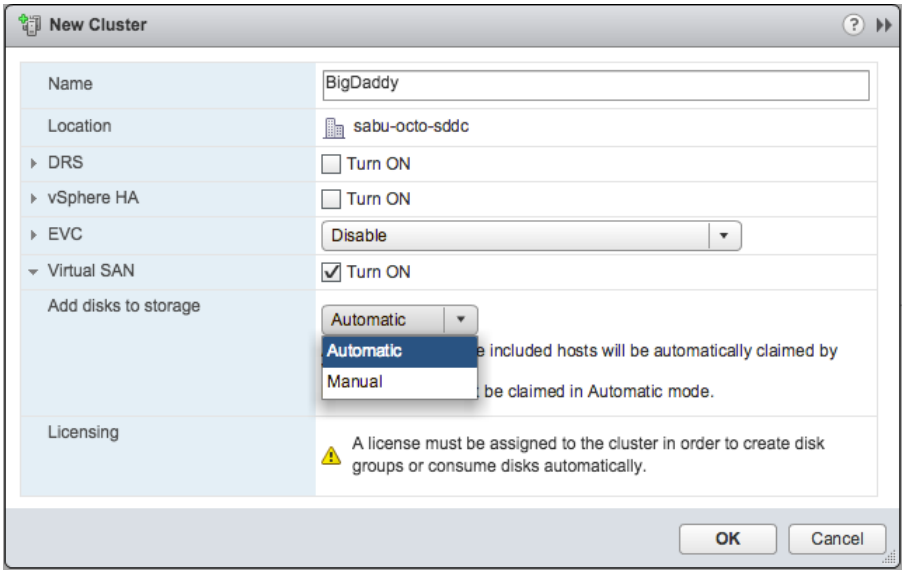

Figure 4: Enable Virtual SAN

Managing Virtual SAN doesn't require any specialized skillset as it can be managed end-to-end through the vSphere Web Client and vCenter Server instances, tools that are already familiar to any vSphere Administrator. It integrates with the entire VMware stack, including features like vMotion, HA, DRS etc.

Virtual machine storage provisioning and day-to-day management of storage SLAs can all be controlled through VM-level policies that can be set and modified on-the-fly. Virtual SAN dynamically adapts to changes in workload conditions to ensure that each virtual machine has the storage resources it needs, as defined by its policy.

This hypervisor integration and per-VM policy-driven approach automates manual storage tasks and makes management of storage in virtual environments easy and seamless.

1.5 Hardware Characteristics

By virtue of being software-defined and delivered from the hypervisor, Virtual SAN is completely hardware-independent and works with any x-86 vSphere compatible server allowing customers to continue leveraging their server vendor of choice.

Figure 5: Virtual SAN Hardware Options

Virtual SAN is the only HCI solution in the industry that allows maximum choice of deployment models, ensuring that customers have the flexibility to choose from a wide selection of hardware vendors and configurations. VMware offers deployment models with multiple options across vendors.

- **The New Virtual SAN Ready Nodes Series**: provides customers with the widest range of configurations and vendor choice for the greatest level of configuration flexibility.

  Customers can pick any of the VMware and OEM-validated server configurations available from a menu of prescriptive options ready for Virtual SAN deployment.

  Each Ready Node is based on a specific series model has specific components and quantity for different use cases/workloads.

Old Profiles Brand	New Profiles Brand
Hybrid Profiles	
Hybrid-Server-Low	HY-2 Series
New Hybrid Profile	HY-4 Series
Hybrid-Server-Medium	HY-6 Series
Hybrid-Server-High	HY-8 Series
Hybrid-VDI-Linked Clones	HY-8 Series
Hybrid-VDI-Full Clones	HY-8 Series
All Flash Profiles	
All Flash-Server-Medium	AF-6 Series
All Flash-Server-High	AF-8 Series
All Flash-VDI-Linked Clones	AF-8 Series
All Flash-VDI-Full Clones	AF-8 Series

Figure 6: New Virtual SAN Ready Node Series

Ready Nodes can be ordered as-is or further customized for different needs. Customers can build their own SDDC by adding VMware SDDC components like VMware NSX, vRealize Suite etc. Ready Node options can be found at: http://www.vmware.com/resources/compatibility/search.php?deviceCategory=vsan

- **Build Your Own:** Component-based approach that allows transforming existing vSphere hosts into Virtual SAN nodes using certified components from a broad HCL for different performance profiles, form factors across vendors.

- **Integrated Systems through the EVO Family** that are powered by VMware's HCI software (vSphere, Virtual SAN, management software etc.) are pre-integrated, pre-configured systems with certified partner hardware. They include additional software to simplify deployment, configuration and lifecycle management.

Regardless of the deployment model of choice, Virtual SAN supported hardware options are based on industry standard storage components.

1.6 What's New VMware Virtual SAN 6.1

Now in its 3rd release, Virtual SAN has been already adopted by a large number of VMware customers in just 15 months. Customers of all industries and sizes, trust Virtual SAN to run their most mission critical applications.

Virtual SAN 6.1 adds greater abilities around availability and management that can perform in the most demanding storage environments.

In this new release VMware continues to focus on delivering high performance and increased scale to provide enterprise-class storage for virtualized workloads, including Tier-1 production and business-critical applications.

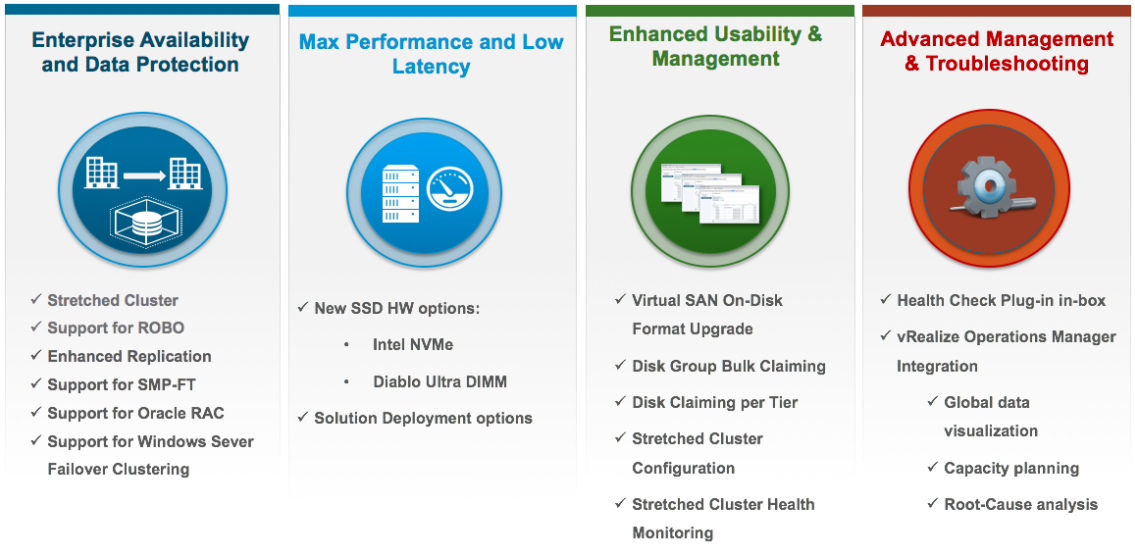

Figure 7: What's New in Virtual SAN 6.1

VMware is expanding the capabilities of Virtual SAN as a platform and is introducing higher levels of enterprise availability and data protection by enhancing the replication capabilities of the solution as well as providing fault tolerance support for multi-processors virtual machines.

A critical aspect of Virtual SAN is ensuring that data is never lost. Virtual SAN 6.1 stretched clusters further enable the ability for no interruption in operations even in the event of a complete site failure.

The most significant new features and capabilities of Virtual SAN 6.1 are described below:

1.6.1 Enterprise Availability and Data Protection

Virtual SAN Stretched Cluster – Virtual SAN 6.1 allows the ability to create a stretched cluster between two or more geographically separated sites, synchronously replicating data between sites.

The stretched cluster feature introduces enterprise level availability where an entire site failure can be tolerated, with no data loss and near zero downtime. In addition, the feature enables for disaster and downtime avoidance by allowing the proactive move of virtual machines from one site to another in order to avoid an impending outage or for planned maintenance purposes.

VMware vMotion, and vSphere DRS are naturally integrated and can help load balance between two active sites, while Virtual SAN insures high performance by making sure virtual machines use the local site copy for read and write operations.

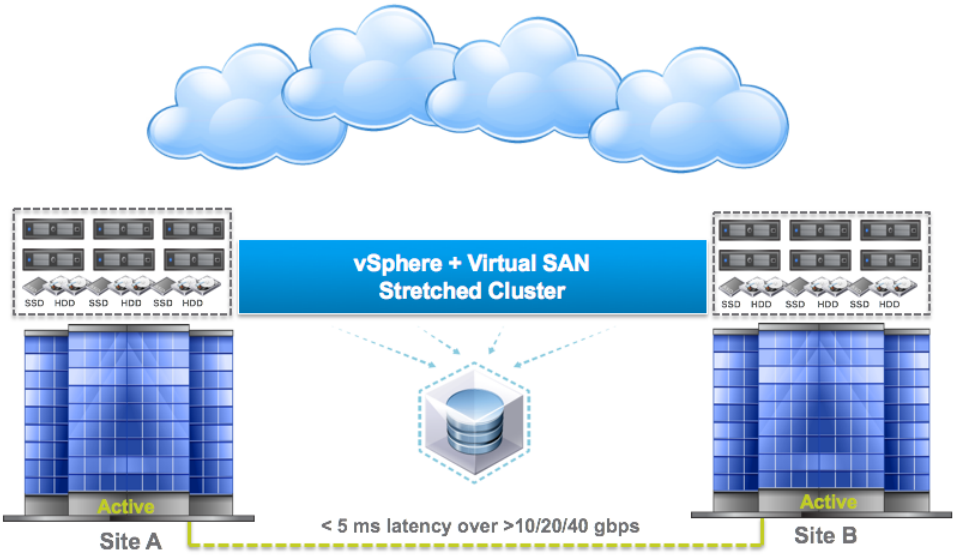

Figure 8: Virtual SAN Stretched Cluster

Because of the tight integration with the VMware stack features like VMware vMotion, vSphere DRS and vSphere HA the solution works seamlessly with no additional overhead. vSphere administrators are already familiar with how to operate and run everything around vSphere - this skillset becomes even more valuable with Virtual SAN.

Virtual SAN removes all the complexity around storage management in a multi-site environment and also simplifies storage tasks to a few clicks in the vSphere web client.

Virtual SAN for Remote Office / Branch Office (ROBO) – Virtual SAN 6.1 provides the ability and support to deploy Virtual SAN clusters for ROBO scenarios. Customers can now deploy a large numbers of 2-node Virtual SAN clusters that can be centrally managed from a centralized data center through an individual vCenter Server instance.

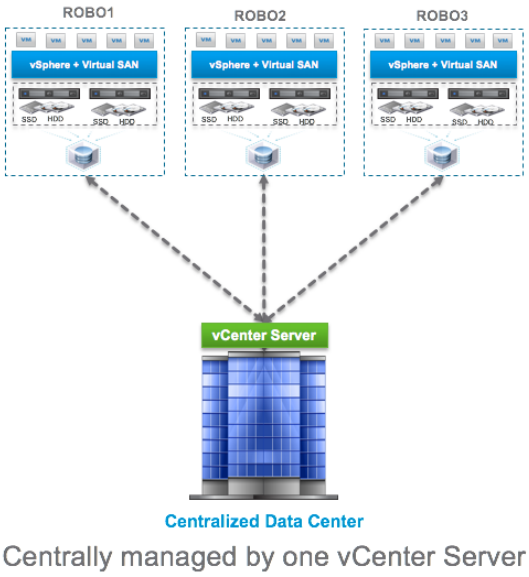

Figure 9: Virtual SAN Remote Office Branch Office (ROBO)

Virtual SAN Replication with vSphere Replication -

Virtual SAN 6.1 utilizes vSphere Replication for its replication data services. vSphere Replication with Virtual SAN supports RPOs from 5 minutes to 24 hours.

VMware Site Recovery Manager (SRM) can be used as part of the solution in order to deliver a fully orchestrated disaster recovery solution.

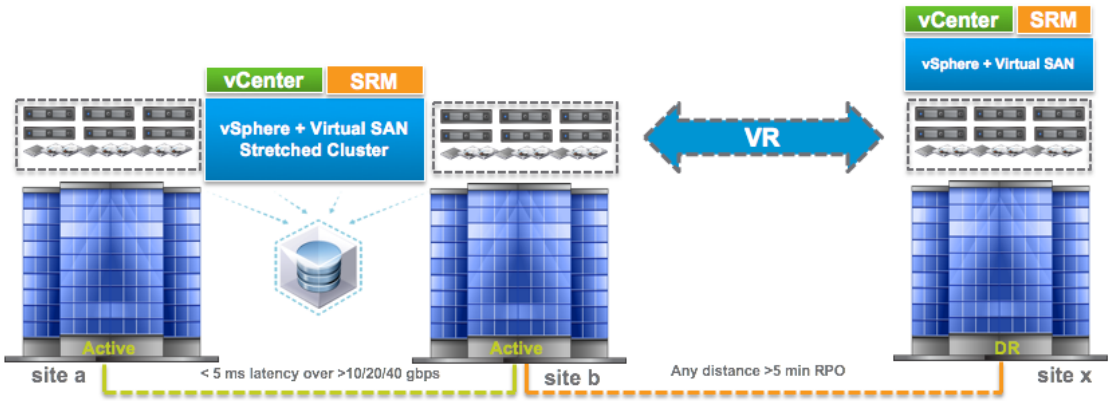

Figure 10: Virtual SAN Replication - 5 Minutes RPO

Support for Multi-Processor Fault Tolerance (SMP-FT) – Virtual SAN 6.1 is now able to support vSphere Fault Tolerance feature and provide continuous availability with zero downtime for mission critical applications in the event of hardware failure.

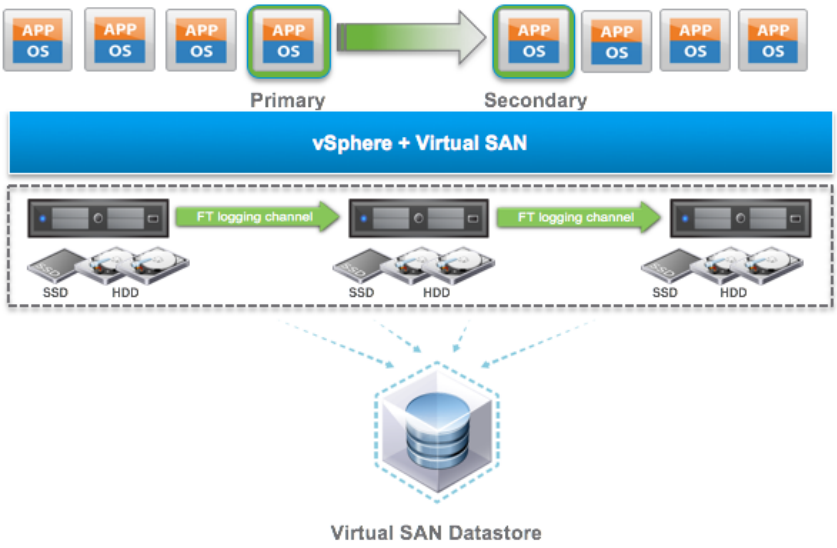

Figure 11: Support for SMP-FT

Support for Windows Server Failover Clustering (WSFC) and Oracle Real Application Cluster (RAC) – Virtual SAN 6.1 now support clustering technologies from both Oracle and Microsoft. With Oracle Real Application Clusters (RAC) customers can run multiple Oracle RDBMS instances accessing the Virtual SAN datastore, delivering better performance, scalability and resilience. Similarly, with Windows Server Failover Clustering (WSFC)) failover clustering functionality is enabled, and can guard against application and service failures.

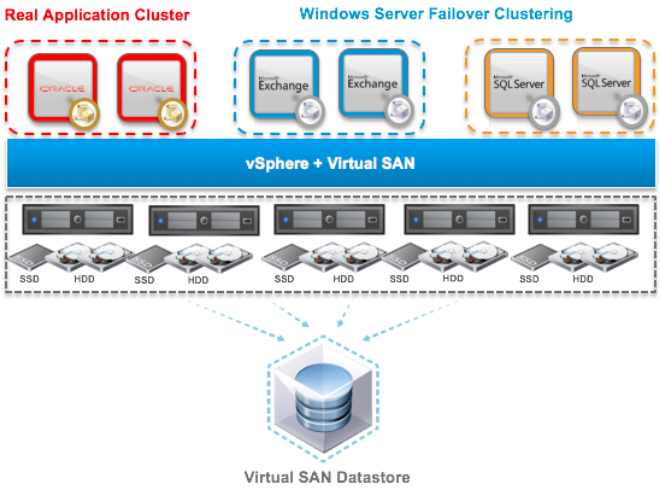

Figure 12: Support for Oracle and Windows Clustering Services

1.6.2 Maximum Performance and Low Latencies

As flash and non-volatile memory technologies continue to evolve rapidly, with the emergence of new flash form factors and new flash interfaces. These technologies have the potential to further reduce storage latencies by an order of magnitude.

In fact, read and write latencies will soon become much faster than doing a network hop. As a result, the only way to really leverage the enormous performance gains of these new technologies is to eliminate that network hop and bring the data much closer to the compute – onto the same server.

VMware continues the performance improvement of Virtual SAN by expanding the support of new type of hardware devices and interfaces.

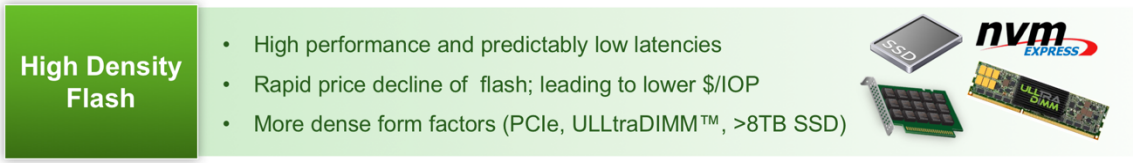

Figure 13: New Supported Hardware and Interfaces

ULLtraDIMM™ – ULLtraDIMM SSDs connect flash storage to the memory channel via DIMM slots, achieving very low (<5us) write latency. Ultra DIMM provides even greater density and performance - for example, it allows for a 12-TB All Flash Virtual SAN host in a thin blade form factor, as well as 3x improvement in latency compared to external arrays.

NVMe – Non-Volatile Memory Express (NVMe) is a new communications interface developed especially for SSDs allowing for greater parallelism to be utilized by both hardware and software and as a result various performance improvements. Leveraging NVMe in a Virtual SAN all flash deployment resulted in 3.2M IOPS measured on a 32-node cluster - ~100K IOPS/Host.

1.6.3 Enhanced Usability and Management

Virtual SAN On-Disk Format Upgrade – Virtual SAN 6.1 introduces the ability to perform the on-disk format upgrade from Virtual SAN 5.5 format version 1.0 to the new VSAN FS format 2.0 with a single click from the UI.

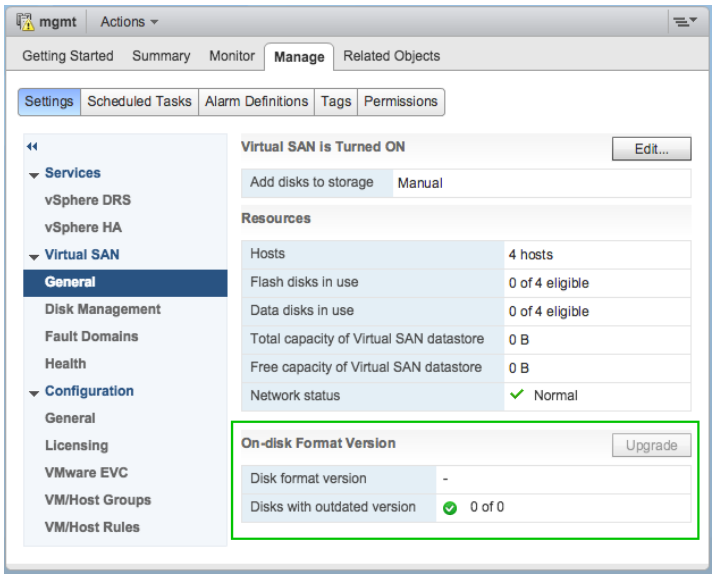

Figure 14: Virtual SAN 6.1 On-Disk Format Upgrade

All-Flash Enhanced usability UI – The Virtual SAN All-Flash architecture introduced in Virtual SAN 6.0 required the use of the CLI to manual configure the flash devices for the capacity tier. In Virtual SAN 6.1, the flash devices for either tier caching and capacity are now claimed in bulk for from the UI.

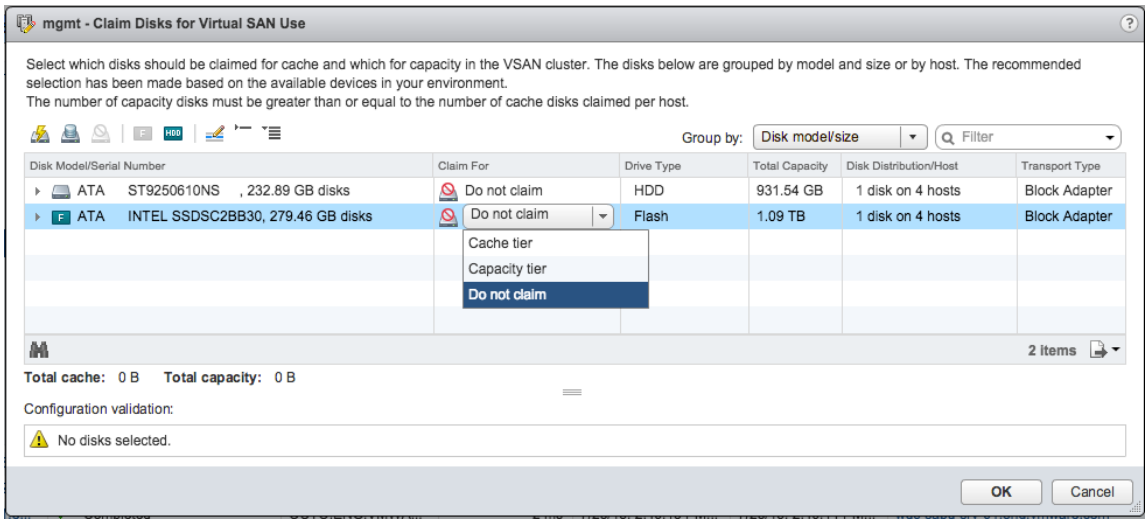

Figure 15: Disk Bulk Claiming

1.6.4 Advanced Management & Troubleshooting

Virtual SAN Health Check-Plug-in – is a vCenter Server plug-in for the vSphere Web Client that is designed to provide and centralized health reporting solution that checks all aspects of the Virtual SAN configuration and health of hardware, data and its limits.

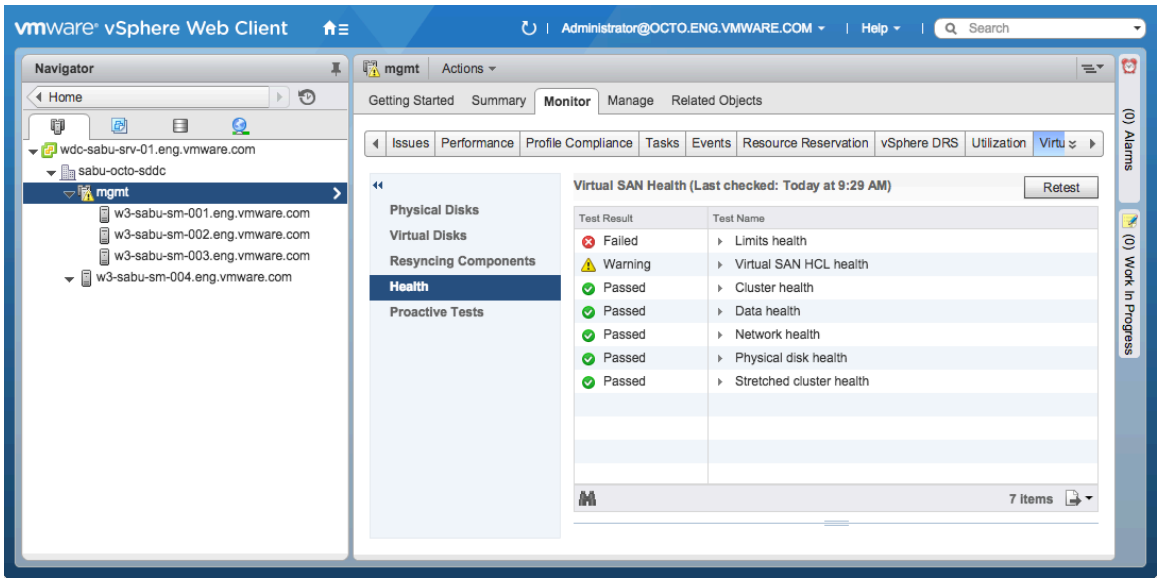

Figure 16: Virtual SAN Health Check Plug-in

- **HCL Compatibility check** – can check the underlying hardware type, firmware and drivers to make sure that it is compatible even before installing Virtual SAN.

- **Hardware Diagnostics** – detects and helps remediate issues with Virtual SAN hardware related issues such as storage device health, network connectivity, and configuration on both host and cluster level.

- **Configuration Management** – makes sure advanced configuration options are uniform across the cluster.

- **Proactive Tests** – provides the ability to perform run real time test of cluster functionalities and dependencies.

    - **VM Creation Tests** – test the ability to successfully create virtual machines in the cluster
    - **Multicast performance test** – access connectivity as well as suitable multicast speeds between hosts
    - **Storage performance test** – tests the stability of the cluster under heavy I/O loads.

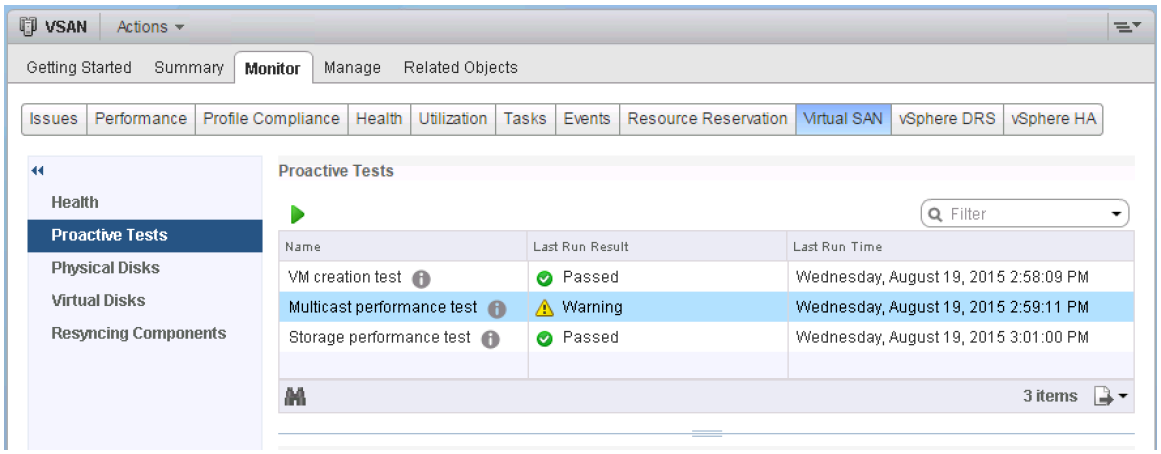

Figure 17: Virtual SAN Health Check Plug-in Proactive Tests

Virtual SAN Management Pack for vRealize Operations – Virtual SAN's 6.1 new integration with vRealize operations delivers a comprehensive set of features to help manage Virtual SAN.

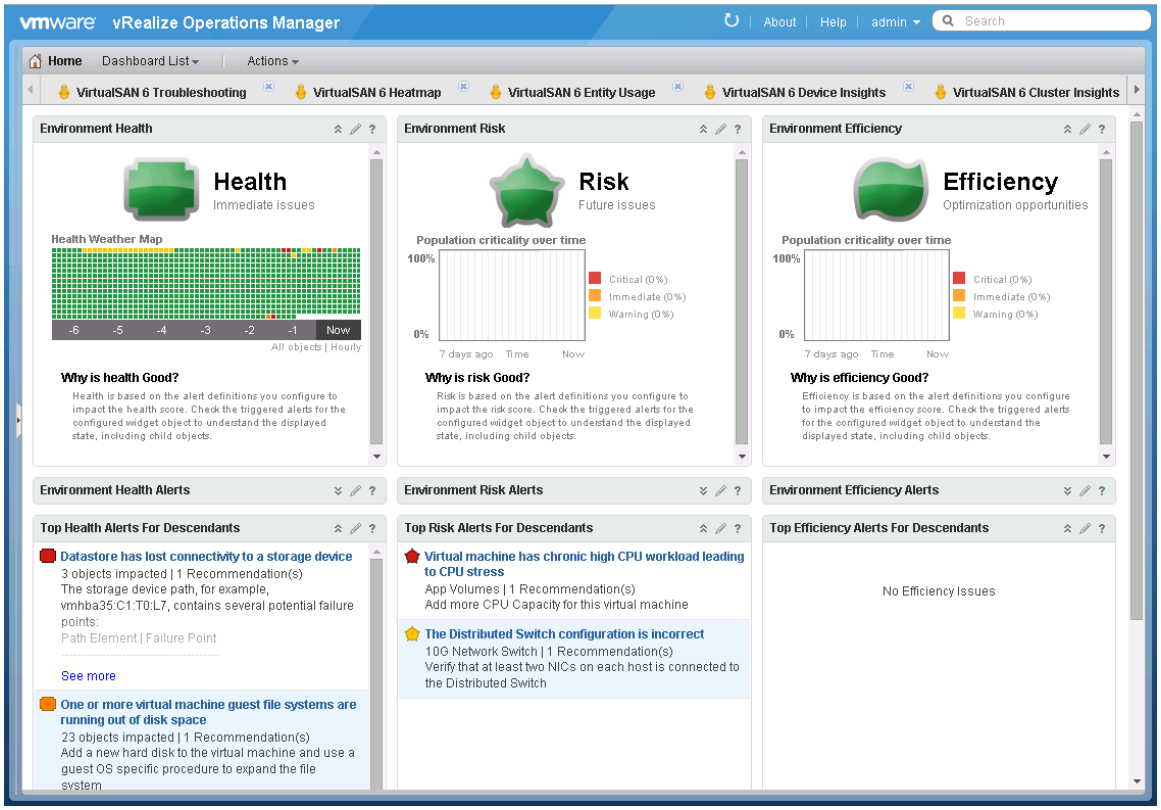

Figure 18: Virtual SAN Management Pack for vRealize Operations

- **Global Data Visualization** – global visibility across multiple Virtual SAN clusters for monitoring, alerts and notifications.

- **Health Monitoring and Availability** – get proactively notified on failures, performance or policy compliance issues on an ongoing basis, with root-cause analysis and remediation strategy.

- **Performance Monitoring** – monitor aggregate performance and latencies on disk or disk group level, as well as CPU, memory and network congestion metrics.

- **Capacity Monitoring and Planning** – monitor disk usage and SSD wear out across all Virtual SAN hosts and use advanced capacity planning based on "what if" scenarios, historic demand and stress trends

Conclusion

The new VMware Virtual SAN 6.1 is a third-generation hypervisor-converged enterprise-class storage solution for vSphere-virtualized infrastructures that combines the compute and storage resources of vSphere hosts.

With it's two supported architectures hybrid and all-flash, and expanded enterprise capabilities in availabilities, data protection, with enhanced usability and management with global monitoring and troubleshooting tools, Virtual SAN 6.1 is able to satisfy the demands of all virtualized applications, including business-critical applications.

Virtual SAN 6.1 is a VMware-designed storage solution that makes software-defined storage a reality for VMware customers that delivers radically simple, hypervisor-converged storage for virtual machines.

Acknowledgements

I would like to thank Rakesh Radhakrishnan Product Manager, Ken Werneburg, Manager of Technical Marketing and Jase McCarty is a Sr. Technical Marketing Manager of the Storage and Availability Business Unit for contributing and reviewing this paper.

About the Author

Rawlinson Rivera is a Principal Architect working in the Office of the CTO in the Storage and Availability Business Unit at VMware, Inc. He specializes in cloud enterprise architectures, Software-Defined storage, Hyper-converged Infrastructures and business continuity / disaster recovery solutions with focus on Virtual SAN, vSphere Virtual Volumes, as well as storage solutions for OpenStack and Cloud-Native Applications. He serves as a partner and trusted adviser to VMware's customers primarily in the US.

Rawlinson is amongst the few VMware Certified Design Experts (VCDX#86) in the world, and author of multiple books based on VMware and other technologies. He is the owner and main author of virtualization blog punchingclouds.com.

- Follow Rawlinson's blogs:

  http://blogs.vmware.com/vsphere/storage
  http://www.punchingclouds.com/

- Follow Rawlinson on Twitter:

  @PunchingCloud

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

Copyright © 2015 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdiction. All other marks and names mentioned herein may be trademarks of their respective companies.