

OUTDATED

# Hyper-Converged with Software Defined Storage

A Review of Scalable Hyper-Converged Technologies

LAST UPDATED: JULY 2014

WHITE PAPER BY CHRIS M EVANS

Independent Consultant  
Langton Blue Ltd.

## Introduction and Executive Summary

As virtualisation technology matures and businesses look to deploy ever more efficient scalable solutions, a new breed of technology has emerged. Hyper-converged infrastructure solutions merge compute, storage and networking together to take things a step further than the simple packaging offered by the first generation of converged infrastructure systems.

The first generation of converged solutions simply packaged components together into a single rack-based deployment that could be ordered under a single SKU (Stock Keeping Unit). Many solutions were static in nature and provided companies nothing more than a single point of contact and some savings in solutions testing.

Hyper-converged solutions merge server and storage into the same form factor, eliminating the need to deploy and manage separate physical components. Bringing these constituent parts together allows simplification of the hardware and the ability to move many features to software, including data protection, data reduction, resiliency and the ability to scale out.

Vendor offerings are not restricted to packaged hardware and many solutions are available in software-only format. These kinds of products can be beneficial to the customer in many ways, fitting in with existing hardware infrastructure standards and practices. Ultimately the features of the software define any hyper-converged solution. Some of the most common vendor products are discussed in this paper.

The conclusion of this white paper is a feature comparison between Virtual SAN from VMware and Virtual SAN from StarWind Software.

VMware Virtual SAN is a version 1.0 product, first shipped with VMware vSphere 5.5 U1 and higher. It provides a distributed storage environment across multiple nodes in a vSphere cluster using local DAS storage. As a first release, Virtual SAN offers limited data efficiency features, supporting only thin provisioning.

StarWind Virtual SAN is the evolution of the StarWind iSCSI Target, which has been in development since 2003. The Virtual SAN software integrates into Windows Server environments (2008 R2 and later) adding data services (availability, performance, efficiency) to complement those offered on the Windows platform.

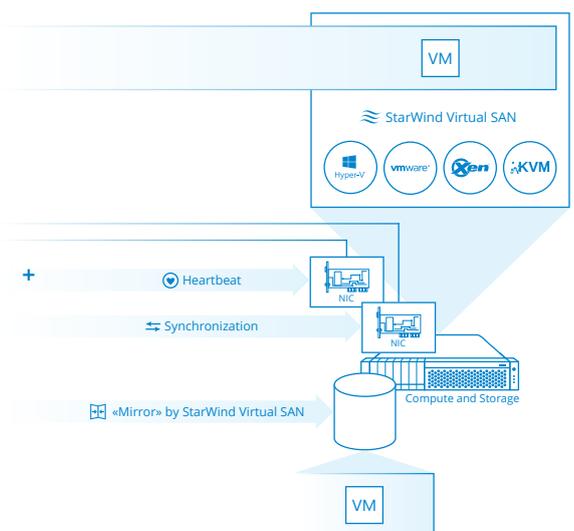
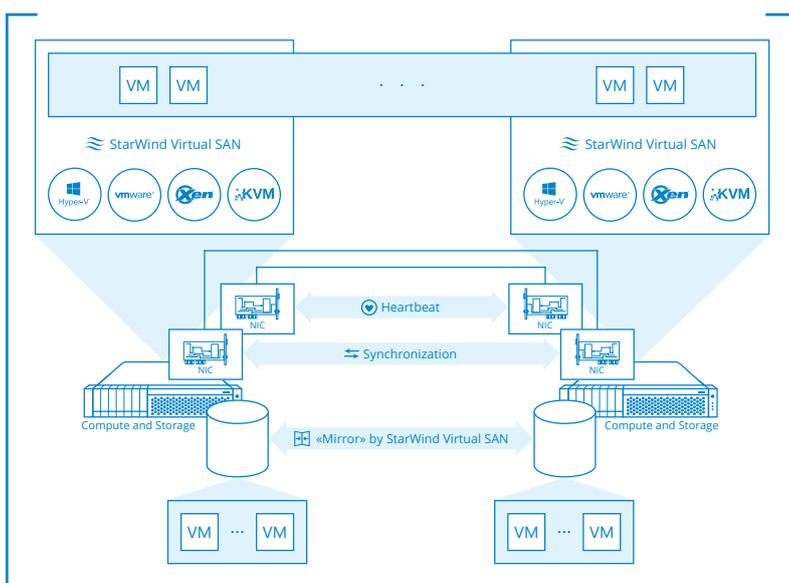
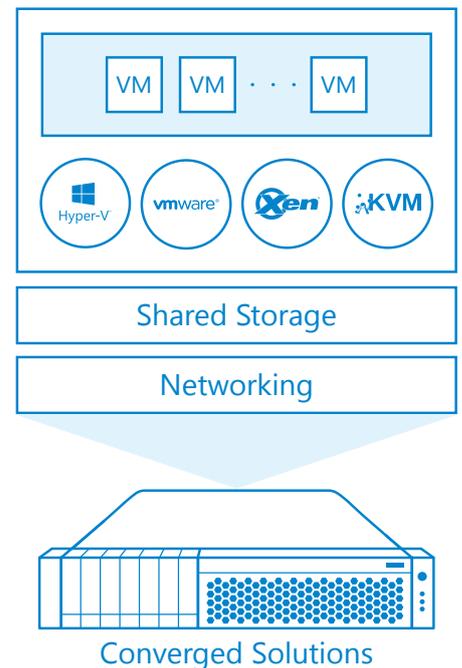
The comparisons showed that StarWind Virtual SAN provides a greater set of features compared to VMware Virtual SAN, with specific advantages in data protection, data efficiency and I/O performance.

## Putting a Definition to Hyper-Converge

Hyper-converged solutions follow on from the first wave of consolidated hardware and software offerings known as Converged Infrastructure. Converged Infrastructure solutions have been around since late 2009 and package servers (compute) with storage and networking to deliver an “all-in-one” offering, usually managed by custom software that provides a layer of orchestration to provision and manage virtual machines.

Many converged solutions are static in nature, with little flexibility to expand the configurations and so do little more than provide customers with a “single throat to choke”, alleviating some of the overhead of certifying multiple components into a single solution. Some converged infrastructure offerings aren’t delivered as a product and are classed as “reference architectures”, requiring the customer to purchase the individual components separately. However all Converged Infrastructure solutions place all of the hardware components into a standard rack configuration that still ultimately requires storage management, networking and server skills.

Hyper-converged solutions take things a step further than simple product packaging by eliminating the separate storage and server components, merging them together into a single form-factor. The “secret sauce” is delivered in the software, which provides the functionality of converged solutions (and more) without the additional hardware and management overhead.



Hyper-converged solutions are divided into two main categories; hardware-based and software-based offerings.

## Hardware-based

Hardware-based solutions take standard off the shelf servers, integrating networking and storage into the chassis to create an appliance. The appliances will typically be configured with a number of different storage options; hard drives for capacity and flash in multiple forms to deliver I/O at low latency. The hardware will be supplemented with software that provides the functionality of a file system with features for data protection & reduction, automated provisioning, monitoring the configuration and managing the expansion of the system.

## Software-based

These solutions deploy directly onto customer hardware, providing the same level of software features as hardware-only solutions. As no hardware is shipped, the customer is responsible for choosing technology that works with the hyper- converged solution and that means ensuring components are both reliable and able to perform at the level of workload required. Software solutions may also have different charging models, as there are no hardware maintenance or upgrades to consider. As before, software-based solutions are all about the “secret sauce” in the software itself, which implements the advanced data management and availability features.

## The Good and Bad Sides of Hyper-Convergence

The idea behind hyper-converged solutions is one of simplification. Large enterprises are able to support significant product engineering teams to validate and build configurations to their requirements. However many small and medium enterprises lack the funding and skills to put in the same level of due diligence and design and of course many simply don't have the economies of scale compared to large enterprises. These smaller companies have relied on vendors to provide that level of support and guidance on their behalf. Hyper-converged solutions offer benefits in terms of simplicity and reduced cost as a result of:

- **Removal of dedicated hardware** – hyper-converged solutions specifically allow the removal of dedicated storage hardware, moving permanent storage media back into the server in what would have been previously described as DAS. Storage capacity and performance growth is managed through scaling out the hardware, distributing data across all nodes and so reducing the shortcomings of legacy DAS-style deployments.
- **Hardware agnostic** – software-only solutions can be deployed on a wide variety of hardware platforms from the major server vendors. We have already seen HP, Dell and others come forward with reference architectures and designs for VMware Virtual SAN, which also has a specific hardware support matrix. Allowing the customer to choose their own hardware means solutions can be deployed on hardware the customer is familiar with (reducing training costs) and already have contracts in place for support.

- **Simplification** – simplifying solutions is more than just removing components such as storage. For hardware offerings, vendors have tested and validated their hardware configurations and understood the scalability capabilities. Many solutions allow for capacity growth by adding new nodes or additional components, thereby easing migration issues. Software solutions provide options to migrate data between tiers or between servers, allowing a more managed growth strategy without application downtime.
- **Increased efficiency** – Hyper-converged solutions potentially offer greater performance improvements than traditional deployments as storage latency is significantly reduced by placing active data closer to the processor. The distributed nature of hyper-converged solutions means all components are being used all of the time, for maximum efficiency.

With hyper-converged solutions, there is a debate on whether the software components should be deployed within the kernel of the hypervisor or within a virtual machine. Kernel-based solutions such as VMware Virtual SAN are more closely integrated with the operating system and vendors will sell the virtues of closer software integration. However these solutions are also tied to the hypervisor for both feature and fix upgrades, potentially making maintenance of a hyper-converged solution more difficult. By comparison, VM-based or driver-based hyper-converged solutions have the benefit of being more easily upgradable as they have little if no dependency on the hypervisor platform, other than to support specific features and configurations. The downside for these configurations is that the storage component is run in parallel with virtual machines serving production workloads and so there is a risk of resource contention, which the user needs to manage.

## Features to Look Out for

What features are important and should be available in hyper-converged solutions? Although vendor solutions vary considerably, there is a set of basic features that should be deliverable in all solutions.

- **High Performance** - hyper-convergence provides the ability to deliver performance in a number of ways. Storage is more closely located to the processor (compared to converged solutions), reducing latency. Applications and storage can also be distributed across a hyper-converged cluster, allowing a flexible use of available resources, reducing bottlenecks and improving overall throughput.
- **High Scalability** – hyper-converged solutions should scale well. Most are built around the concept of grouping multiple nodes (i.e. physical servers) together, with each node bringing processing capacity, storage capacity/performance and network bandwidth. As solutions scale, more nodes can be added, with the restriction that the maximum size of a virtual machine will be limited to the capacity of one node and the overall size of a cluster may be limited to the restrictions in the size of the cluster.
- **Data Efficiency** – most solutions offer savings in storage capacity through the use of a number of data reduction techniques including thin provisioning, compression and de-duplication. The savings

achieved depend on the type of data being processed, with VDI and virtual server environments seeing the most benefit. Figures quoted by vendors range from 2:1 for structured to 10:1 for VDI, with a typical average 6:1 savings ratio.

- **Data Protection** – Hyper-converged solutions should offer data protection features. These can be divided into local resiliency (managing component and node failures) and site resiliency (through replication). The resiliency that was previously built into external storage arrays has to be delivered by the hyper-converged platform. Scale out solutions use features such as distributed file systems to achieve this, whereas the more basic offerings simply mirror data between physical disk devices. Remote replication provides for geographic resiliency and can be used as part of a wider data protection strategy.
- **Management/Reporting** – Effective management is a key component of any hyper-converged solution. Management tool features include visualisation of the infrastructure, automated provisioning of VMs, automation/balancing of cluster resources and alerting.
- **Security** – Security is a concern for any solution; where hard drives and flash storage is used in the server, there are issues of data encryption at rest and of course for data in flight, encryption across WAN links. In traditional infrastructure deployments using shared SAN or NAS storage, vendors offer encryption for physical disk and in most cases will certify the destruction of hard drives removed from arrays. However with hyper-converged solutions, the onus is on the customer to destroy drives or ensure data at rest is encrypted on disk in the first instance.
- **Licensing** – cost is a major factor in implementing hyper-converged solutions. The cost model for hardware is typically based on the hardware cost itself, with annual maintenance. Pricing may be based “per node” or CPU socket and vary according to the node specifications. Software-only solutions can be licenced differently, including per TB of storage deployed or per node configured (irrespective of the node capabilities). As with traditional infrastructure deployments, understanding the TCO of a solution can be complex and even hyper-converged solutions need to be carefully validated.

## Vendor Roundup and Comparisons

Which vendors are offering hyper-converged solution in the market place today?

### Hardware

- **Nutanix** was founded in 2009 and released their first product (called the Virtual Compute Platform) to the market in 2011 under the marketing slogan “No SAN”. A VCP deployment consists of at least three physical nodes, combined in a cluster. Each node implements a virtual machine to host storage resources that are combined to create a cluster-wide scale-out file system (called SOCS – Scale Out Converged Storage). Hypervisor hosts access the storage through iSCSI LUNs presented from the virtual machines. VCP offers a range of data reduction features and supports VMware vSphere, Microsoft Hyper-V and KVM platforms.

- **Scale Computing** – Scale is a startup company focusing on the SMB end of hyper-converged solutions. Their platform, known as HC3, combines compute, networking and storage using open source KVM as the hypervisor and their own in-house developed object store technology. Scale have developed an entire ecosystem around KVM with their own Software Defined Storage object layer known as SCRIBE (Scale Computing Reliable Independent Block Engine) and a state engine which is used to manage the status of hardware and software components in the ecosystem. This makes HC3 very much a software-based solution.
- **SimpliVity Inc** – SimpliVity is a startup company, also founded in 2009 that offers a hardware-based scale out solution. Their offering is known as OmniCube and encompasses, compute, storage and networking in a single hardware platform with additional features such as WAN optimisation for distributed nodes, de-duplication and global scale out. OmniCube achieves high I/O performance levels through the use of a dedicated PCIe module that manages inline compression and deduplication of data. The data deduplication feature is global, allowing virtual machines to move to remote nodes by transferring only the uncommon de-duplicated data.

## Software

- **Atlantis Computing** – Atlantis initially focused on the VDI market with their original releases of ILIO (Inline Image Optimisation). The company recently released ILIO USX, a version of the software focused on using flash and system memory to accelerate virtual machine deployments on VMware vSphere, although the technology is hypervisor agnostic and will work across hypervisor types.
- **DataCore SANsymphony** – DataCore Software has been in the market since 1998 and recently released version 10 of their flagship SANsymphony storage virtualisation platform. SANsymphony offers a wide range of features, including asynchronous replication, storage pooling and tiering, caching, thin provisioning and CDP. As yet the product does not offer compression or data deduplication.
- **Maxta MxSP** – Maxta Inc was founded in 2009 and is a startup company building a distributed storage solution for virtual server environments called Maxta Storage Platform (MxSP). The MxSP software implements a shared storage environment with features including thin provisioning, inline compression and de-duplication and high availability through virtual machine failover.
- **StarWind Software** – StarWind have been providing storage solutions for Windows Server environments since 2003. Their iSCSI target software has evolved into their current flagship product, StarWind Virtual SAN. The software extends and enhances the features of Windows Server with features including scale-up and scale-out deployments, fault tolerance and high availability, caching, asynchronous replication, automated storage tiering, deduplication and compression. Virtual SAN can be deployed in both Windows Hyper-V and VMware vSphere environments to create hyper-converged infrastructure solutions.

- VMware Virtual SAN** – VMware originally announced Virtual SAN at VMWorld 2013 in San Francisco. The platform was released in beta in September 2013 and went GA in March 2014. Virtual SAN implements a distributed datastore across a number of clustered vSphere ESXi hosts, aggregating disk storage and using flash as a read and write I/O accelerator. Virtual SAN is only at release 1.0 and this is highlighted in some of the features that are (and are not) implemented. Data protection is achieved solely through mirroring (with additional resiliency achieved through more mirrors, also known as replicas). Data deduplication is currently not supported. Flash storage is used as a cache only and not as a permanent store for data.

## Virtual SAN Comparisons – VMware and Starwind

### VMware Virtual SAN

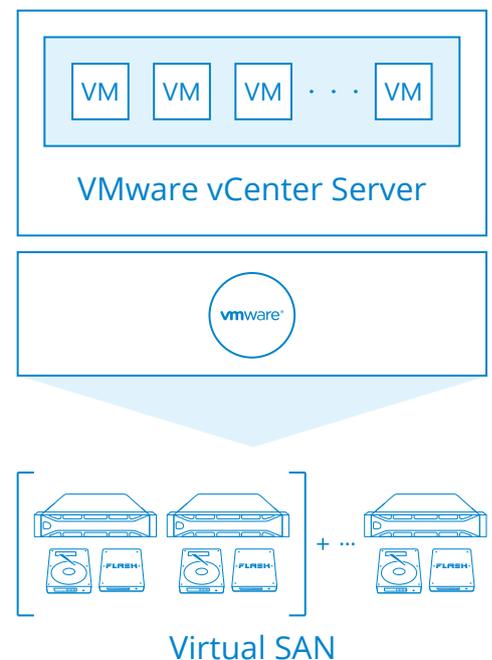
Virtual SAN, more commonly known as Virtual SAN from VMware has been in general availability since the release of VMware vCenter Server 5.5 Update 1 on 11 March 2014. Virtual SAN was added as a feature to the vSphere ESXi hypervisor but requires vCenter for operations and licensing.

A Virtual SAN cluster comprises of a minimum of three vSphere ESXi 5.5 hosts in order to satisfy the minimum requirements of data redundancy. This is broken down into two servers which store mirrors (or replicas) of virtual machines in a cluster-wide distributed datastore, plus a third server to act as a Witness, in the event that networking problems make it impossible to communicate between the data hosts – otherwise known as the classic “split brain” scenario. Operationally, Virtual SAN can potentially create more Witnesses depending on the number of hosts in a cluster and the layout of data across the cluster members.

Virtual SAN uses local resources from each host, building pools of disks and flash storage into a single distributed datastore. Each contributing host must have a minimum of one hard disk and one solid state disk drive, although not all hosts in a cluster need to contribute to the storage resources but may still consume them.

Virtual SAN storage policies are applied to virtual machines through the use of rules; a policy is constructed using one or more rules and vSphere then applies those rules when storing the virtual machine in the distributed Virtual SAN datastore. Currently five rules (or capabilities) exist:

- Number of Failures to Tolerate** – number of host/network/disk failures that may occur while still ensuring VM availability.



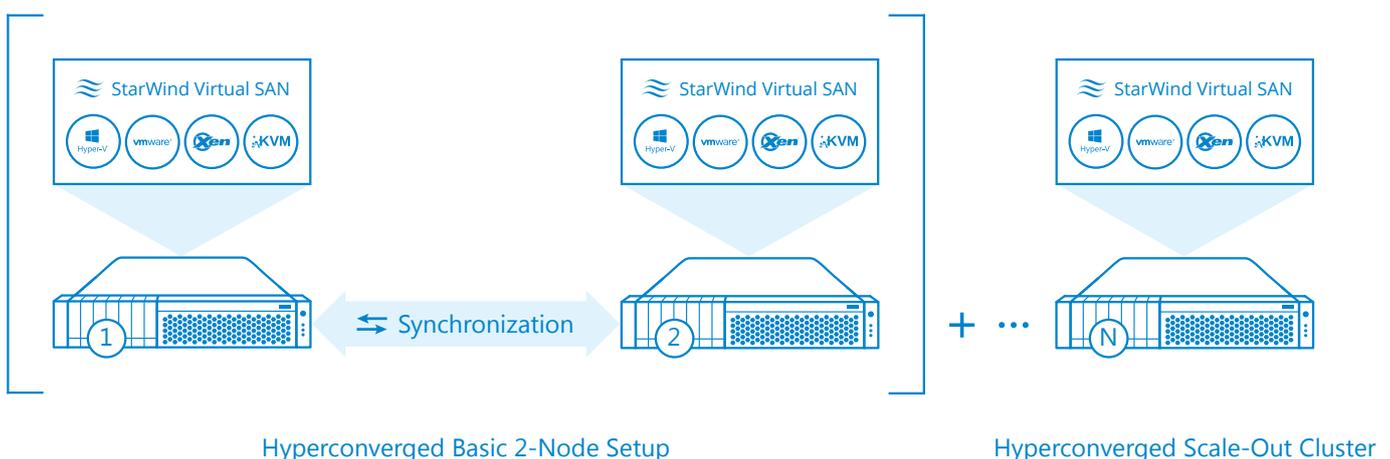
- **Number of Disk Stripes Per Object** – number of physical disks across which to stripe each replica of a virtual machine (depending on physical disks available).
- **Flash Read Cache Reservation** – a percentage figure of the virtual machine VMDK that should be retained in flash.
- **Object Space Reservation** – a percentage figure indicating how much of the configured capacity of a VM should be reserved on disk (a figure of 100% results in a fully provisioned “thick” VM).
- **Force Provisioning** – indicates whether provisioning of a VM should still be done, even if it breaks the VM Storage Policy rules for availability.

Virtual SAN is very much a version 1.0 product; there is no support for data compression or de-duplication and the only policy for data protection is multiple mirroring. Virtual SAN Capabilities are infrastructure rather than service-based, focusing on physical aspects of the configuration rather than logical metrics like Quality of Service.

Virtual SAN standard licensing is based on the number of CPU sockets in a cluster. Each socket is priced at \$2,495 at the time of writing.

## StarWind Virtual SAN

Virtual SAN from StarWind Software is an evolution of the iSCSI target product originally developed by Rocket Division Software. StarWind was spun off as a separate company in February 2009 and has continued to focus on developing the iSCSI target software into today's Virtual SAN.



Virtual SAN is natively supported on Windows 2008 R2 and later, including the current flagship releases of Microsoft's Windows Server platform, 2012 R2. Deploying Virtual SAN on Windows allows the software to use native features of the Windows operating system, such as Storage Spaces for data protection. Virtual SAN may also be used in Hyper-V environments, creating hyper-converged solutions. Virtual SAN is also supported as a VM guest on VMware vSphere.

Virtual SAN contains many data reduction and resiliency features, including inline data de-duplication and compression. Data within a Virtual SAN LUN can be stored using a Log Structured file layout, which helps to reduce the “I/O blender” effect of random I/O produced by virtual environments.

Local data protection is achieved through synchronous mirroring, allowing a minimum configuration of only two servers to be deployed for high availability. Remote data protection is achieved through asynchronous replication, which can target either another geographic location or replicate into public cloud infrastructure such as Microsoft’s Azure or Amazon AWS.

| Feature                         | VMware Virtual SAN  | StarWind Virtual SAN (with Hyper-V)   |
|---------------------------------|---|---|
| Minimum Node Count              | Three (four recommended to ensure compliance during server maintenance).  | Two   |
| Maximum Node Count              | 32 (limited to vSphere cluster Size).   | No limit (proprietary clustering feature)   |
| Minimum Storage Configuration   | One HDD and one SSD per host.   | One HDD per host.   |
| Maximum Storage Configuration   | 5 Disk groups, each with 7 HDDs and one SSD (35 HDD, 5 SSD).  | No practical maximum.   |
| Data Protection (intra cluster) | Multiple replicas (mirrors) per VM, fault tolerance based on node count. Disk or node failure triggers reconfiguration and data copying to regain resiliency level. | LUN-level RAID protection; can also use the data protection features of Windows for intra-node data protection. |
| Data Protection (inter cluster) | Metro clustering not supported.   | Synchronous and Asynchronous data replication supported.  |
| Data Protection – Snapshots     | Uses native ESXi Snapshot functionality and capabilities  | CDP-level snapshots (block level granularity) with unlimited rollback points.                                   |
| Performance – Write Ordering    | No native write-ordering capability.  | Sequential write ordering achieved through Log Structured File System.  |
| Performance – Caching           | One SSD required per host/disk group, used as read (70%) and write (30%) cache. Cache writes replicated between nodes.  | Caching supported in DRAM (Level 1) and SSD (Level 2). Cache contents replicated between multiple hosts.        |

| Feature                  | VMware Virtual SAN   | StarWind Virtual SAN<br>(with Hyper-V)  |
|--------------------------|--|---|
| <b>Storage Tiering</b>   | No. Virtual SAN creates a single datastore across all local disk resources. Storage DRS not supported with Virtual SAN.                    | Tiering can be enabled through use of multiple hypervisor LUNs with multiple disk characteristics (windows). Node level tiering for offload of cold data snapshots. |
| <b>High Availability</b> | Virtual SAN is supported with vSphere HA (with issues/ restrictions around multi-server failure). VMware Fault Tolerance not supported.    | Active-Active, Active-Active-Active & Active-Passive with failover/failback and data synchronisation.   |
| <b>Data Reduction</b>    | Thin provisioning only, no data deduplication, no data compression.  | Thin Provisioning, data duplication are supported.  |
| <b>Licensing</b>         | \$2495 per CPU socket for every Virtual SAN cluster member contributing or consuming storage; \$50/user in VDI solutions. No free edition. | Feature limited free edition is available. VMware edition free of charge (2-node config, unlimited capacity, correct at time of document creation).                 |
| <b>Management</b>        | Requires vCenter Server (at additional cost) for management and operations (Virtual SAN only supported in Web client)                      | Dedicated management tool for multi-server configuration, management and monitoring.  |
| <b>Platform</b>          | vSphere 5.5 U1 and higher only supported. No cross-vendor support.   | Native support for Hyper-V on Windows 2008 R2 onwards and for vSphere running as a guest VM.  |



Chris M Evans has worked in the technology industry since 1987, starting as a systems programmer on the IBM mainframe platform, while retaining an interest in storage. After working abroad, he co-founded an Internet-based music distribution company during the .com era, returning to consultancy in the new millennium. In 2009 he co-founded Langton Blue Ltd ([www.langtonblue.com](http://www.langtonblue.com)), a boutique consultancy firm focused on delivering business benefit through efficient technology deployments. Chris writes a popular blog at <http://blog.architecting.it>, attends many conferences and invitation-only events and can be found providing regular industry contributions through Twitter ([@chrismevans](https://twitter.com/chrismevans)) and other social media outlets.

## Langton Blue Ltd

Langton Blue Ltd is a hardware and software independent consultancy company, working for the business value to the end customer. We provide consultancy services to end-users and to other businesses looking to understand today's technology challenges. Contact us to discuss how we can help you bank the benefits in your technology investments.

Post: Langton Blue Ltd 78 York Street Londonz W1H 1DP

Web-site: [www.langtonblue.com](http://www.langtonblue.com)

Twitter: [@langtonblue](https://twitter.com/langtonblue)

E-mail: [enquiries@langtonblue.com](mailto:enquiries@langtonblue.com)

Phone: 0845 275 7085

## StarWind Software

| US Headquarters  | EMEA and APAC   |
|--|---|
|  1-617-449-7717<br> 1-617-507-5845 |  +44 20 3769 1857 (UK); +49 302 1788 849 (Germany);<br>+33 097 7197 857 (France);<br>+7 495 975 94 39 (Russian Federation and CIS)<br> 1-866-790-2646 |

Support Portal: <https://www.starwind.com/:support>

Sales: [sales@starwind.com](mailto:sales@starwind.com)

Support Forum: <https://www.starwind.com/forums>

General Information: [info@starwind.com](mailto:info@starwind.com)

### In 2016, Gartner named StarWind "Cool Vendor for Compute Platforms".

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.