VMware Virtual SAN

Mixed Workloads in HCI

Eliminate the Complexity of Multiple Infrastructure Silos

A common approach to building out compute and storage infrastructure for varying workloads has been dedicated resources based on the applications that run in each environment. For example, architects often dedicate server and storage hardware to server environments and a completely separate infrastructure to virtual desktops. In larger environments, it is not uncommon to see hardware stacks dedicated to specific applications and services. This approach is typically a result of application owners requiring dedicated hardware to help ensure performance and availability service level agreements (SLAs) are met.

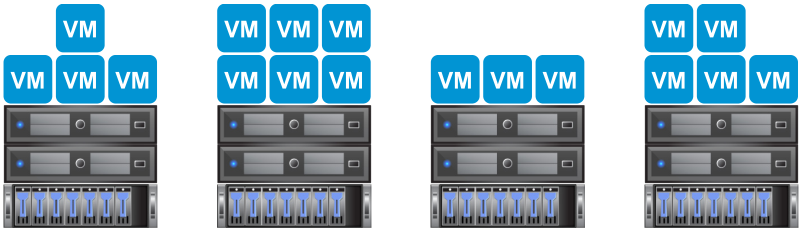

Dedicated infrastructure hardware for multiple environments and applications also leads to complexity, wasted resources, and higher costs. It is common for an application owner to budget and procure a certain amount of compute and storage capacity based on software vendor recommendations only to find that the application actually requires much less. The result is excess server and storage resources with low utilization. In other words, excessive capital was spent on hardware and it is difficult to subdivide and provision resources to other application owners and projects.

Hyper-converged infrastructure (HCI) makes it easier to consolidate dedicated hardware silos, reduce complexity, provision the right levels of compute and storage resources, and reduce capital and operational expenditures. However, HCI hardware is only part of the equation. The solution must also have the right software to manage resources, allocate capacity, and run a variety of applications while ensuring performance and availability requirements for every workload are met. Resources must be organized into flexible, easily-consumable pools that are governed by policies and usable by any kind of workload.

VMware Hyper-Converged Software

VMware Hyper-Converged Software (HCS) is VMware's premier solution for HCI. The natively integrated software combines radically simple VMware Virtual SAN™ storage, the marketing-leading VMware vSphere® hypervisor, and the vCenter Server unified management solution with the broadest and deepest set of HCI deployment choice. VMware HCS converges all datacenter functions as software running on industry-standard x86 building blocks, enabling a modular, building-block architecture. This not only provides scale-out capabilities at a fraction of the cost of traditional, siloed infrastructure, but also simplifies operations by enabling automation of all datacenter functions.

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. 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. Provisioning is dramatically faster - compute and storage are provisioned together on a per-virtual machine basis with a few clicks.

Storage Policy-Based Management

Instead of dedicating storage arrays to specific environments and workloads, virtual machine-centric storage policies are created and assigned for various workload types. A storage policy consists of one or more rules that define the levels of availability and performance provided by Virtual SAN. For example, a policy can be created that stripes data across four disks for improved read performance and utilizes RAID-6 erasure coding to ensure data loss is avoided if there is a simultaneous loss of two disks or hosts.

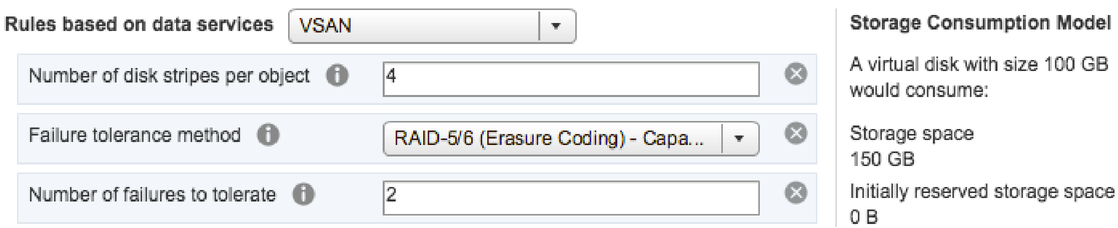

Multiple policies with a variety of rule combinations can be configured in a single environment. Policies are assigned to virtual machines or even individual virtual disks, which provides precise control of performance and availability on a per-VM basis. Virtual SAN also supports a policy rule that guarantees a virtual machine will not run out of storage capacity by implementing an object space reservation. IOPS limits can also be imposed to avoid “noisy neighbor” performance impacts on other virtual machines in the Virtual SAN datastore. Storage policies can be assigned to new and existing virtual machines. Policies can be reassigned and modified with no downtime. The ability to configure multiple storage policies and assign them to individual virtual machines in a single Virtual SAN datastore eliminates the need to build separate storage infrastructures and manage complex LUN and volume configurations for various workloads.

Virtual SAN Resiliency

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.

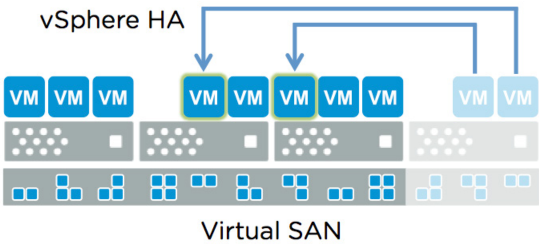

Since services such as VMware vSphere vMotion®, VMware vSphere High Availability, and VMware vSphere Fault Tolerance can be utilized at the disaster recovery site to protect workloads failed over from the production site. This is especially important if an organization needs to run production workloads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

With these availability services natively integrated into the platform, it is easy to provide specific levels of resiliency to each individual workload within the same environment.

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. All-flash Virtual SAN configurations will naturally provide the highest performance with predictable low latencies. A combination of magnetic and solid state disks is another option to enable flash-accelerated hybrid architectures. Virtual SAN provides exceptional performance with numerous configuration options, which are managed by per-virtual machine policies all on a single datastore.

Compute Resource Management

Workloads vary widely when it comes to CPU and memory requirements. With vSphere, granular provisioning and management of these resources can be accomplished with just a few clicks. When a virtual machine is provisioned, virtual CPU and memory resources are configured for the virtual machine. In most cases, the underlying physical CPUs and memory are shared across the virtual machines running on the host. It is possible to oversubscribe these resources to increase utilization and minimize the amount of physical infrastructure required to run the workloads.

Some workloads might require dedicated amounts of CPU cycles and memory to help ensure application performance and availability. One possible solution is to configure reservations. For example, it is possible to guarantee a virtual machine will have exclusive access to 1000 MHz of CPU capacity and 2GB of memory.

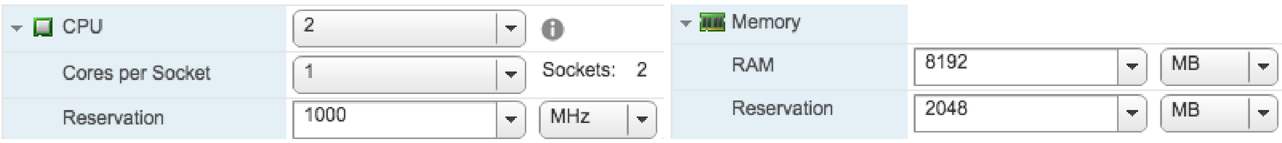

Another option, which is preferred, is the use of “shares”. Shares basically determine the relative importance of a virtual machine when there is contention for a resource. For example, consider a case where there are two virtual machines running on a host and Virtual Machine 1 has twice as many shares configured for CPU. When there is no contention for CPU on the host, both virtual machines are able to consume any amount of CPU cycles, as needed. In a situation where there is contention for shared CPU resources, Virtual Machine 1 with twice the number of shares will have access to the physical CPUs twice as often as Virtual Machine 2. This approach effectively enables prioritization at the virtual machine component level.

Reservation	0	MHz
Limit	Unlimited	MHz
Shares	High	4000

Virtual Machine 1

Reservation	0	MHz
Limit	Unlimited	MHz
Shares	Normal	2000

Virtual Machine 2

Similar to IOPS limits with Virtual SAN, it is also possible to limit the amount of resources a virtual machine can consume with vSphere. While limits should be used sparingly, they can be an effective means to limit the impact a “noisy neighbor” virtual machine has on other virtual machines running on the same host.

Resource Pools

Managing resources at the virtual machine level is fairly simple in smaller environments, but becomes more cumbersome for the larger enterprise. That is why vCenter Server includes the ability to configure resource pools. Administrators can create pools of resources at the cluster level with shares, reservations, and limits. Virtual machines added to the resource pool are governed by the resource pool configuration, which is easier than configuring resource settings on a per-virtual machine basis. Multiple resource pools can be used to manage an environment with mixed workloads ensuring proper resource allocation and prioritization.

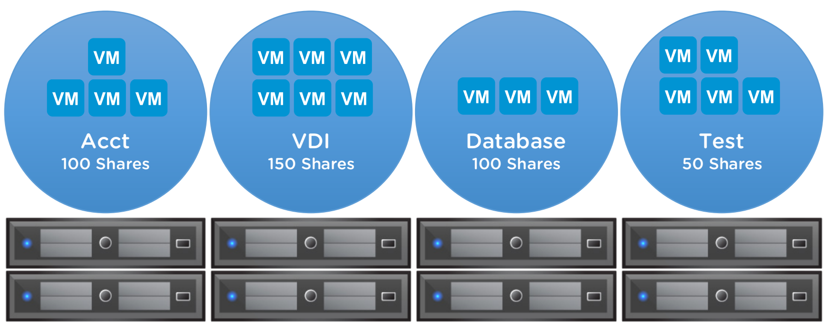

Virtual machines can be moved from one resource pool to another in cases where resource demands and priorities change. Consider the example of accounting applications processing year-end data. The virtual machines running the accounting applications could be moved to a resource pool with more CPU shares to help ensure the processing is completed on time.

vmware®

Virtual Network Management

VMware vSphere Distributed Switch™ includes Network I/O Control to provide precise management of networking resources using shares, reservations, and limits. Similar to compute resource shares, network services can be prioritized to ensure specific traffic types such as Virtual SAN and virtual machine traffic are not impacted when there is contention for available bandwidth. It is also possible to configure precise bandwidth reservations and limits for these services.

Traffic Type	Shares	Shares Value 1 ▼	Reservation	Limit
Virtual SAN Traffic	High	100	5,000 Mbit/s	Unlimited
Fault Tolerance (FT) Traffic	Normal	50	0 Mbit/s	Unlimited
vMotion Traffic	Normal	50	0 Mbit/s	Unlimited
Virtual Machine Traffic	Normal	50	1,000 Mbit/s	Unlimited
Management Traffic	Normal	50	0 Mbit/s	Unlimited

Traffic shaping is also available with vSphere virtual switches, which enables administrators to place limits on average bandwidth, peak bandwidth, and burst size.

Summary

VMware HCS is the premier solution for HCI featuring precise controls for compute, network, and storage resource administration. Allocate the exact amount of resources required for an application and easily adjust them without downtime as resource demand changes. Utilize local storage devices to create a resilient storage platform with extreme performance with Virtual SAN. The integration of Virtual SAN and vSphere simplifies administration through policy-based management. Build highly available clusters using vCenter Server and vSphere with standard x86 hardware to manage CPU and memory as pools of compute capacity. A variety of mixed workloads such as web sites, applications, databases, virtual desktops, file services, and communications can benefit from a single HCI platform including shared storage with VMware HCS.

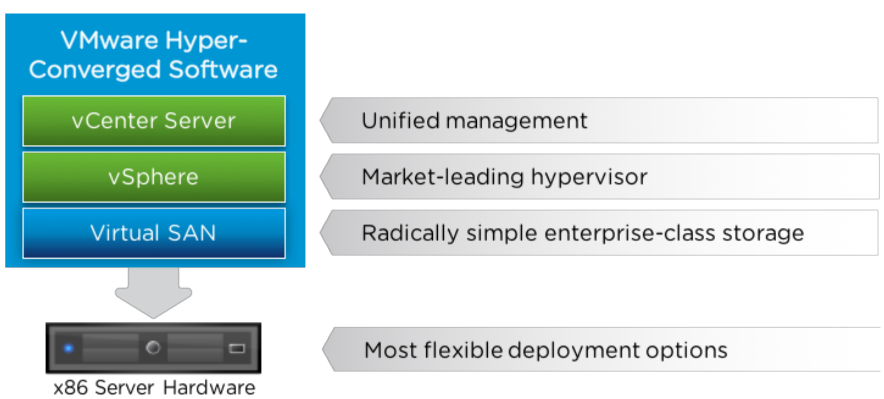

Learn More

VMware Hyper-Converged Software

Virtual SAN

Virtual Blocks Blog

Customer Stories

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