VMware vSAN for Disaster Recovery

Lowering the Risk and Cost of DR

"The new design of IT at the university starts from a fundamental choice for the modernization of data centers. It is a choice that has seen in the virtualization of systems a fundamental piece and that is based, for all the most critical server and database systems, also, on VMware technology."

MAURIZIO DAVINI
CTO UNIVERSITY OF PISA

Solution Summary

- VMware vSAN reduces storage costs and complexity
- VMware vSphere supercharges performance and enhances operational efficiency
- VMware vCenter provides a centralized platform for controlling vSphere environments
- VMware vSphere Replication enables per-VM asynchronous replication
- VMware Site Recovery Manager automates orchestration of failover and failback to minimize downtime

The United States Federal Emergency Management Agency (FEMA) estimates that "about 25 percent of businesses do not reopen after disasters." Even if organizations manage to reopen after a disaster, they are susceptible to significant losses, as shown in numerous online news articles.

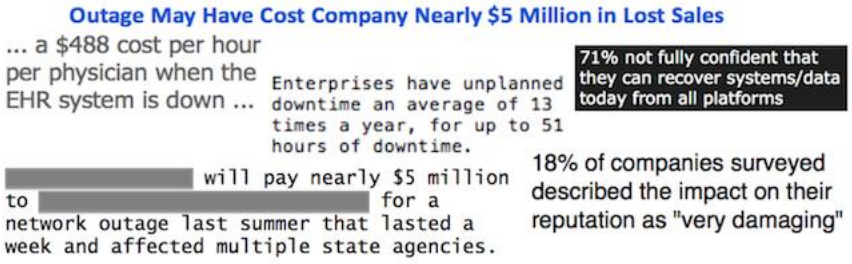

Many businesses are considering or using cloud-based disaster recovery as a service (DRaaS) offerings such as VMware Cloud Disaster Recovery, but some are keeping disaster recovery sites on-premises. This might be because of organizational preference or regulatory requirements. The expense of a disaster recovery site might be cost-prohibitive for some organizations. As a result, many 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 infrastructure, including server hardware, storage, and replication software. It was common in the past for organizations to deploy similar systems at production and disaster recovery sites. This expensive approach was made to help prevent failure in recovering applications and data due to differences in the infrastructure.

Virtualization removes the need for the same infrastructure hardware at 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 differs 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 must be protected. Performance is also a factor that must be considered, as it directly impacts the reliability of recovery and the recovery time. An organization may 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 vSAN for Disaster Recovery?

VMware vSAN provides a cost-effective storage solution that is fast, reliable, and resilient to run business critical workloads in the event of a disaster recovery.

VMware vSAN™ is VMware's enterprise-class storage solution for hyperconverged infrastructure (HCI). vSAN and VMware vSphere® provide a complete, natively integrated platform consisting of compute, network, and storage resources for various use cases, including disaster recovery. Deploy inexpensive industry-standard x86 server components to remove significant, upfront investments. Since disks internal to the vSphere hosts are used to create a vSAN 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.

vSAN is built on an optimized I/O data path 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 with little or no local IT staff, such as a disaster recovery site.

Per-VM Protection with vSphere Replication

VMware vSphere Replication™ is a feature included with vSphere Essentials Plus and higher editions. It 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 providing 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.

VMware vSphere Replication provides efficient per-VM asynchronous replication with RPOs as low as five minutes.

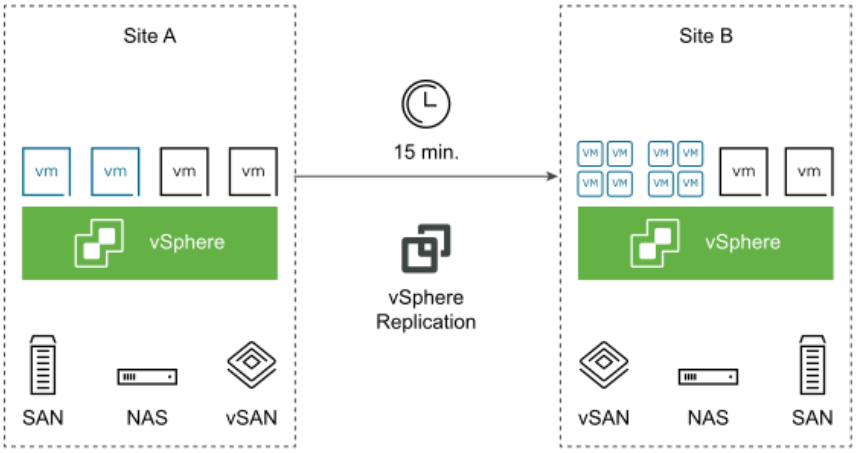

For example, consider four 200GB virtual machines on a single LUN at the production site. Disaster recovery protection is needed only for two of the virtual machines. The entire LUN (all virtual machine data) is replicated with array replication. vSphere Replication can replicate just the two virtual machines needing protection, which reduces capacity requirements at the disaster recovery site and wide area network (WAN) bandwidth consumption.

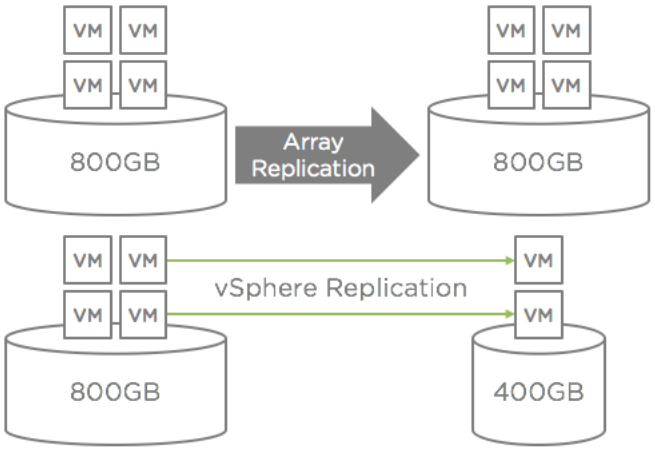

Virtual machine-centric storage policies can be created and assigned for various workload types. Policies are based on the availability and performance services provided by vSAN. These policies can be modified and reassigned, as needed, with no downtime. vSphere Replication supports storage policies.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

Local Disks Are Not Enough

The use of local disks without vSAN 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 sometimes 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, further increasing risk and recovery time.

vSAN addresses these challenges by aggregating local disks into a shared datastore distributed across hosts in the cluster. vSAN features a storage policy rule called "Failures to tolerate," which defines how many hosts can be offline before data becomes inaccessible. Data is distributed across multiple hosts to achieve this resilience. For example, if "Failures to tolerate" is set to "1 failure – RAID-5 (Erasure Coding)", a host can be offline, and all VMs with that policy assigned will remain accessible.

vSAN Enables vSphere Cluster Features

With shared storage, cluster services such as VMware vSphere vMotion® and VMware vSphere High Availability™ (vSphere HA) can be utilized at the disaster recovery site to protect recovered workloads. This is especially important if an organization needs to run production workloads for an extended period at a disaster recovery site, as these services help minimize planned and unplanned downtime.

vSphere cluster features such as vMotion and vSphere HA require shared storage. vSAN provides that shared storage.

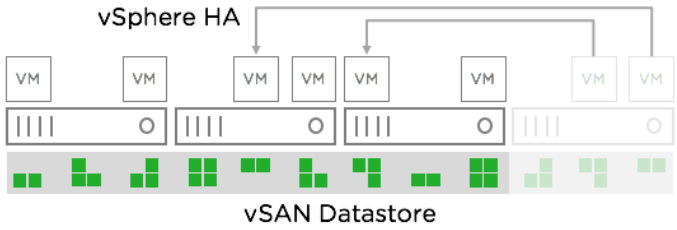

Various data protection solutions are available to back up and recover virtual machines and applications in a vSAN cluster. Many of these solutions include the capability to replicate backup data to a disaster recovery site. The backup data at the production and disaster recovery sites facilitates recovery from various downtime scenarios.

Depending on business requirements, some virtual machines could be protected from disaster by vSphere Replication and others by replicating backup data to a remote site. 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 restores from backup. Tier-2 workloads can be backed up locally, and the backup data can be copied to the disaster recovery site. It will take longer to restore a virtual machine from backup data. Still, the backup data will consume less storage capacity due to various deduplication and compression features built into some data protection solutions. Storage

capacity consumed by vSphere Replication replicas can be reduced using vSAN deduplication and compression.

Automation with Site Recovery Manager

VMware Site Recovery Manager™ can be utilized with vSAN and vSphere Replication to orchestrate the recovery of multiple virtual machines. Automation reduces recovery times and minimizes risk by eliminating manual, error-prone processes. Site Recovery Manager includes the ability to control the startup order of virtual machines precisely, 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 increases 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.

vSAN Stretched Clusters and Site Recovery Manager

For higher levels of resiliency across three sites, consider using a vSAN stretched cluster with Site Recovery Manager. For example, two production locations 100 kilometers (about 62.14 mi) apart could each house one-half of a stretched cluster to protect against the failure of either site. A third location farther away hosts a second vSAN cluster to supply compute, storage, and network resources for recovered virtual machines and any workloads that run regularly at the disaster recovery site.

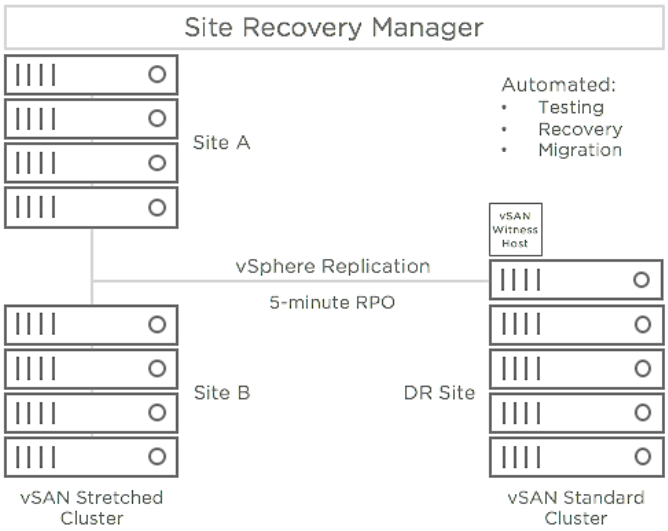

A vSAN stretched cluster requires a “witness host,” which is vSphere running on a virtual machine. The witness host serves as a tiebreaker in certain situations, such as losing network connectivity between the two locations comprising a stretched cluster. The witness host cannot be located within the 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 utilize synchronous replication between the two locations, an RPO of zero is achieved. That means no data loss 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.

vSphere Replication facilitates replication from the stretched cluster to the disaster recovery site. As mentioned, per-virtual machine RPOs for replication between two vSAN 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

vSAN Performance

vSAN is uniquely embedded in the vSphere hypervisor kernel. It can deliver the highest performance levels without taxing the CPU or consuming considerable amounts of memory resources, compared to other solutions requiring storage virtual machine appliances that run separately on top of the hypervisor. An all-flash vSAN configuration will naturally provide the highest performance, leading to lower recovery times.

The latest configuration offering, vSAN Express Storage Architecture (ESA), offers higher performance and space efficiency levels. These are achieved through a new log-structured file system and by more effectively utilizing NVMe devices. More details on vSAN ESA can be found here.

Summary

HCI with VMware vCenter Server™, vSphere, and vSAN collectively create the best platform for running and managing virtual machine workloads requiring predictable performance and rapid recovery during a disaster. The integration of vSAN with vSphere simplifies administration through storage policy-based management. Business-critical workloads such as websites, 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 and vSphere Replication to minimize unplanned downtime. Site Recovery Manager automates 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 failback operations. The health and performance levels of a vSAN datastore are constantly monitored to lower risk before, during, and after a disaster recovery. It is simple to add capacity using a scale-up or scale-out approach without downtime if more is needed.

Resources

- VMware vSAN Product Page
- VMware vSphere Product Page
- VMware vCenter Product Page
- VMware Site Recovery Manager Product Page
- VMware vSphere Replication Feature Page

Item No: vmw-sol-brief-temp-uslet-word-2021 8/21