

SFD6 - StorMagic

Wednesday PM

Co-founders

Hans O'Sullivan - CEO

Chris Farey - CTO

\*Hans

Founded StorMagic after Eurologic and Elipsan (both sold to Adaptec)

Founded 2006

Launched SvSAN in 2009

StorMagic refocuses on multi-site customers - 2011

SvSAN v5 released - 2012

Rebrands in line with growth - 2014

Trend to centralise - solving problems, you create others.

Not all apps run well in "the centre" - performance, availability, compliance

Distributed enterprise

- virtualising remote infrastructure
- introducing new remote services
- systems in harsh environments
- no local IT
- seeking to reduce support costs
- experiencing downtime of critical, remote apps

The remote site

- 2TB avg data capacity
- 7 or 8 key apps
- separation of servers on site
- no computer room
- set and forget mentality

Need

- HA
- Centralised Management
- Small IT footprint
- Simple, automated deployment

Traditional storage does not fit at the remote site

- SPoF
- Complexity
- Specialist staff
- depreciate in value
- high capex and opex
- tied in to hardware vendor

Data protection is also a problem - but they don't solve that problem yet, although

they're working on some enabling technology for the future

The focus on the edge is currently a marketing consideration - there's a lot less focus there than in the DC

Business took off in 2011 when VMware released their VSA.

VSAN?

Doesn't scale down to the edge

Scale out design - can't control data locality

3 node requirement

Mix of hdd and ssd

vSphere only

SvSAN Overview

Virtual SAN - shared storage using internal / DAS, synchronous mirroring between nodes

Flexible - runs as VSA independent of storage hardware

HA - withstands server or storage failure

Scalable - 2 nodes to many nodes

Support for vVOLs is being worked on

57PB being managed. 1226 customers, 30000 licenses, 2200 sites. Mainly Retail, Energy, Manufacturing, Govt.

\*Real-life use cases\*

Retail - inventory control systems, customer and staff management, PoS

Govt - diplomatic communication platforms

Defense - battlefield control systems

Manufacturing - process control

Fin services - customer transactions

Restaurants - reservation and online order systems

Transportation - vehicle positioning and monitoring

Energy Production

Medical - PACS

\*Hans

US Retailer

Big box US retailer

2000+ sites

2+ servers per site

100 store outages per year (avg 6 hrs per outage) - keep in mind these are low-cost storage devices

Solution is truly active / active - can read from both parts of the mirror at the same time

Working on SSD write-back caching, working on read caching for a future release

E-ON energy - renewable energy  
100+ sites  
2+ servers per site

Safekick - oil rig monitoring software  
Real-time monitoring and management of drilling operations

German Army - Bundeswehr  
Requirements:

- little or no space in the field
- free of excess weight
- portable
- drop anywhere in the world
- easy to deploy and manage
- offer high availability

Licensed per capacity and per site  
Entry level is 2 servers and 2TB per site - \$2000 List - pricing per server pair  
After 16TB go unlimited (\$10K)  
Largest capacity installation is about 200TB

If you add a 3rd server, it's another \$1000

\*Chris

SvSAN stack provides storage services  
Based on Linux

Currently supports vSphere and Hyper-V  
Any storage (HDD, SSD, RAID, Flash)  
Presents as virtual disks - virtual iSCSI SAN

Writes are committed to both sides of the mirror before acknowledged (mirrored on disk)

A lot of CPU cycles get spent on the TCP stack (using software iSCSI) - add cores to help with this.

Video demos

Backup the VMs, not the VSA (you can take a backup of the VSA config separately)

Technical Background

Layered Stack Architecture

iSCSI Target  
SCSI Target  
Sync Mirroring  
SSD Caching

Migrator  
Remote and Local I/O

Users create virtual disks from storage pools, then presented over iSCSI

They wrote their own iSCSI stack due to a number of shortcomings in existing offerings when they first started (2006)  
Can take advantage of TCP offloads, but not iSCSI offloads, etc.

Synchronous mirroring  
Mirror to any pair - with 3 VSA, A, B and C, you could have A-B, B-C and C-A  
Active / active - data can be read or written from either side  
Presented as multi-ported SCSI target device  
SCSI ALUA protocol reports on state  
Supports vSphere NMP and Windows MPIO DSM

Quorum requirement - tiebreaker is a software service - can run local or remote, windows or linux, there's even an ARM version - why not run it on a Raspberry Pi?  
For a VSA to come on-line, there must be a majority in agreement - 2 VSAs, 1 VSA and quorum service or all 3.  
Quorum service can act as a quorum for 1000s of virtual disks. If it goes down, as long as both sides of the mirror are up, you're okay.

Mirror resync  
Full only required for initial and hw replacement

Fast re-sync - changes tracked in a meta-data journal

Mgmt and monitoring  
Integration with vCenter, and MS System Center  
Scripting tools to automate deployment (useful for many remote sites)  
Monitor using standard tools

- SNMP
- MS SCOM (Management pack)

Usage stats - not configurable - distinguish between 2 halves of pairs, different hosts, different network ports  
Daily stats report I/O per minute (stored for 24 hours)  
Monthly stats - I/O per hour (stored for a month)  
Yearly stats report I/O per day

If you want these to last longer - you can manually download these every day and dump to excel (csv). This could be scripted, potentially.

Scalability Options  
Scale up - start small and grow, adding storage, CPU and RAM  
Scale out - add nodes and re-balance workloads, dynamic migration allows virtual disks to be migrated between VSAs, non-disruptive upgrades

Last section got canned by Hans - he's glad that they ran out of time to talk about the

competition

Centralised mgmt

Quorum mgmt

“You always know where your data is”