StorMagic

WHITE PAPER

StorMagic SvSAN
TECHNICAL OVERVIEW

OVERVIEW

StorMagic SvSAN is a hyperconverged storage solution designed to run on commodity x86 servers, with two servers being the minimum cluster size to provide high availability. It has been designed with lightweight system requirements to deliver highly available, shared storage with centralized management.

This white paper describes SvSAN's technical specifications, deployment options and use cases. For a broader outline of SvSAN including system requirements, compatibility tables and support levels, please refer to the **SvSAN Data Sheet**.

SvSAN eliminates the need for traditional physical SANs which are costly, complex and present a single point of failure. With high availability out-of-the-box, business-critical applications and IT services suffer no downtime and the removal of a SAN along with the low cost of commodity servers and storage lowers IT acquisition and operating cost.

The flexibility of SvSAN allows it to be deployed in a number of ways to suit different requirements. Deploy in a typical hyperconverged configuration for compute and storage or as a storage-only cluster, known as a server SAN.

Hyperconverged

The hyperconverged deployment model consolidates the compute, storage and hypervisor functions into a single server footprint, utilizing the server's internal storage to create a pool of highly available shared storage. This is the most common SvSAN deployment option, with thousands of installations (see fig. 1).

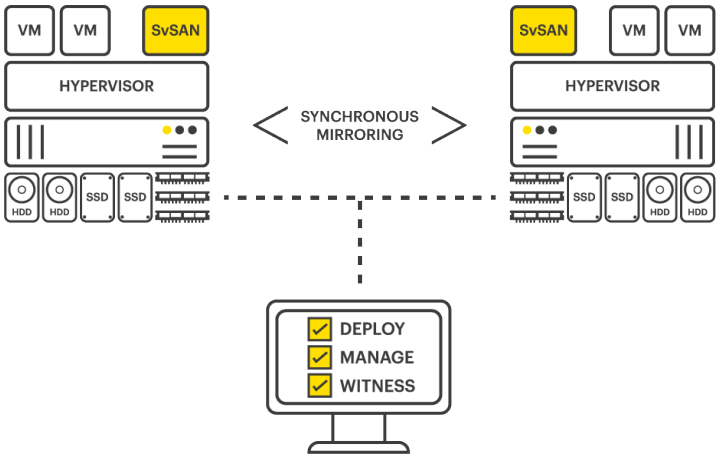

Fig. 1 *A 2-node hyperconverged deployment.*

In hyperconverged configurations, each server requires enough system resources (CPU and memory) for the applications and virtual machines as well as deliver storage and hypervisor features.

Hyperconverged configurations are direct replacements for traditional storage arrays and are ideally suited to environments that require a small footprint and high availability. These include, remote and branch office (ROBO) locations, edge computing environments, or internet-of-things (IoT) applications, where the data needs to be stored and processed close to where it is created or used by a customer.

Server SAN

A server SAN is a direct alternative to a traditional storage array, delivering shared storage and similar data services.

Like hyperconverged configurations, server SANs leverage internal (or external) server storage. However, the fundamental difference between the two is that the compute and storage functions are separated in a server SAN. This provides segregation between the compute and storage functions, enabling them to be scaled independently. However, this configuration requires more servers to deliver compute and storage components (see fig. 2).

A server SAN is a direct replacement for traditional storage and can therefore be used in similar environments where traditional storage has been prevalent. These include datacenters, small or medium-sized enterprises (SME) or environments that have non-virtualized workloads such as legacy operating systems, but require high availability and shared storage.

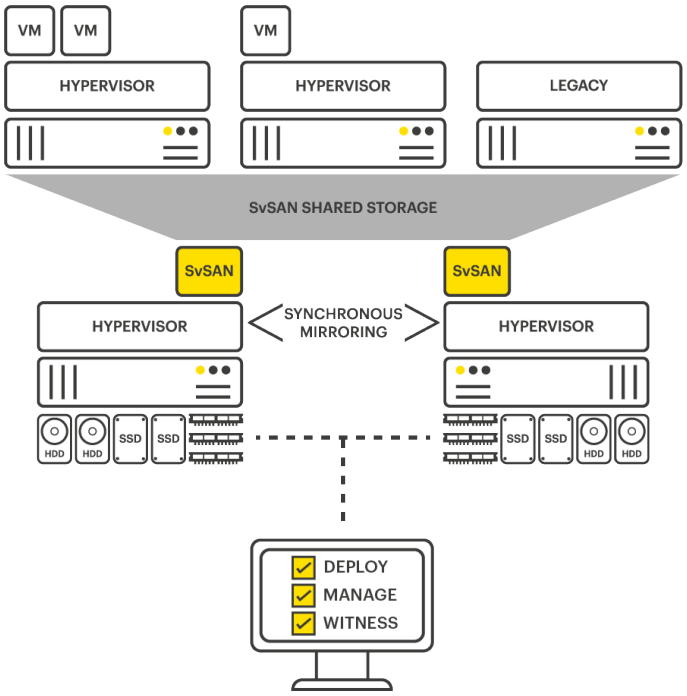

Fig. 2 *A typical server SAN deployment.*

Hyperconverged Benefits	Server SAN Benefits
Small footprint • Can be built using just two servers for high availability • Servers can be sized to the exact requirements • Ability to scale up (adding resources to existing servers) or scale out (adding more servers to the configuration)	**Lower entry cost than traditional storage arrays** • Using industry-standard hardware removes the cost of proprietary components from vendors • Decommissioned servers can be re-purposed • Simplified infrastructure reduces the need for components such as network switches • Advanced caching across flash and memory available with commodity components
Lower solution complexity • Starting with just two servers simplifies network topology and cabling • Requires IT staff with fewer specialist skill sets	**Flexibility** • Easily scale up and out from a small footprint • Deploy as a single server SAN or, for high availability, as a pair to increase resiliency • Avoid vendor lock-in by separating management from hardware and thus choose any industry-standard server
Highest performance • No additional latency as compute and storage are contained within the same server • Ability to use PCIe flash, SSD or server memory to act as a cache in front of the hard disks, dramatically improving IOPS and throughput	**Maintain compute and storage separation** • Allows hardware to be dedicated to delivering storage features or running VMs/applications without having to contend for the same resources • Enables the environment to scale compute and storage capacity independently, preventing the need for an entire refresh when new servers become available • Workloads/applications that are not virtualized can access highly available shared storage using industry standard protocols such as iSCSI. • Older servers can be re-purposed to deliver storage functionality. This negates the need for new servers and reduces the number going to landfill
Flexibility • Data management features such as caching, snapshots, asynchronous replication and encryption can be easily be added in the future • Can utilize any server storage (PCIe flash, solid-state disk or hard disks) enabling all-disk, all-flash or hybrid storage solutions	

Redundancy is ridiculous! SvSAN gives us five 9s, the gold-standard of redundancy. We've built a solution with SvSAN that offers a 10-fold improvement in performance at half the cost

Anders Kringstad, Solutions Architect
ITsjefen

SvSAN ARCHITECTURE

What is SvSAN?

SvSAN is the virtual SAN that is simple, flexible and reliable. The technology is based on software-defined storage that eliminates the need for physical SANs whether deployed as part of a hyperconverged solution or as a server SAN.

SvSAN delivers an iSCSI shared storage stack along with other storage features such as synchronous mirroring, read and write caching and centralized management. Its flexible architecture allows additional storage features to be added in the future, through software upgrades, without the need to swap out the servers.

How does it work?

SvSAN supports VMware vSphere, Microsoft Hyper-V and Linux KVM hypervisors and can utilize any commodity x86 servers. It is installed as a Virtual Storage Appliance (VSA) requiring minimal server resources to provide the shared storage necessary to enable the advanced hypervisor features such as High Availability/ Failover Cluster, vMotion/Live Migration and VMware Distributed Resource Scheduler (DRS)/Dynamic Optimization.

SvSAN can be deployed as a single server, simple 2-node cluster, or multi-node cluster, with the flexibility to meet changing capacity and performance needs. This is achieved by adding additional capacity to existing servers or by growing the SvSAN cluster, without impacting service availability.

A typical 2-node SvSAN configuration is shown in fig. 1.

Any server	Any storage	Scale up/Scale out
SvSAN can be deployed onto any commodity x86 server[1]	SvSAN can utilize any type of disk storage[1]	Avoid over-provisioning
• Use any server from any vendor. Avoid vendor lock-in and enables a multi-vendor purchasing policy. • Create clusters using servers from different vendors. • Repurpose older servers to deliver storage capabilities. Extend the useful life of a server and reduce acquisition/CAPEX costs.	• Use any type of block based storage device, including PCIe Flash, Solid-state disks (SSD), or hard disks (HDD). • Disks can be internal server storage or an external JBOD array • Use server memory for caching, ensuring the most frequently accessed data is located on the "fastest" storage tier. • Build all-flash, all-disk or hybrid storage solutions, based on cost and performance requirements.	• Build solutions to fit the exact business requirements. • Start small and scale the solution as required. • Increase compute and/or storage capacity independently of one another. • Scale up: Add more disks to an existing server or replace smaller capacity disks with larger capacity ones. Allocate more memory to SvSAN to meet the performance demands of the applications. • Scale out: Add more servers non-disruptively to a SvSAN cluster to increase the compute or storage capacity.

[1] Servers and components must be included on the VMware vSphere or Microsoft Hyper-V hardware compatibility lists.

DEPLOYMENT MODELS

SvSAN is typically deployed as a 2-node cluster. However, it can also be installed as a single node solution or as a multi-node cluster.

2-node cluster

Installing SvSAN onto two servers is the most common installation option and the minimum requirement to provide highly available shared storage. This installation option can tolerate a single server failure and provides the shared storage required to enable advanced hypervisor features to be used (see fig. 3)

An optional witness is used to act as a cluster quorum and eliminate split-brain scenarios from occurring. See the "Cluster Witness" section for more details on witness requirements and deployment options.

Single server

SvSAN can be installed on a single server to deliver shared iSCSI storage keeping costs to the bare minimum. Although this installation option does not provide any protection against a server failure, data can be protected using a RAID controller to provide disk mirroring or parity RAID (see fig. 4).

This installation option is best suited to environments where a loss of server is not important or with non-critical data that can be easily replaced, such as test and development environments. Additionally, this installation option can be used where high availability is not required, but the highest levels of performance are. In this scenario, SvSAN acts as a caching device utilizing its advanced predictive caching features.

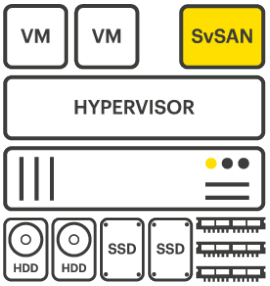

Fig. 4 *A single server configuration.*

NOTE: *This installation option does not prevent additional servers being added to the configuration that use the SvSAN synchronous mirroring capabilities at a later time. This can be done non-disruptively.*

Multi-node clusters

It is possible to scale out an SvSAN cluster beyond two servers. This introduces an added level of resiliency while a server is undergoing maintenance or suffers a failure. When deploying only two nodes per cluster, if one should fail, the applications and services will remain operational on the surviving node, but data cannot be replicated and is no longer protected. Introducing an additional node into the cluster protects against this scenario, maintaining high availability. More information about multi-node SvSAN clusters can be found in the **corresponding white paper**.

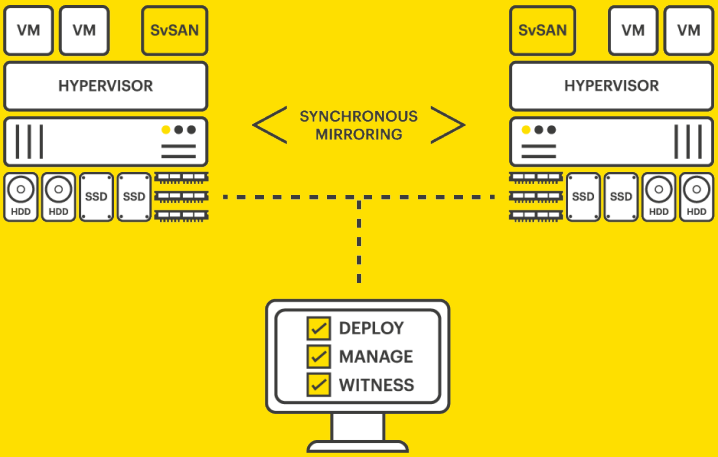

Fig. 3 *A 2-node cluster with witness.*

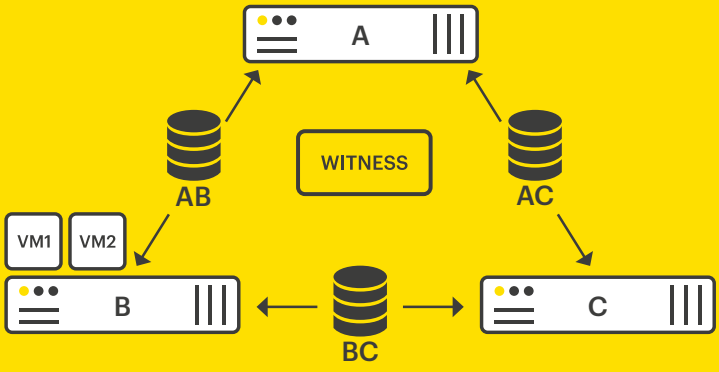

Fig. 5 *A multi-node cluster.*

In Fig 5, the three servers in the cluster are labelled “A”, “B” and “C”, with three synchronously mirrored volumes between the servers:

- Volume AB *(mirrored between servers “A” and “B”)*
- Volume BC *(mirrored between servers “B” and “C”)*
- Volume CA *(mirrored between servers “C” and “A”)*
- The witness provides quorum for each of the three volumes

Stretched clusters

Stretched clusters provide additional levels of resiliency for clusters with two or more nodes.

Stretched clusters are created when the servers running SvSAN are geographically separated, allowing copies of data to be stored in two discrete locations, this can mean residing in different server racks, at opposite ends of a datacenter or between buildings on a campus or across a city or urban area (see fig. 6).

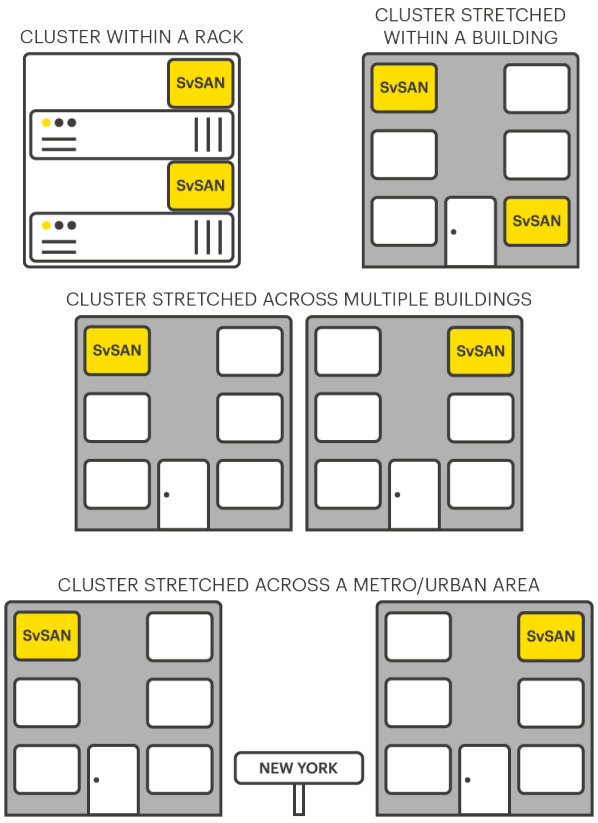

Fig. 6 *Different ways of leveraging a stretched cluster deployment*

They reduce the risk of all servers being affected by natural disasters, for instance earthquakes or hurricanes, or localized or “rolling” disasters such as fire, flooding, power outages or theft.

The main factor determining how far the servers can be separated is the network round-trip time latency, typically longer distances have higher latencies and it is important to understand the effect it has on the performance of synchronous write operations. If the network round-trip time latency becomes too great the write performance will degrade. SvSAN supports stretched clustering and this feature is included with all SvSAN licenses. More information is provided in **this white paper.**

NOTE: *Read operations will be unaffected as they will always be serviced from the local copy of data.*

Cluster witness

StorMagic SvSAN utilizes a witness which acts as a cluster quorum, providing a majority vote in the event of a failure that results in a cluster leader election process. The witness prevents SvSAN clusters from getting into a state known as split brain. To read more about the witness and for its minimum server requirements, please refer to **the SvSAN witness white paper.**

The witness has multiple configuration options to suit different requirements, these include:

- Shared remote witness:
  This enables the minimum amount of IT infrastructure equipment that provides high availability (two servers) to be located at the remote site. The witness is located at a central location (datacenter/HQ) and accessed over a WAN connection. Due to the lightweight protocol, it is possible for a single witness to be shared between, and act as a witness for thousands of SvSAN clusters in remote locations.

StorMagic SvSAN enables us to cost effectively implement simple, 2-server highly available systems with minimal hardware requirements

Uwe Fischer, Head of Asset Information Systems
E.ON Climate & Renewables

- Local witness: Similar to the previous witness configuration, there are only two SvSAN servers at each remote site. However, the witness is hosted on a third physical server or virtual machine (outside the SvSAN cluster) located at the same site. This witness configuration is for totally isolated locations and have limited or no external network connectivity.
- No witness: The final option is to deploy two servers at the remote site with no witness. This is for customers that do not have space or money to implement a witness (local or remote). In this configuration, it is possible to enter a split-brain scenario, for example, in the event of loss of network connectivity between servers or if the servers are rebooted simultaneously. To reduce the chance of split-brain occurring, best practices should be followed, these include providing resilient network connections between servers, using quality components and using multiple, redundant power supplies. For best practices in this area, **refer to this white paper**.

Performance and security enhancements
Depending on the type of deployment, the hardware used and the organization's requirements, it may be necessary to consider enhancements to the SvSAN cluster for performance or security reasons. For performance, SvSAN has a suite of caching features, including write-back caching to SSD, predictive read-ahead caching to SSD and memory, and data pinning that maps frequently-repeated operations which can be loaded into memory from disk ahead of time.

Collectively these features are known as **Predictive Storage Caching** and can help an organization better leverage its existing hardware rather than pay significant amounts for upgrades or a complete refresh.

For security, SvSAN can be enhanced by encrypting the mirrored storage through the data encryption feature add-on. SvSAN's data encryption feature delivers ultra-secure encryption using a FIPS 140-2 compliant algorithm. This ensures an organization's data held in edge sites is protected from the vulnerabilities often seen in these less secure environments. More information on encryption, and SvSAN's data encryption feature can be found in the **corresponding white paper**.

Furthermore, the key management necessary to operate SvSAN's data encryption feature can be provided by StorMagic SvKMS encryption key management. SvKMS is a highly flexible and robust key manager with enterprise-grade features that allow it to integrate with any workflow and be located anywhere - whether onsite, in the datacenter, or cloud. A complete explanation of SvKMS is available in the **product data sheet**.

USE CASES

The flexible SvSAN deployment options enable it to be used in a variety of different scenarios, including remote sites, small and medium enterprises and branch locations, and datacenters. Each use case has different characteristics such as number of sites, workloads and capacity.

Remote sites
SvSAN can be deployed into remote sites that have the following characteristics:

- Located in isolated or harsh environments. Generally, there are approximately one to fifteen interrelated or dedicated applications that contribute towards a single workload. For example factory automation or wind turbine monitoring.

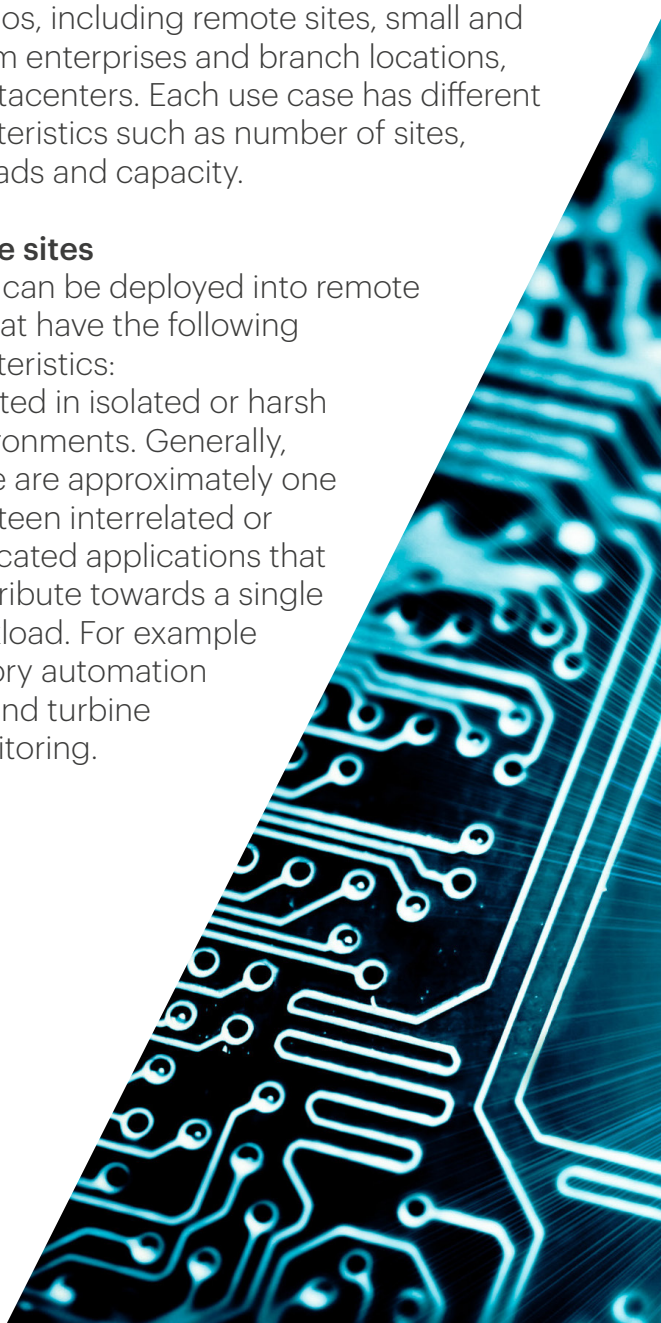

- Storing and processing data at the edge
- Deliver sub-second response times for latency sensitive applications
- Drastically reduce the network bandwidth requirements needed to transfer the vast amounts of data being generated.
- Requires simple, centralized management. Typically these sites are unmanned.
- Low complexity: simple,repeatable and easily deployable solution.
- Low cost, small installation footprint (low system resources).

The Internet of Things (IoT) is a prime example of this use case, where data is generated by machines and sensors rather than by human interaction. Examples include **oil rigs (drill monitoring)**, **wind farms**, **manufacturing plants** or **factories (SCADA systems)**.

Small and medium enterprises and branch locations

SvSAN can be deployed into small and medium enterprises (SME) or branch locations which generally have the following characteristics:

- Organizations with tens to thousands of autonomous remote customer-facing locations, each with a micro datacenter.
- Each site runs a variety of different workloads, from ten to thirty applications or services, that are business-critical or are latency-sensitive. This can include line of business applications (LoB), back office (email, messaging, file and print) and network services (directory services, DHCP, firewalls)
- A need for high availability, but without the cost and complexity of a traditional SAN/NAS solution.
- Requires simple management as each location has limited or no IT management skills onsite and will usually be managed remotely from a central location.
- Low cost. Typically these organizations have low budgets, therefore the cost of the solution needs to be kept to a minimum, as a change will be multiplied by the number of sites.

Example applications of this use case include **big box retailers**, **hotels**, **legal practices**, **banking institutions**, **universities**, and **hospitals/medical practices**.

SME datacenter

Small datacenter environments are another use case for SvSAN. It can be seen as a cost-effective replacement for traditional SAN and NAS storage arrays that doesn't require any specialist fibre channel, networking or storage skills.

The datacenter use case has the following characteristics:

- Single site or stretched cluster deployments.
- Large storage capacity requirements.
- High availability is required to meet strict service level agreements (SLAs).
- Support for many (20+) disparate workloads with varying performance requirements.
- The ability to increase storage or compute capacity independently of one another.

Example applications of this use case include **hosting environments**, **airports**, and **banking institutions**.

The following table highlights which deployment options are applicable for each use case:

	Remote sites	SME and branch locations	Small data-center
Workloads or applications	1-15	10-30	20+
Hyper-converged	Applicable	Applicable	Applicable
Server SAN	Not applicable	Somewhat applicable	Applicable

Example configurations

The following table shows some typical server configurations. These can vary and depend on each user's specific requirements, applications/ workloads and performance requirements.

	Remote sites	SME and branch locations	Small data-center
CPU	1 or 2	2	2+
Memory	32GB to 96GB	64GB to 256GB	>128GB
Disk	≤4TB	≤12TB	>8TB
Network	1GbE	1GbE or 10GbE	10GbE

FURTHER READING

StorMagic SvSAN is a simple solution, but we understand its features and details cannot be explained completely in a single white paper. Therefore, we have compiled a wide range of documents and resources to assist storage architects and interested parties, and these are outlined below.

Additional details on SvSAN are available in the **SvSAN Data Sheet** which summarizes system requirements, compatibility and support levels.

Dive deeper into the individual features of SvSAN, such as **Predictive Storage Caching**, the **SvSAN witness** or **Data Encryption.** These features and more can be accessed through the **extensive collection of white papers** on the StorMagic website.

Furthermore, explore how SvSAN integrates with third party software and hardware with solution briefs on everything from **Lenovo ThinkSystem servers** to **Citrix Workspace Appliances**. Access all of the solution briefs on **the StorMagic website**.

Learn how others have made use of SvSAN in a wide range of environments and applications **with case studies and configurations** covering many industries and locations around the world.

Finally, why not evaluate SvSAN with a free trial license, available to download immediately from the StorMagic website. Simply complete the form and begin testing the software in your own environment, with StorMagic support available if required. **Start your free trial now.**

StorMagic
The Quadrant
2430/2440
Aztec West
Almondsbury
Bristol
BS32 4AQ
United Kingdom

+44 (0) 117 952 7396
sales@stormagic.com

www.stormagic.com