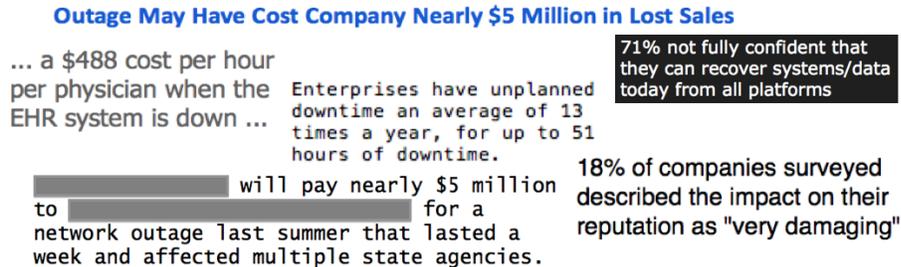


VMware Virtual SAN

Disaster Recovery

Lowering the Risk and Cost of Disaster Recovery

According to the United States Department of Homeland Security, roughly 40 to 60 percent of small businesses never reopen their doors following a disaster. www.fema.gov Large organizations are also susceptible to significant losses from a disaster as shown in numerous online news articles.



The cost of a disaster recovery site can be cost prohibitive for many organizations. As a result, a number of businesses have an inadequate or no disaster recovery plan, which introduces considerable risk. One of the more significant costs of a disaster recovery site is the IT systems infrastructure including server hardware, storage, and replication software. It was common for organizations to deploy similar systems at both the production site and the disaster recovery site. This costly approach was done to help prevent failure in recovering applications and data due to differences in the infrastructure.

Virtualization removes the need to have the same infrastructure hardware at both the production and disaster recovery sites by abstracting the physical hardware from the applications. Furthermore, a virtual machine's configuration, operating system, applications, and data are stored as files making it much easier to copy a virtual machine from one site to another. The virtual machine can be reliably powered on and used at the disaster recovery site even if the underlying physical hardware is different from the production site.

However, storage is needed just the same at a disaster recovery site to hold the virtual machine files. Many of these files can be quite large ranging from just a few gigabytes (GB) in size up to hundreds of GB or even a few terabytes (TB). Storage capacity requirements can add up quickly depending on the number and size of the applications and databases that need to be protected. Performance is also a factor that must be considered as it has a direct impact on the reliability of a recovery and the recovery time. It is possible that an organization will have to run business-critical applications at a disaster recovery site for several days or weeks. Storage that does not meet ongoing business demands can impact productivity and revenue. In summary, storage provisioned at the disaster recovery site must meet capacity and performance requirements at a reasonable cost.

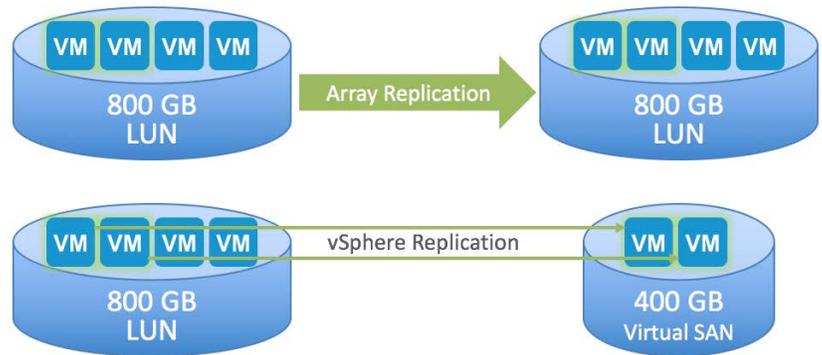
Why Virtual SAN for Disaster Recovery?

VMware Virtual SAN™ is VMware's radically simple storage solution for hyper-converged infrastructure (HCI). Virtual SAN and VMware vSphere® provide a complete, natively integrated platform consisting of compute, network, and storage resources for a wide variety of use cases including disaster recovery. Deploy on inexpensive industry-standard x86 server components to remove large, upfront investments. Since disks internal to the vSphere hosts are used to create a Virtual SAN datastore, there is no dependency on external shared storage hardware. This helps reduce the total cost of the solution while providing sufficient capacity, reliability, and performance.

Virtual SAN is built on an optimized I/O data path in the vSphere hypervisor for exceptional performance. It is managed as a core component of a vSphere environment meaning separate administration tools and connections are not required. This simplifies management particularly in locations that have little or no local IT staff such as a disaster recovery site.

VMware vSphere Replication™ provides asynchronous virtual machine replication with recovery point objectives (RPOs) as low as five minutes. Replication is configured on a per-virtual machine basis enabling precise control over which workloads are protected. This approach avoids the need to provision excess capacity at a disaster recovery site to accommodate an all-or-nothing replication solution. Furthermore, there is no requirement to have the same type of storage at both sites enabling more deployment options.

As an example, consider four 200GB virtual machines on a single LUN at the production site. Disaster recovery protection is needed only for two of the virtual machines. With array replication, the entire LUN (all virtual machines) is replicated. vSphere Replication could be used to replicate just the two virtual machines needing protection, which reduces capacity requirements at the disaster recovery site and wide area network (WAN) bandwidth consumption.



Virtual machine-centric storage policies can be created and assigned for various workload types. Policies are based upon the availability and performance services provided by Virtual SAN. These policies can be modified and reassigned, as needed, with no downtime. vSphere Replication supports storage policies. When configuring replication, a storage policy is selected and the configured storage policy is automatically assigned to the virtual machine when it is recovered.

Select Target Location

Select a datastore where the replicated files will be stored.

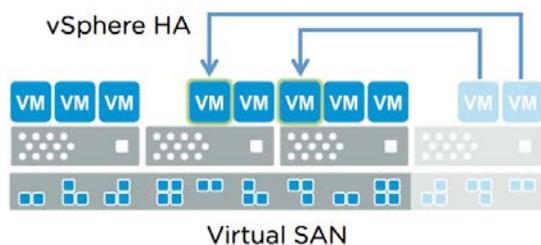
Filter datastores by name:

VM Storage Policy:

Name	Capacity	Provisioned	Free
Compatible			
 vsanDatastore	12.06 TB	16.76 TB	8.51 TB
Incompatible			
 hdd01-02	838.25 GB	977 MB	837.3 GB
 hdd01-03	838.25 GB	977 MB	837.3 GB

The use of local disks without Virtual SAN introduces risk to application uptime. For example, only one copy of a virtual machine's files is stored on a local disk. If that disk fails, the virtual machine files must be restored from backup media, which is time consuming and, in some cases, unreliable. It is possible to create a second copy of virtual machine files on another disk, but the process is not automatic and must be performed frequently to minimize data loss. The recovery of a second copy would also be a manual process in further increasing risk and recovery time.

Virtual SAN addresses these challenges by aggregating local disks into a shared datastore distributed across hosts in the cluster. Virtual SAN features a storage policy rule called "Number of failures to tolerate" or "FTT", which defines the number of copies of a virtual machine's files to distribute across the physical nodes in the cluster. The formula for determining the minimum number of hosts required to support an FTT rule is $2n+1$. For example, five hosts are required for $FTT=2$.



Since services such as VMware vSphere vMotion®, VMware vSphere High Availability, and VMware vSphere Fault Tolerance can be utilized at the disaster recovery site to protect workloads failed over from the production site. This is especially important if an organization needs to run production workloads for an extended period of time at a disaster recovery site as these services help minimize planned and unplanned downtime.

A variety of data protection solutions are available to back up and recover virtual machines and applications in a Virtual SAN cluster including VMware vSphere Data Protection™. Workloads can be backed up by vSphere Data Protection at the production site. This backup data can then be replicated to a vSphere Data Protection appliance at a disaster recovery site. Having the backup data at the production and disaster recovery sites facilitates recovery from a variety of disaster scenarios.

vSphere Data Protection features efficient, variable-length deduplication and compression to help minimize capacity consumption on Virtual SAN and other datastore types. Backup data replication is also very efficient as only unique data segments after deduplication has been applied are replicated. This minimizes the amount of bandwidth consumption. The replication stream is encrypted for secure transmission.



vSphere Data Protection and vSphere Replication can be used in the same environment. Depending on business requirements, some virtual machines could be protected from disaster by vSphere Replication and others by vSphere Data Protection. This approach provides flexibility in recovery times and capacity consumption at the disaster recovery site. For example, Tier-1 workloads can be replicated with vSphere Replication, which offers faster recovery times than restore from backup, but consumes more capacity at the disaster recovery site. Tier-2 workloads can be backed up by vSphere Data Protection locally and the backup data replicated to the disaster recovery site. It will take longer to restore a virtual machine from backup data, but the backup data will likely consume less storage capacity due to vSphere Data Protection deduplication and compression. As a footnote, the capacity consumed by vSphere Replication replicas can be reduced using Virtual SAN deduplication and compression in all-flash configurations.

Automation with Site Recovery Manager

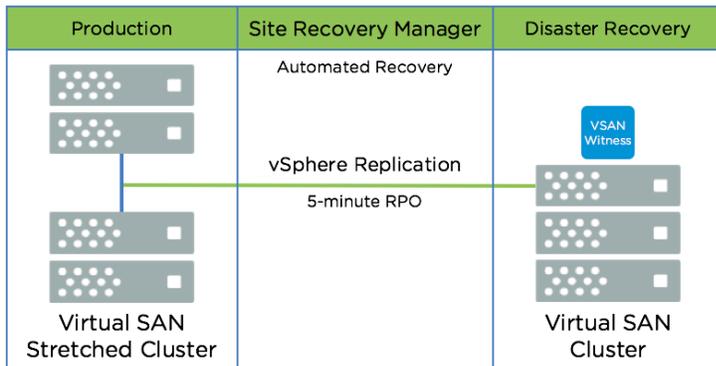
VMware Site Recovery Manager™ can be utilized with Virtual SAN and vSphere Replication to orchestrate the recovery of multiple virtual machines. Automation further reduces recovery times and minimizes risk by eliminating manual, error-prone processes. Site Recovery Manager includes the ability to precisely control the startup order of virtual machines and it automates IP address changes when virtual machines are failed over. Testing recovery plans with Site Recovery Manager is non-disruptive, which enables frequent testing. Frequent testing leads to higher levels of confidence that recovery will work as planned when needed. History reports are generated with every test and failover event providing documentation to satisfy organization and regulatory requirements.

Recovery Plan History Report VMware Site Recovery Manager

Plan Summary	
Name:	oracle12c
Description:	
Protected Site:	wdcpod06vm01.pml.local
Recovery Site:	pod02vm01.pml.local

Run Summary	
Operation:	Test
Storage Options:	Synchronize storage when plan runs
Started By:	VSPHERE.LOCAL\Administrator
Start Time:	2016-03-14 14:14:12 (UTC 0)
End Time:	2016-03-14 14:16:23 (UTC 0)
Elapsed Time:	00:02:11
Result:	Success

Virtual SAN Stretched Clusters and Site Recovery Manager



For higher levels of resiliency across three sites, consider the use of a Virtual SAN stretched cluster with Site Recovery Manager. For example, two production locations 100 kilometers apart could each house one half of a stretched cluster to protect against the failure of either location. A third location farther away hosts a second Virtual SAN cluster to supply compute, storage, and network resources for recovered virtual machines, as well as, any workloads that run on a regular basis at the disaster recovery site.

A Virtual SAN stretched cluster requires a “witness”, which is vSphere running in a virtual machine. The witness serves as a tie-breaker in certain situations such as loss of network connectivity between the two locations that make up a stretched cluster. The witness cannot be located within the same site as the stretched cluster so the disaster recovery site is the natural place to host this virtual machine appliance. Other workloads running at a disaster recovery site might include test and development, virtual desktops, email, directory services, and DNS.

Since stretched clusters essentially utilize synchronous replication between the two locations, an RPO of zero is achieved. That means no loss of data if one of the locations in the stretched cluster is offline. vSphere HA automates the recovery of virtual machines affected by an outage at either location in the stretched cluster. Recovery time for these virtual machines is typically measured in minutes.

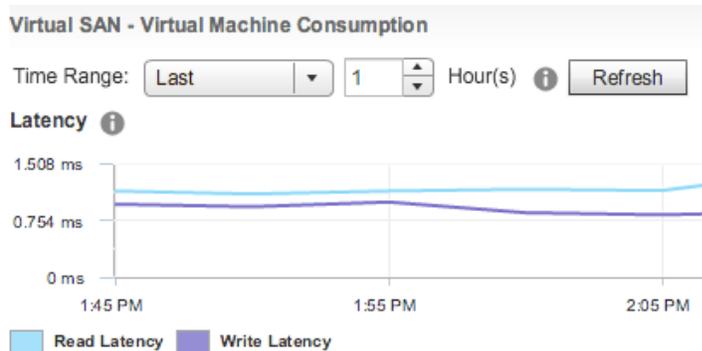
Replication from the stretched cluster to the disaster recovery site is facilitated by vSphere Replication. As mentioned previously, per-virtual machine RPOs for replication between two Virtual SAN datastores can be as low as five minutes. Site Recovery Manager automates the failover and fail-back processes between the stretched cluster and the disaster recovery site.

Virtual SAN Performance

Virtual SAN is uniquely embedded in the vSphere hypervisor kernel. It is able to deliver the highest levels of performance without taxing the CPU or consuming high amounts of memory resources, as compared to other solutions requiring storage virtual machine appliances that run separately on top of the hypervisor. An all-flash Virtual SAN configuration will naturally provide the highest performance, which translates to lower recovery times. A combination of magnetic and solid state disks is another option to enable a flash-accelerated hybrid architecture for a disaster recovery site. Virtual SAN provides exceptional performance with numerous configuration options.

Built-in Health Checks and Performance Statistics

Virtual SAN includes a health check feature to monitor items such as network connectivity, cluster health, disk capacity, component metadata, and compliance with the VMware Compatibility Guide. These configuration items are automatically checked once per hour by default. If a problem is detected, an alert is raised proactively informing administrators of the issue.



Multiple performance metrics are available in the vSphere Web Client including throughput, IOPS, and latency. It is possible to look at data on various levels such as the Virtual SAN datastore, host, virtual machine, and virtual disk (VMDK) levels. The time range for these metrics can be set to the last x number of hours or a custom date and time range.

Summary

VMware vCenter Server™, vSphere, and Virtual SAN collectively create the best platform for running and managing virtual machine workloads requiring predictable performance and rapid recovery in the event of a disaster. The integration of Virtual SAN with vSphere simplifies administration through storage policy-based management. Business-critical workloads such as web sites, e-commerce applications, databases, employee remote access, and communications can benefit from shared storage without the cost and complexity of dedicated storage hardware. vSphere includes availability features such as vSphere HA, vSphere Replication, and vSphere Data Protection to minimize unplanned downtime. Site Recovery Manager can automate virtual machine migration and disaster recovery through tight integration with vSphere Replication. This includes precise virtual machine startup orders, IP address changes, and the generation of history report documentation for testing, failover, and fail-back operations. The health and performance levels of a Virtual SAN datastore are constantly monitored to lower risk before, during, and after a disaster recovery. If more capacity is needed, it is simple to add using a scale up or scale out approach without incurring downtime.

Learn More

[Virtual SAN](#)

[vSphere Replication](#)

[Site Recovery Manager](#)

[vSphere Data Protection](#)

[Virtual Blocks Blog](#)

[Customer Stories](#)

