www.Storage-Insider.de

A Publication of:
Storage Insider

Software-defined storage

eBOOK

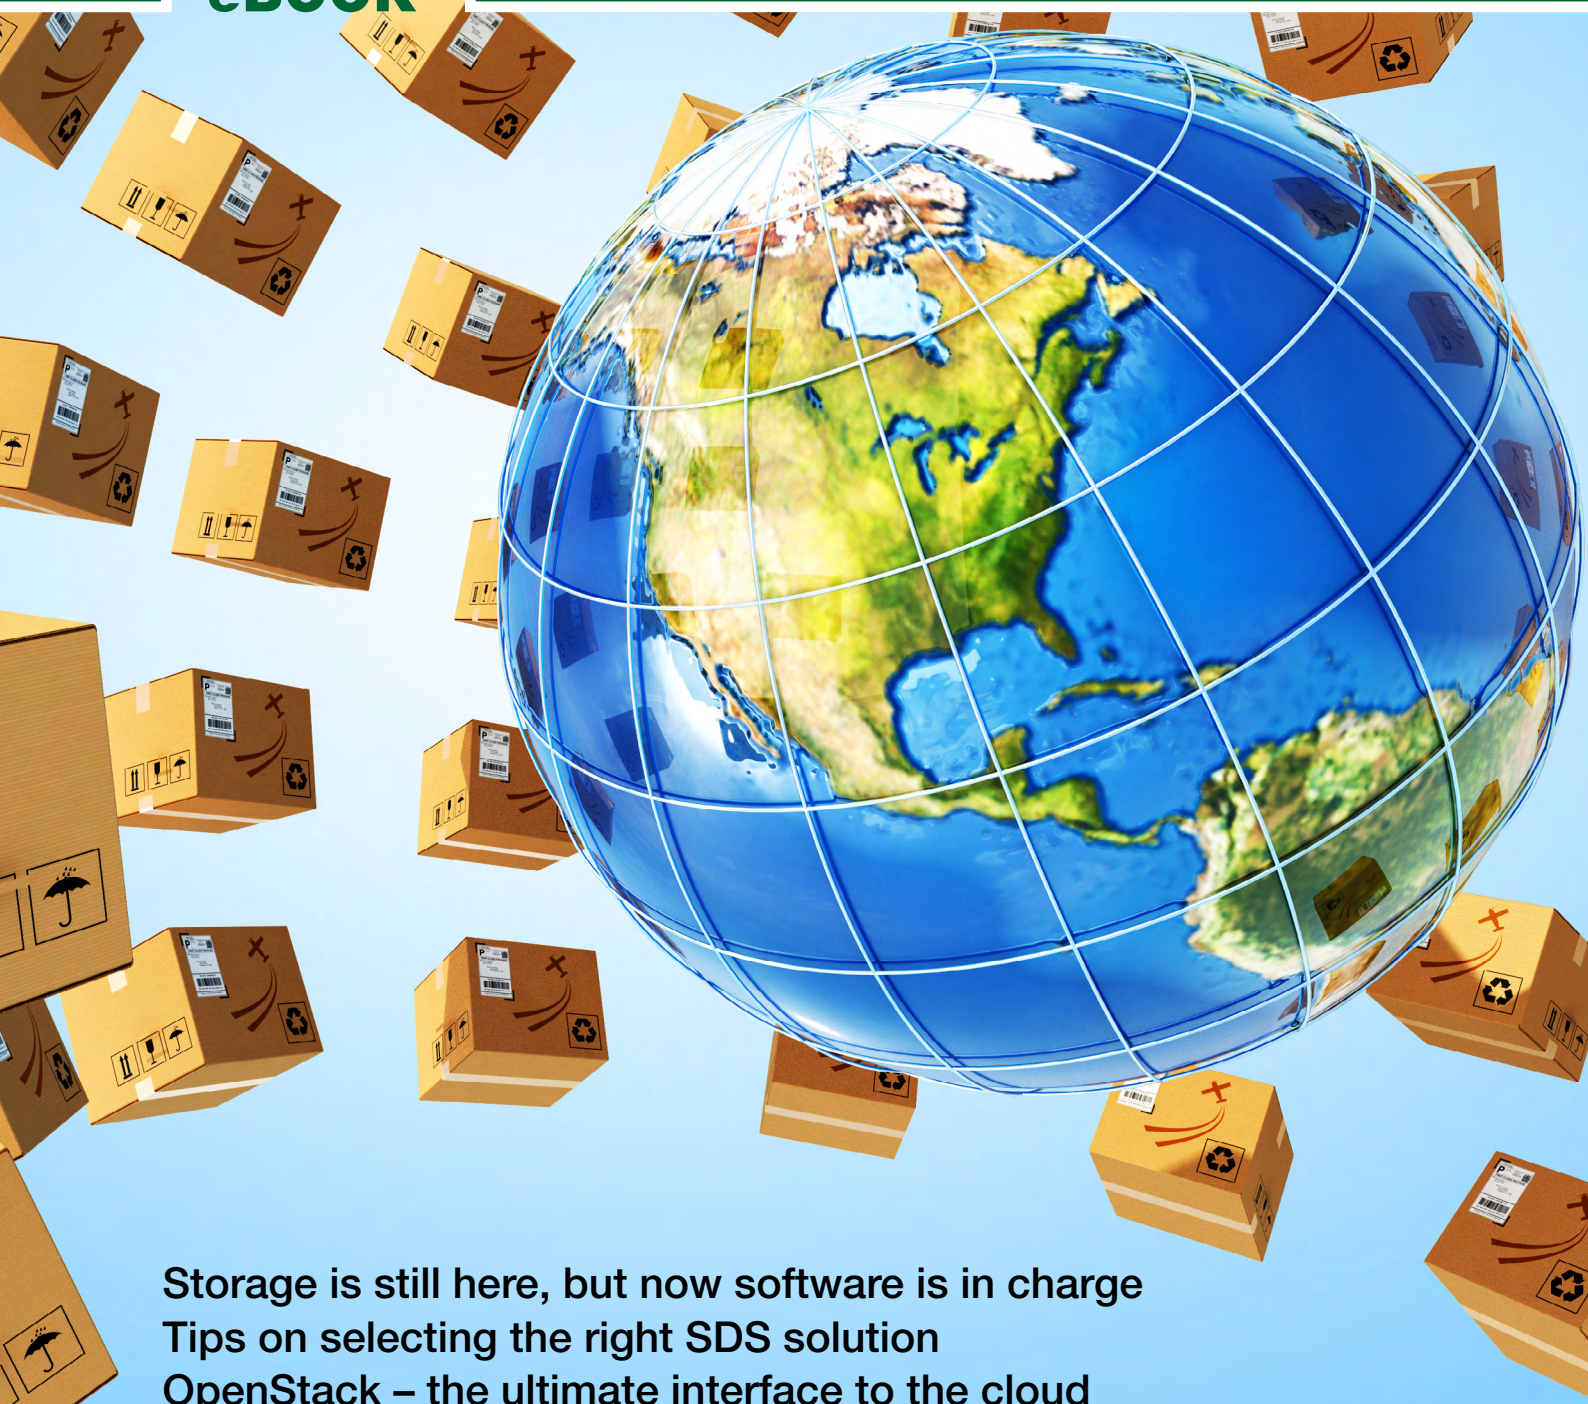

Storage is still here, but now software is in charge
Tips on selecting the right SDS solution
OpenStack – the ultimate interface to the cloud
SDS solutions – an overview

Powered by:
Tintri

Vogel Business Media

Content

Powered by:

Tintri

Tintri
Fountain House, 10th Floor
130 Fenchurch Street
London EC3M 5DJ
Phone +44 (0) 203 053 0853
Email emea@tintri.com
Web www.tintri.com

Vogel IT-Medien GmbH
August-Wessels-Str. 27, 86156 Augsburg, Germany
Phone +49 (0) 821/2177-0
E-mail redaktion@storage-insider.de
Website www.Storage-Insider.de
General manager: Werner Nieberle
Editor in chief: Rainer Graefen, responsible as per press laws, rainer.graefen@vogel-it.de
Publication date: September 2016

Title image: Cybrain - Fotolia.com

Vogel Business Media

The evolution of storage

Software-defined storage is just a small aspect of the vast OpenStack offering that is currently taking shape with strong support from the storage industry.

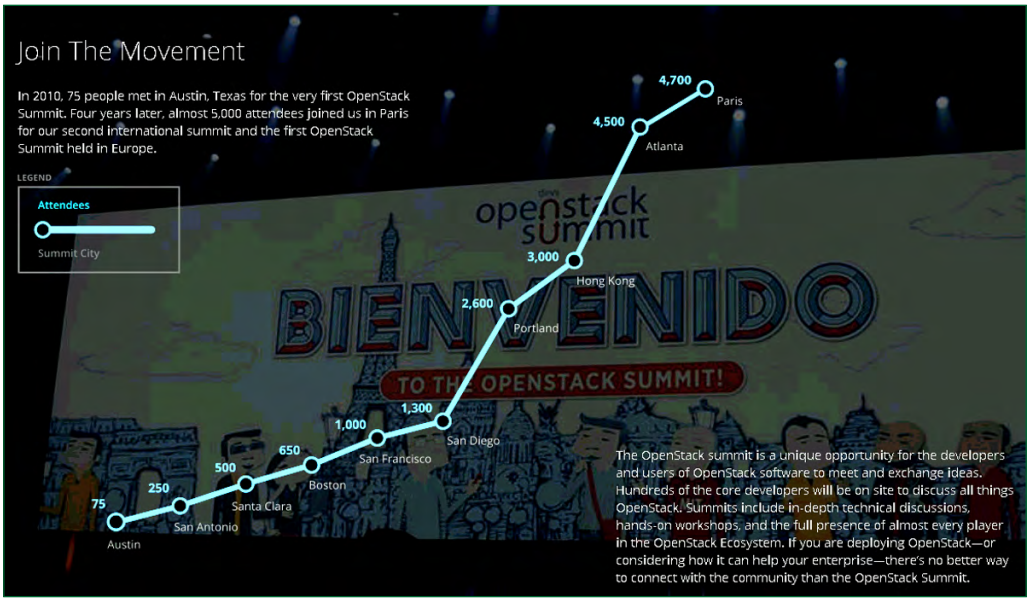

OpenStack is forging ahead: only 75 people took part in the first OpenStack summit in 2010, but some 6,000 visitors attended the 2015 event in Vancouver. Image: OpenStack Foundation)

It is a full twelve years since the storage industry was taken by the magnificent idea of creating a storage pool. Data went into the pool, while metadata was managed externally. As with so many innovations, however, the timing proved to be a big stumbling block. Another sticking point was that some operators linked server and storage via a hub, while others used a switch for the purpose. And last but not least, the interface software and the manufacturers' willingness to achieve interoperability left a lot to be desired.

Initially things worked fairly well when using products from mutually certified network and storage manufacturers. Thanks to the Fibre Channel (FC) protocol, the SAN has now achieved a level of maturity that our Ethernet colleagues struggled to attain with Fibre Channel over Ethernet (FCoE). Despite all the innovations in server, network and storage virtualisation, thin provisioning and other areas besides, it is only users who are still interested in the SAN.

Prospect of declining sales?

Software-defined storage (SDS) is being touted as the solution to the perennially over-expensive storage of ever-growing quantities of data. Such solutions can be found in a scalable cluster of low-cost storage media, with data flows, as well as backup/deduplication, replication, snapshots, synchronous and asynchronous mirroring and disaster recovery being managed by software. The manufacturer-specific and – depending on application area – specialized controller hardware of the storage systems is being rendered obsolete. This poses a threat to steady sources of income, as storage is becoming a commodity – just as EMC, IBM and HP have long feared.

Now it is up to Open Source to master the business-critical operation of the software-defined data centre via software. SDS is just a small aspect of the vast OpenStack offering that is currently taking shape with strong support from the storage industry.

But all the delightful project names such as Cinder (block storage) and Swift (object storage) cannot disguise the fact that the virtualisation of compute, network and storage layers in an heterogeneous environment is an extremely complex undertaking.

Rainer Graefen
Editor-in-chief
Storage-Insider.de

Nobody wants complex IT

So it's wise to exercise a degree of caution when seizing upon techniques designed to consolidate storage and make it simpler to operate. The history of storage shows all too well that involving too many manufacturers results in a far from standard platform. It took a number of years for people to grasp even the basic concepts of storage virtualisation. In fact some users still shy away from it, as the virtualisation layer can complicate many a simple concept.

SDS and OpenStack harbor a great deal of potential. But, as with virtualisation, it could easily take ten years to realign an IT infrastructure.

Best regards, Rainer Graefen

Storage is still here, but now software is in charge

The storage industry faces a huge challenge. In future, the aim is for every hardware-specific functionality of a storage system to be controlled by software. The functional repercussions of this break with the past are not as severe as its implications for manufacturers.

Anyone who has ever installed a hard disk knows that several steps have to be completed before a computer is able to write and read data to and from it. First of all it has to be partitioned to let the operating system, e.g. Windows or Linux know in precisely which sections data can be saved.

Several partitions can be created on a hard disk, meaning that a hard disk can incorporate several operating systems. Each operating system manages data in a specific format. This involves the creation of a file system such as FAT32, NTFS, extFS, AFS, GPFS and so on.

Demarcation

While blocks are addressed on the logical drive at partition level (a process also known as block storage), documents are stored on the file system (file storage). Blocks are numbered consecutively and are the prerequisite for the storage of data in file or object form. A Logical Volume Manager (LVM) can be used to combine several logical drives into a larger unit in cases where more storage space is required than is available on a hard disk. With a RAID controller, this happens at hardware level.

Operating systems format the blocks available to them. In the simplest scenario, if a document or file requires one or several blocks, the operating system notes the number of the occupied block in a file directory. In cases where more than one block is required, a pointer written in the start block during the storage process serves as a reference to the next block. This process continues until the stop sign is placed in the final block. It works in similar fashion to the table of contents in a book.

Higher storage functions

Snapshots: With higher storage functions such as snapshots, this table of contents serves as the basis for virtual freezing of the data storage.

It only takes a fraction of a second for any software to register the status of the

storage medium. Active applications are either stopped beforehand, so that they are unable to write any changes to the block, or the software registers the write access to the block and diverts new information to a free block. Alternatively, the old status can also be copied to a new block.

Thin provisioning: In the past, another storage management problem – expanding or reducing the size of a partition – could in most cases only be solved in idle state and with the help of external software. The difficulty here lies in identifying the unused blocks, extracting them from the partition and moving the partition boundaries accordingly to free up the resulting storage space. This used to be the IT equivalent of open-heart surgery. These days, most operating systems can cope with this task while up and running.

Several startups then had the idea of making the partition or logical drive as small as possible and expanding the provisioning of block storage at runtime. Thin provisioning was the answer: it solved the problem by making logical drives so big that the entire project data of a specialist department would presumably fit inside. This assumption was based on historical experience and had the effect of blocking considerable storage media capacity that was often never used despite costing money.

SDS – much is possible without hardware

This simple and incomplete outline of the essential prerequisites for the storage and, in particular, copying of data demonstrates all too clearly that a lot of programming work is involved in storing data on a storage medium and, where required, finding it again via file directory and folder structures or a search engine. However, many standard functions are “cast” in hardware by means of processors, microcode, so-called firmware and hard-wired circuits. This is a precise operation involving timed processes of predictable length.

With software-defined storage (SDS), the control of all storage and copying processes suggested above is supposed to take place solely via programs executed by a processor. The storage controller, NAS head, filer or whatever the controller is called is replaced by an x86 processor. This industry standard is the platform on which the basic functions can be called up via Application Programming Interfaces (APIs), and capacity and robustness are determined by the choice of processor used.

So SDS is not about making storage systems work better. The focus is more on avoiding the negative aspects: storage costs eat up too great a share of the IT budget when data growth is taken into account. The idea behind SDS is to turn storage into a commodity, while also paving the way to cloud services via OpenStack.

Asking whether SDS is the future of storage is in fact the wrong question. What is currently taking place is a change of form. Hardware processes are being replaced by software because the entire storage industry has no idea how they are to get into the cloud unscathed. Many failures have already resulted from the decision to rely on specialised hardware. SDS provides a certain measure of future security thanks to its programmable flexibility.

Walter Schadhauser

6 tips on choosing the right SDS solution

There are plenty of software-defined storage solutions to choose from. So sifting through the vast range of offers to identify the differences between the individual approaches and products is by no means an easy matter. But what important points should IT managers bear in mind when considering the solutions available on the market?

© Cybrain - Fotolia.com

Comparing apples and oranges is an impossible task. But that's just what you have to do when assessing SDS product offerings: there are significant differences between the various approaches and applications, and users have to examine them closely. So what exactly should you be looking out for? Six tips, pitfalls and criteria that you should bear in mind when choosing an SDS solution.

1. Not all SDS solutions are the same

Many manufacturers are currently competing in the SDS market with their ostensibly software-defined products. Their ideas, approaches and applications can differ widely, however. Even if they are all adhering to a similar basic concept, namely decoupling hardware from software, their means of doing so are sometimes worlds apart. Some of them virtualise storage arrays and combine these in a single resource pool, while others incorporate widely available servers based on standard hardware into their solutions. Because of the wide range of options on the market, it's wise to scrutinize each manufacturer's technical characteristics and USPs closely with a view to establishing whether the concept in question is compatible with corporate IT and whether the solution will be beneficial in the near future.

2. Low-cost standard hardware can prove expensive in the long run

For many years, storage managers have bemoaned the fact that proprietary hardware was just too expensive and not worth it in the long term, all things considered. Of course, SDS solutions that exploit the potential of standard hardware can initially appear attractive in terms of costs. Experts point out, however, that their use may call for a completely new infrastructure. The virtualisation of available

storage arrays may prove to be a cheaper solution.

3. The pitfall of integration

Integrating the selected SDS solution may also prove problematic. Although the majority of manufacturers test and validate their commodity hardware in advance to ensure that it is fit for purpose, it is still the internal IT department that is left to pick up the pieces when it comes to integrating it properly. And sometimes the human and financial resources simply aren‘t available to do that. Make sure that the method of implementing the standard hardware in the proprietary environment is spelled out clearly.

4. Testing hardware compatibility

Theoretically, SDS software stacks can run on any x86 server designed for universal use. But this is not the case in practice, as most software-defined storage solutions come with manufacturer-defined hardware compatibility lists (HCLs).

Thus, not every new hardware product – whether current hard disk model, flash solution or Host Bus Adapter (HBA) – will automatically interact with the SDS software. So problems are almost inevitable at some stage or other. It may be that the solution simply fails to detect the hardware, that disk failures occur intermittently, or that hardware problems are difficult to identify. Companies wishing to launch an SDS solution on their own should therefore exercise caution and test the prospective solutions for compatibility.

5. The application workload is crucial

The storage world likes to believe that storage basically serves a single purpose: to keep data available without disruptions, providing users and applications with quick and easy access at any time.

All too often a detailed analysis of the individual application workloads is simply forgotten. Given the availability on the market of such a wide range of SDS offerings geared to special workloads, analysts regard this as a common error.

They advise customers to pay greater attention to the workloads and to consider which solution is the most suitable as regards technical aspects and price-performance ratio.

6. Individual requirements are the benchmark

As mentioned above, SDS can be implemented in a variety of different ways. Since no single approach can cover all application scenarios, the planned utilisation scenario determines which solution is chosen. Thus, a dedicated array, a hyperconvergence platform or software installed on a server may serve to provide software-based data services, or function as native cloud services.

The applications, existing environment, internal resources and business objectives are the key factors in determining the right SDS solution. So the first thing to do is to establish precisely which storage functions are of primary importance, and thus indispensable. *Tina Billo*

The ultimate interface to the cloud

Standardised interfaces to the public cloud are a basic requirement if the hybrid cloud is to succeed. But it quickly becomes clear that this alone is not enough. The services of a company's own data centre should also be accessible from the cloud. This is far from easy to implement.

The IT industry lives with and feeds off its own hype. The hype proves that progress is being made and new opportunities are arising. This is currently borne out by the large number of supporters of the OpenStack Foundation, whose membership figures have risen from 20 to almost 2,000. The results of hard programming work are another indicator of progress – this can be measured in lines of programming code, nearly two million of which have now been written following a modest start with just 30,000 lines or so.

As if that wasn't enough, the now six-monthly OpenStack Summits, such as that held in Vancouver in May 2015, attract 6,000 interested participants. And last but not least, dozens of manufacturers are now competing to provide APIs (Application Programming Interfaces) for the various OpenStack components – and these are already being actively used by their customers.

Proceed with caution

To the longstanding industry observer, the formidable public relations machine is reminiscent of the hype formerly surrounding Java as an internet programming language: everyone uses it, many don't like it.

OpenStack is not at this stage yet. The protagonists still firmly believe that you only have to dismantle a data centre into its three separate layers (compute, network and storage) in order to be able to roll out this tried-and-tested infrastructure concept as Open Source software throughout the IT industry.

This already gives rise to a number of questions:

- Can such a complex infrastructure be operated with Open Source software?

- Should my company be worried that large IT suppliers will cease to offer many services as a result of OpenStack?
- How should the special requirements of business-critical IT systems be implemented and financed on an Open Source platform?
- What will happen to important business processes if unexpected disruptions cause IP connections to fail?

The OpenStack core

As mentioned above, the only difference with OpenStack is the sheer scale of the challenge. It's no longer a matter of local IT with metro clusters, but of IT for global dimensions – known as the cloud. The structure of an on-premise data centre is adopted. All the user has to do is to become familiar with the project names, with Nova used for the compute layer, Neutron for the network layer and Cinder as a general term for the storage layer. Cinder is further subdivided into:

- Cinder for the basis of all storage, block storage. Familiar media such as hard disk, CD, DVD and flash storage can be virtualised as blocks via specific parameters.
- Swift for object storage. This project serves to program the replication of objects. The object itself contains the information on its degree of importance to its user. Depending on its importance, it is copied twice, three or four times, or even more.
- Glance, an image service that provides images of virtual machines, to give just one example.
- Trove, the database service in OpenStack that is designed to act as a Database as a Service for SQL as well as NoSQL databases.

This short list just for the storage layer (other projects are in the pipeline for the other two layers) demonstrates that, even at data centre level, complex material has to be adapted for internet-capable operation. This is undoubtedly not a challenge that will result in a robust, highly available infrastructure within just a few short years.

Mixed operation

Statistically it takes between ten and twenty years to rebuild an infrastructure and make it universally operational. Linux, for example, is now approaching its 25-year anniversary. Even if you admit that OpenStack has enjoyed dynamic development and the first analysts are recommending that you get to grips with the concept as a matter of urgency, there is no reason to rebuild a functioning IT system with it in mind.

It's always the same with innovative products: you have to identify the elements that can be implemented productively and then analyse which of them can be integrated in your own IT system and whether they offer potential for the future. Because even if you concede that several thousand OpenStack developers can and do react dynamically to new business requirements, the reverse argument is always that code that just about works in one release will inevitably function differently in the next.

Walter Schadhauser

Different types of storage, countless manufacturers, confusing choice of products

Hardware is useless without software. In this respect, we have to ask ourselves what is really new or different about the concept of "software-defined storage". If you're looking for an answer to this question and also want to distinguish between offerings that tend to be more marketing-driven and those that really deserve to be known as SDS, you will need to invest plenty of time. Here's an overview of classification options and representatives of the individual types of SDS, as well as several manufacturers and their products.

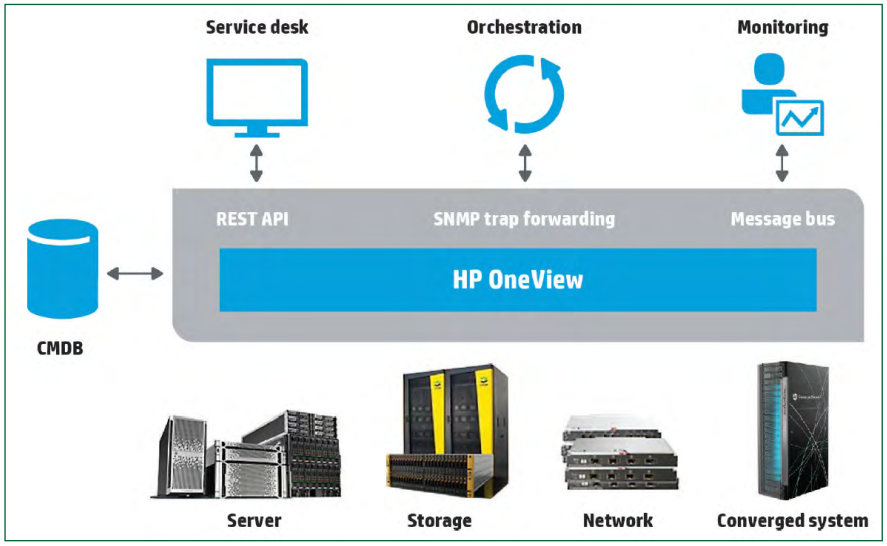

An example of a "vendor-defined" SDS solution: the HP OneView system management platform was developed for the central management of existing HP infra-structures. (Image: HP)

Dell, EMC, HP, IBM, Netapp, DataCore and Red Hat – who hasn't heard of these major IT players? But former startups such as Nexenta, Nimble Storage and Tintri, have now also made a name for themselves. On the other hand, it is harder to classify businesses of the likes of Aetherstore, Elastifile, Izumobase and Prophetstor with any degree of precision. This gives rise to several good quiz questions, such as: "How does the company Cleversafe (now IBM) earn its money?" A. Safes, B. Learning toys, C. Software-defined storage solutions or D. Bicycle locks? Of course we're talking about SDS here, so the answer is obvious. But without this background knowledge, the question might prove tricky for many.

This should not come as any great surprise, as SDS is a case of history repeating itself: an emerging market with the potential for strong growth is given a catchy buzzword. If it catches on, manufacturers – the established protagonists and a host of newcomers alike – try to jump on the bandwagon as quickly as possible and appropriate the term for themselves. However, the lack of standards means that what lurks behind the name – mere packaging for technologies and solutions – can differ greatly.

The spectrum ranges from traditional solutions geared to SDS all the way to completely newly designed products. The range of providers who are active in the vast field of software-defined storage is equally large.

"Vendor-defined" or "software-only"

The SDS solutions available today are categorised using various criteria, which depend largely on the perspective of the observer. The simplest approach is to determine whether the provider belongs in the hardware manufacturer or software manufacturer camp. The terms frequently used in this context are "vendor-defined" or "software-only" solutions.

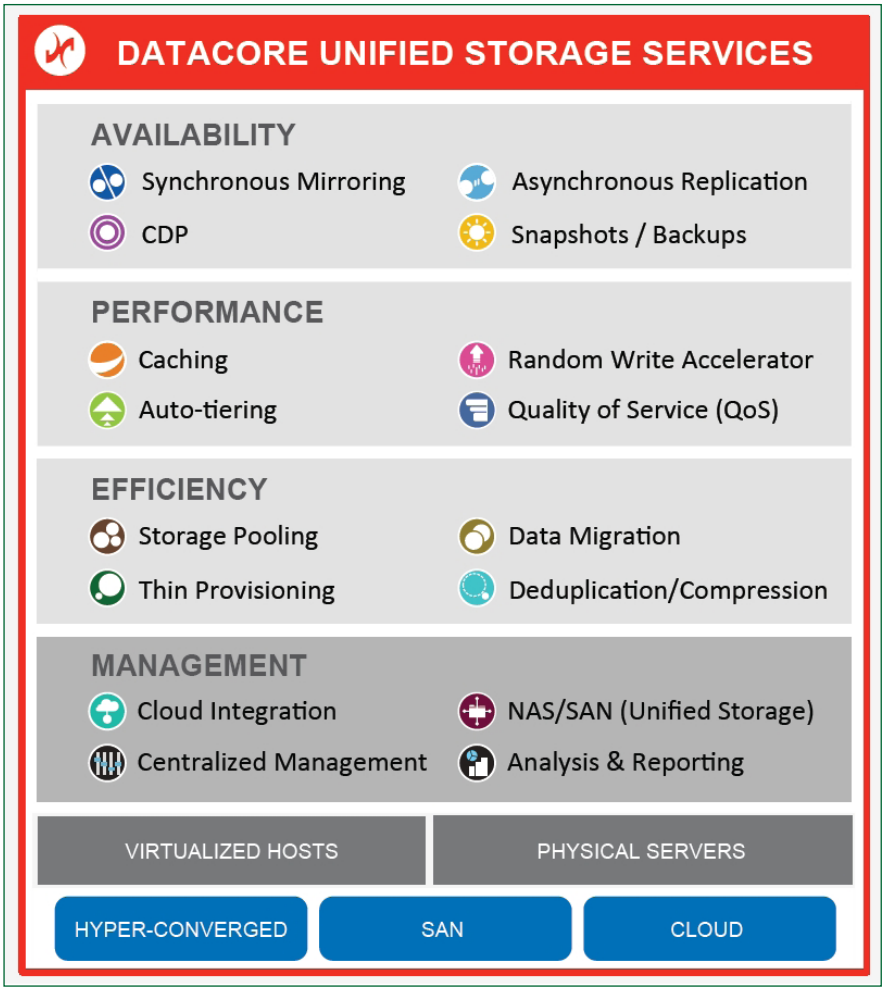

Datacore has without doubt been one of the leading SDS companies from the outset. SANsymphony-V is based on an exclusively software-oriented approach. (Image: Datacore)

The first category comprises traditional names such as Dell, EMC, HDS, HP, IBM and Netapp, as well as companies that spice up their arrays or appliances by means of supplementary SDS software. These include Nimble Storage, Nutanix, Simplivity, Tintri and Pure Storage. It is often suggested that this group only uses SDS as a vehicle via which to maintain its hardware business and sales of its often closed systems, without really offering "thoroughbred" SDS solutions.

On the other side of the coin are the manufacturers who have dedicated themselves to pursuing an exclusively software-oriented approach. Many of these, including DataCore, Falconstor, Red Hat and Veritas, have been active in the storage segment for many years. However, this group also includes several younger market participants – Atlantis Computing, Nexenta, Scality, Springpath and Starwind among them – who have now extended some of their offerings to include appliances based on commodity hardware. There are also many and varied software houses whose products are designed for special applications in the SDS environment. These may be based on Open Source, work in object-oriented fashion or as virtual appliances, be scale-out capable and perform special backup and disaster recovery, data reduction, analytics and performance management tasks. This category includes companies such as Aetherstore, Caringo, Infinio, Permabit, Rozo Systems, Virtual Instruments, Zerto and many others besides.

When implementation determines classification

Another option is to classify SDS solutions as hardware-centred and software-centred. And yet another is to take the form in which the products are applied as a basis. Thus products can be subdivided into Virtual Storage Appliances (VSAs), virtual controllers implemented on the server, storage-agnostic software operated at control level, and integrated platforms promising media-independent data storage and integrated, cross-site management.

Virtual Storage Appliances provide a broad range of virtualised storage services, not

restricted to SDS alone. However, since VSAs decouple data management from the physical systems and grant virtual machines shared access to resources bundled in storage pools, they are marketed by some manufacturers as SDS solutions. This category includes Falconstor, HP StoreVirtual, Nexenta and StorMagic, as well as VSAN from VMware.

Virtualised storage controllers implemented on physical or virtual systems on the server take advantage of the server processing speed instead of using the slower physical storage controllers. These products are therefore especially recommended when SDS is to be implemented in environments with high I/O performance requirements. Virtual controllers can be managed via a central management console. Companies such as Gridstore and Tintri are among the providers of this type of product.

SDS software solutions that operate at control level are suitable for use in heterogeneously designed storage landscapes. They combine the resources of physical storage systems in a single pool and provide all the tools required for management as well as a wide range of additional intelligent functions entirely independently of the basic storage. As the physical data path is not changed in any way, the orchestration services can be introduced without any significant modification of the storage infrastructure. Examples of this type of product include ViPR from EMC, the applications of the Spectrum family based on the features of the Virtual SAN Controller recently presented by IBM, and software from SDS pioneer DataCore.

Another category of product comprises platforms that render storage systems of various types SDS-capable across different sites. Local, remote or cloud-operated tape-, disk- or flash-based storage can be integrated. Tarmin‘s GridBank is a good example of a solution used for the end-to-end organisation of data processing and data management and also capable of supporting eDiscovery and Analytics.

Manufacturers and their offerings – a small selection

Manufacturer-specific SDS definitions, a variety of classification approaches and the sheer number of products still mentioned in connection with software-defined storage make it difficult to present a relatively clear overview of the entire offering. There‘s also the fact that several companies are already starting to shy away from using the term.

Thus the following table is intended solely for guidance and could doubtless be supplemented by a number of additional solutions.

Summary

It is currently extremely hard to grasp the sheer range of manufacturers who are vying for pole position in the SDS market or hoping to gain a foothold in it through the judicious labelling of their solutions. Thus the list below is merely a small, non-representative selection. Presumably, however, the SDS hype will kick-start a wave of consolidation (a development that can already be seen to some extent).

Tina Billo

Software-defined storage solutions

Conventional hardware manufacturers		
Manufacturer	**Product**	**Main focus**
Dell	Dell XC Series	Nutanix-based SDS hardware appliances
	Dell/Nexenta SDS	Nutanix-based SDS hardware appliances
	Dell storage with MS Storage Spaces	Pre-validated SDS storage systems
EMC	EMC ScaleIO	Server-SAN software
	EMC ViPR Controller	Storage automation software
	EMC ViPR SRM	Storage resource management software
	EMC ECS Appliance	SDS cloud storage platform (software or hardware appliance)
HDS	Hitachi Storage Virtualisation Operating System (SVOS)	Operating system
	Hitachi Virtual Storage Platform (VSP)	Storage systems
HP	HP Oneview	SDS system management platform
	HP Storevirtual VSA	Virtual Storage Appliance (VSA)
	HP Storeonce	SDS backup and deduplication software
IBM	IBM Spectrum Accelerate	SDS software platform
	IBM Spectrum Control	Data and storage management software
	IBM Spectrum Scale	Storage management software
	IBM Spectrum Virtualize	Virtualisation software
	IBM Spectrum Protect	Backup and recovery software
	IBM Spectrum Archive	Archiving software
	IBM Flashsystem	Software-defined flash storage and all-flash storage systems
	IBM Storewize	Virtualised and flash-optimised storage systems
NetApp	NetApp Clustered Data Ontap	Operating system
	NetApp Oncommand	Management software package including management integration tools
	NetApp FAS Series	Flash storage systems
	NetApp Flexarray	Storage virtualisation software
Providers of storage systems and appliances with an “SDS” flavor		
Manufacturer	**Product**	**Main focus**
Coho Data	Datastream	Hybrid and all-flash storage arrays
Gridstore	Gridstore Hyperconverged Infrastructure (HCI)	Hyperconvergent infrastructure software
	Gridstore Hyperconverged Appliance	Hybrid and all-flash arrays
Nimble Storage	Nimble Storage CS	Storage array family
	Nimble Storage SmartStack	Pre-validated reference architectures
Nutanix	Nutanix Acropolis	Virtualisation and mobility software platform
	Nutanix Prism	Infrastructure management software
	Nutanix Xtreme NX-Serie	Software-defined appliance
Pivot3	vSTAC	Operating system
	vSTACK & All-Flash Enterprise HCI Appliances	Hyperconvergent infrastructure appliances
Pure Storage	Purity	Operating system/flash-optimized storage software
	Flasharray//mm & Flasharray 400 Series	All-flash storage systems
Simplivity	Simplivity Omnistack	Hyperconvergent SDS infrastructure software
	Simplivity Omnicube	Hardware appliances
Tegile	Intelliflash	Operating system
	Intellicare	Cloud analytics-based support platform
	Intelligent Flash Arrays T3000 Series	Hybrid and all-flash storage systems
Tintri	Tintri OS	VM-sensitive SDS operating system
	Tintri Global Center	Central management software
	Tintri Vmstor T800 & T5000	Hardware appliances

Software-defined storage solutions

Software-only SDS providers incl. companies that have supplemented their offering with appliances		
Manufacturer	**Product**	**Main focus**
Amplidata (HGST)	Himalaya	Object-based storage software
Caringo	Caringo Swarm	Object-based storage software
	Cloudscaler	Cloud plug-in for Amazon S3
	Filefly	Global namespace solution
	SwarmFS for Hadoop	Hadoop connector for Caringo Swarm
	Swarm for OpenStack	Object storage for Swift
Ceph	Ceph	Open Source SDS software platform
DataCore	Sansymphony-V	SDS software platform
	DataCore Virtual SAN	Hyperconvergent SDS software
	DataCore Parallel I/O	Performance acceleration
Elastifile	Elastifile (Beta)	SDS software platform
Atlantis Computing	Atlantis USX	SDS software platform
	Atlantis Hyperscale	Hyperconvergent all-flash appliance
	Atlantis ILIO	In-memory storage optimization/VDI
Cleversafe (as of 10/2015 IBM)	Dispersed Storage Network (dsNet)	Object-based web-scale storage software
	dsNet Manager	Storage manager (appliance)
	Accesser	Storage router (appliance)
	Slicestor	Storage server (appliance)
FalconStor	Freestor	SDS software platform
	FalconStor Network Storage Server (NSS)	Storage virtualisation & business continuity
	FalconStor Continuous Data Protection (CDP)	Backup and recovery solution
	FalconStor Recovertrac	Disaster recovery solution
	FalconStor Optimized Backup & Deduplication	Backup and deduplication solution
Formation Data Systems	FormationOne Dynamic Storage	SDS software platform
Hedvig	Hedvig Distributed Storage Platform	SDS software platform
Infinio	Infinio Accelerator	Server caching software
Izumobase	IzumoFS	SDS software platform/scale-out NAS
Nexenta	Nexentastor	Open Source SDS software platform
	Nexentaedge	Object storage management
	Nexentaconnect	Storage acceleration
	Nextentafusion	Central management, reporting, automation and analysis software
Maxta	MxSP	SDS software platform
	Maxdeploy	Hardware appliances
Mirantis	Mirantis OpenStack Distribution	OpenStack-based SDS software platform
	Mirantis Unlocked Appliances	Hardware appliances
OpenStack	OpenStack	Open Source software
Pernixdata	Pernixdata FVP	Storage and VM acceleration software
Quobyte	Quobyte Unified Storage Plane (USP)	Distributed parallel file system
Red Hat	Red Hat Cluster Storage	SDS data management software
	Red Hat Ceph Storage	Object-based SDS software platform
Rozo Systems	rozofs	Open Source scale-out NAS file system
Sanbolic (Citrix)	Sanbolic SdX	SDS software platform
Scality	Scality RING	Object-based storage software
Springpath	Springpath Data Platform	Hyperconvergent software solution
StorMagic	SvSAN	Virtual Storage Appliance (VSA)

Software-defined storage solutions

Software-only SDS providers (continued)		
Manufacturer	**Product**	**Main focus**
Suse	Suse Enterprise Storage	Ceph-based Open Source SDS software platform
Swiftstack	Swiftstack Object Storage	OpenStack Swift- and object storage-based software platform
Veritas	Infoscale Enterprise	Storage management and business continuity SDS software platform
	Infoscale Storage	Central storage infrastructure management
	Infoscale Foundation	Online management of heterogeneous storage landscapes
	Infoscale Operation Manager	Role-based management interface
StarWind	StarWind Virtual SAN	Hypervisor-centric virtual storage software
	StarWind Hyperconverged Appliances	Hardware appliances
Other providers of SDS solutions		
Manufacturer	**Product**	**Main focus**
Aetherstore	Aetherstore	Distributed data store solution
Microsoft	MS Scale-out File Server	Cluster software
	MS Storage Spaces	Storage virtualisation technology
Permabit	Albireo Sanblox	Data reduction appliance
	Albireo VDO	Virtual data optimization software
	Albireo SDK	Software development kit for deduplication solutions
Prophetstor	Prophetstor Federator	Virtual storage controller
	Prophetstor Flexivisor	Storage hypervisor
	DR Prophet	Business continuity and disaster recovery software
Virtual Instruments	Virtualwisdom	Analytics/infrastructure performance management platform
VMware	Virtual SAN	SDS software for hyperconvergent infrastructures
	Evo:Rail	Hyperconvergent infrastructure appliance
	vsphere Virtual Volumes	Integration and management framework
	Site Recovery Manager	Disaster recovery software
Zadara	Zadara Storage Cloud	Storage cloud software
	Virtual Private Storage Array (VPSA)	Virtual Software Appliance (VSA)
	Zadara Container Service (ZCS)	Container service
Zerto	Zerto Virtual Replication	Disaster recovery and business continuity software for virtualised data centres and cloud environments

A milestone on the way to the software-defined data centre

A whole range of new SDS products is now available to help solve the storage problems of our age. But a close look at these products reveals that the vast majority of them focus on solving just one of these problems. The transformation required to implement the software-defined data centre, on the other hand, calls for an holistic approach that integrates server, storage and network. What can VM-aware storage contribute to this approach?

VM-aware storage – a milestone on the way to the software-defined data centre. Pictured here, Tintri's T880 hybrid storage. (Image: Tintri)

Server virtualisation constituted the first wave of the software-defined revolution in the data centre. It decoupled servers from hardware and transformed them into simple VMs or applications. Network virtualisation such as virtual switches, VLANs or other abstract concepts moved the data centre further in the right direction. However, the promise of an entirely Software-Defined Data Centre (SDDC) has to date continued to be hindered by legacy storage and its inflexible storage concepts. What is needed is storage that runs on intelligent software capable of solving the very real storage problems of our age. Tremendous data growth, smaller IT teams, virtualisation and cloud computing are the challenges that Software-Defined Storage (SDS) has to find answers to. A key element of the SDDC must surely be the capacity to operate seamlessly alongside the already virtualised part of the data centre, namely with VMs.

VM-aware storage – an emerging concept

What VM-aware storage actually means depends largely on who is defining it. Some SDS manufacturers who do not sell hardware themselves claim that it is sufficient to dovetail storage provisioning and management more closely with the hypervisor, thereby merely shifting the responsibility up to hypervisor level. Hardware manufacturers have even said that their arrays are already VM-aware because they can identify which VMDK files are associated with which VMs.

The possibility of linking big data and cloud applications with LUN and volumes, and using the performance data of these volumes to deduce VM performance is also frequently mentioned. VMware itself has introduced the VVol API in the form of vSphere 6, thereby enabling hardware manufacturers to adapt their products more closely to the vSphere workflow. None of these approaches go far enough. But what they all show is that the challenges facing us today cannot be solved using yesterday's concepts, and that the current innovations are largely being developed by innovative software providers. What the ultimate solution ends up running on seems to be of secondary concern. As with all challenges in the data centre, the transformation to the SDDC is more evolution than revolution, the long road marked by individual technological milestones. The realisation that storage has to be VM-aware is one of these milestones.

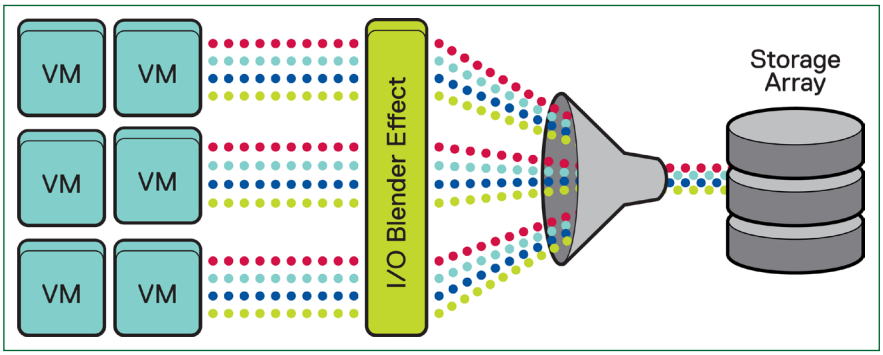

Virtualisation mixes I/O streams. Only VM-aware storage can cope with these mixed workloads. (Image: Tintri)

Which challenges can VM-aware storage solve?

The reality is that storage in today's data centre is still extremely simple in nature. IT managers worry about performance, costs, and simpler storage management – they are always on the look-out for new technologies that are capable of solving their problems more efficiently than the approaches used to date. But a close look at new developments reveals that the vast majority of them focus on solving just one of these problems. The transformation required to implement the SDDC, on the other hand, calls for an holistic approach that integrates server, storage and network. What can VM-aware storage contribute to this approach with a view to delivering at all levels while simultaneously solving all the three of the problems faced by IT managers?

Server virtualisation has revolutionised the data centre far more comprehensively than the mere decoupling of servers from hardware – it has completely changed the rules of the game. These days, over 80 percent of all workloads are already virtualised, and organisations are changing their perception of storage, thinking in terms of costs per VM, rather than costs per gigabyte. This approach is more relevant for them, as they are now judged on whether enough performance is available for every application at all times. Capacity is regarded as a given. However, in order to manage storage at VM level and work seamlessly with hypervisors, the storage itself has to be VM-aware right down to the smallest operations such as IO scheduling and the physical saving of files in the file system.

Before more intelligent storage systems were launched on the market, the old storage manufacturers had a simple solution to the performance problems of their customers: they just added more and more spindles to the system, with the sum of their performance data eventually combining to attain the desired performance objective. Flash storage helped cope with these performance problems to a certain extent.

Adding more and more flash to a system is not particularly efficient, however. It would be more important for the system's in-built intelligence in the form of algorithms to

ensure that practically all workloads arrive via flash, thus making optimum use of the flash that is already available. Only special environments that require top performance at all times justify the expense of All-Flash.

Mark Young, Senior Director, Systems Engineering, EMEA at Tintri.

Only software or appliance?

A whole range of new SDS solutions is now available to help solve the storage problems of our age. Ultimately data is always saved on hardware in some form – so you have to decide which package of software and hardware to opt for, even if you choose a software-defined option.

The purist's approach is to select software that, in principal, can run on any storage server or as part of a hyperconvergence platform in an appliance. Solutions of this kind, especially if they are paired with cheap standard servers, can result in huge cost savings, and there has been great hype surrounding this option over recent years. In reality, however, this approach has failed to satisfy expectations, as it proved to be extremely difficult to render a company's proprietary software compatible with any server. It quickly came up against its limits in the face of constantly changing hardware driver sets. Software manufacturers responded by starting to offer their software only on certified hardware, which of course takes the core concept behind their approach to absurd extremes. Other SDS approaches, such as those based on clustered file systems, have suffered massive scaling and data integrity problems, even though they are easy to install and use.

Some providers have taken the more pragmatic step of offering proprietary appliances optimized for their own SDS approach. The benefits of this concept are obvious: customers get a solution that works immediately and does not require any further attention.

VM-aware storage keeps the promises made by the SDDC

The SDDC may not quite be the finished article yet, but the parts of it that are available to date are already providing solutions for the most urgent problems of many organisations – especially if they want to invest in new technologies such as the cloud and virtualisation as efficiently as possible. Older approaches simply cannot compete, even if the systems appear to have solved the performance problem via flash in the medium term. The SDDC, on the other hand, is an holistic concept that requires the dovetailing of all of its elements. Virtualisation at the heart of this concept calls for all the other parts to be able to operate in VM-aware fashion. SDS can only succeed in a system of this kind if it uses flash performance intelligently, is easy to scale and is completely VM-aware. Moreover, every SDS approach has to support all hypervisors.

VM-aware storage is nothing less than an important milestone on the way to the software-defined data centre. Integration of server and storage virtualisation with intelligent software: that's VM-aware.

Mark Young