

StorMagic's SvSAN – Targets HCI's ROBO Sweet Spot

Over the past ten years hyperconverged infrastructure (HCI) has gone from leading edge technology to become the fastest growing method for building an IT infrastructure. As HCI acceptance has grown, some vendors have targeted the enterprise data center adding functionality like data deduplication and erasure coding that are only effective in larger clusters.

As Robert A. Heinlein famously wrote *There Ain't No Such Thing as A Free Lunch* so while data deduplication can let a system store more data on a given amount of flash, that magic requires CPU and memory resources that are in short supply in the ROBO and SMB environments that are a natural fit for hyperconvergence.

One HCI pioneer, StorMagic, has remained focused on these use cases concentrating on the challenges of servicing these smaller, if not less critical, environments. StorMagic is focused on providing reliable storage in remote environments where communications lines may not offer five nines reliability, and it may take four days to replace a failed hard drive or SSD, not the four hours that's the standard for enterprise IT.

In a classic “do more with less” catch-22, solutions for remote offices have to provide this higher level of resiliency with a limited budget and make it easy enough to use that the store or plant manager can replace that failed hard drive with a little help from central IT on the phone.

StorMagic's SvSAN operates through a lightweight VSA that runs under VMware's ESXi or Microsoft Windows Hyper-V and can consume as little as 1 vCPU, 1GB of RAM and 20.5GB of storage, significantly less than the 8 vCPUs and up to 64GB of RAM used by a leading enterprise HCI vendor with all their features enabled. This means smaller offices can use single-socket servers with Xeon E3 or Xeon-D processors.

Using single socket servers saves not just the extra cost of enterprise-class dual socket servers but also the cost of the software for each host that's licensed by the CPU socket including ESXi and Windows that can easily add up to thousands of dollars per host.

A typical SvSAN site has two servers running virtual machines including the SvSAN VSA which can be allocated 2-32GB of RAM to use as a high performance read cache. Storage capacity is provided by an LSI or OEM RAID controller and off-the-shelf drives. The SvSAN VSA synchronously replicates data between the nodes and publishes iSCSI LUNs the hypervisor can then use to house VMs. Read I/Os are satisfied by the VSA that's local to the VM making the request minimizing network traffic and latency.

SvSAN consumes storage resources from Hyper-V or ESXi allowing administrators the flexibility to create all-flash, hybrid and all-HDD data stores without dedicating entire physical devices to those pools.

In the recently released version 6.1 of SvSAN, StorMagic boosts performance by using additional RAM allocated to the VSA as a read cache. SSD handling has evolved from using the SSD as a read cache to a predictive auto-tiering mechanism that promotes data proactively into the SSD read/write cache. In addition to relatively common

techniques like reading ahead of sequential workloads, SvSAN has a function that will identify and pin data from events like VDI boot storms to ensure the VDI boot drives are in cache for tomorrow morning.

A typical HCI solution stores one copy of each data object on each node which requires a minimum three node cluster to retain a minimal level of resiliency during those four days while a replacement drive is Fedexed from headquarters. By using a local RAID controller, which typically costs \$500-700, SvSAN can survive multiple failures and continue serving up VMs via iSCSI.

Two-node clusters present the special challenge of resolving possible split-brain situations when a node can no longer communicate with its partner. Some other HCI solutions require three nodes to resolve the conflict.

Rather than insist on three nodes, StorMagic provides the Neutral Storage Host which acts as a witness. When a node can't communicate with its partner, the NSH acts as a tiebreaker allowing the surviving node to maintain data services without the concern of split brain or corruption. Of course, as long as the two nodes can communicate with each other they maintain a quorum and will operate normally even if their connection to the witness is lost. The system can even continue to serve data after the loss of both a server and the NSH in some scenarios.

Customers have a great deal of flexibility in implementing the NSH from a centralized witness that can support hundreds of multiple remote locations: NSHs running as virtual machines, Windows services, a Linux service and even on a Raspberry Pi. So a branch office with just two servers can have an external witness at an affordable price.

StorMagic has also paid attention to the management and deployment issues that come with many remote offices. The IT team at *\$1video* can run the StorMagic VSA installation wizard to install their first cluster and at the same time create a PowerShell script to automate configuring the next 700 units to a baseline config. They can then set IP addresses and other site specific details with a second Powershell script that even a store manager could run.

StorMagic licenses SvSAN by capacity in 2, 6 and 12TB tiers plus an unlimited license.

While some HCI vendors have dreams of datacenter domination StorMagic has concentrated on the places where HCI has an inherent advantage remote offices and other small IT installations like SMBs. Their lightweight VSA, and even lighter weight NSH witness, help keep hardware costs down while the RAM and SSD enhancements in version 6.1 deliver the performance to your remote users happy.