

WHAT IS SOFTWARE-DEFINED STORAGE?

ABSTRACT

Like any hot new technology, there are many diverging definitions of what software-defined storage (SDS) means, and many attempts by vendors to lay claim to the term in creative ways. In this white paper, we explore recent technology and operational trends, discuss the benefits of software-defined storage, and distinguish between the various types of SDS solutions available today.

March, 2017

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA 03/17 White Paper H15845

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

TABLE OF CONTENTS

EXECUTIVE SUMMARY	4
BACKGROUND	4
DATA CENTER TECHNOLOGY TRENDS	4
Virtualization: abstract, pool, and automate	4
Web-scale: simplify management and operations, reduce total cost of ownership (TCO)	5
TECHNOLOGY ENABLERS FOR SOFTWARE-DEFINED STORAGE	5
Networking	5
Compute	6
Economics	6
Hypervisor	6
Scale-out Computing	6
Flash Media	6
WHY SOFTWARE-DEFINED STORAGE?	6
OPERATIONAL AGILITY	6
Why Not Direct-Attached Storage?	6
PREDICTABLE PERFORMANCE AT SCALE	7
ELIMINATE SILOS	8
REDUCED TCO	8
WHAT DOES SOFTWARE-DEFINED REALLY MEAN?	8
ARE ALL SDS SOLUTIONS THE SAME?	8
TYPE 1 SDS: SOFTWARE DELIVERY	8
TYPE 2 SDS: STACK INTEGRATED	8
TYPE 3 SDS: INFRASTRUCTURE SERVICE	8
THE PROMISE OF SOFTWARE-DEFINED STORAGE	9
DEPLOY	9
MANAGE	12
REFRESH	12
CONCLUSION	13
RESOURCES	13

EXECUTIVE SUMMARY

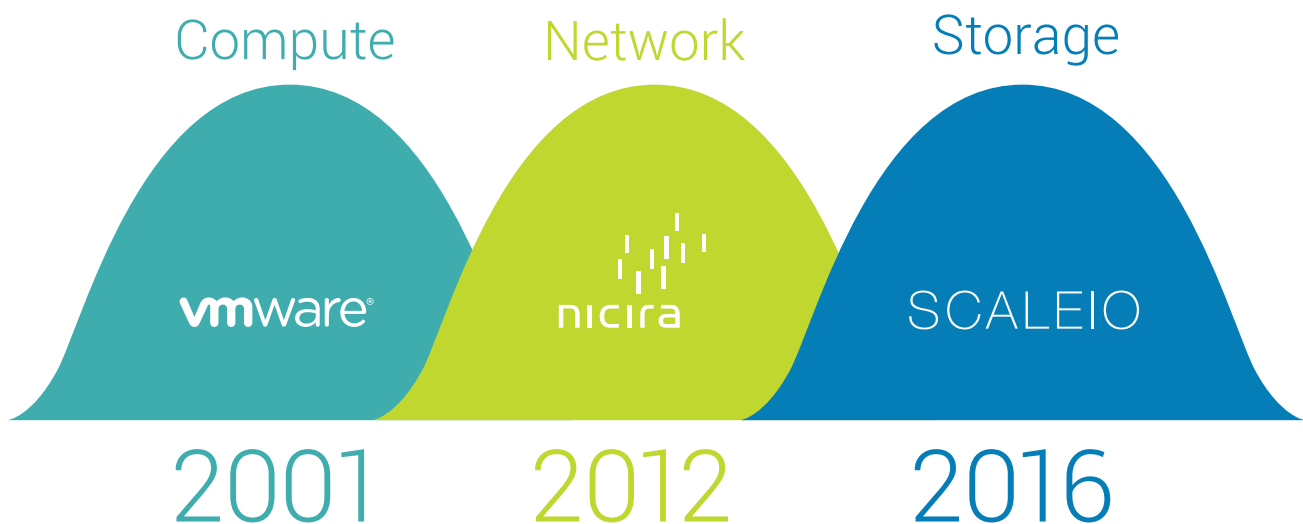
The benefits of server virtualization are well understood in the modern data center. By abstracting, pooling, and automating compute resources, companies have achieved significant savings. In the past decade, web-scale companies like Google and Amazon have demonstrated the ability to operate data centers with ruthless efficiency. Dell EMC ScaleIO Software-Defined Storage applies the principles of abstraction, pooling, and automation to local storage in standard x86 servers, creating a high-performance shared storage service without the need for conventional storage arrays. ScaleIO also enables a highly efficient data center operating model. Combined, these capabilities deliver a radical reduction in storage lifecycle costs. This white paper explores the benefits of abstraction, pooling, and automation of resources and the operational efficiencies of web-scale companies. Next, the paper discusses what software-defined storage really means and its impact in a modern data center.

BACKGROUND

DATA CENTER TECHNOLOGY TRENDS

Of the many new data center trends of the last decade, two of them stand out: virtualization and web-scale.

VIRTUALIZATION: ABSTRACT, POOL, AND AUTOMATE



Compute

VMware pioneered the concept of “abstract, pool, and automate” for CPU and memory resources. Simply put, rather than managing compute and memory resources on physical servers, VMware enabled the compute and memory across a number of servers to be abstracted away from the physical server itself. These resources could be combined together into a resource pool and then provisioned back to virtual servers in any quantity desired. This whole process was automated, which made it easy to reassign resources on the fly.

Virtualization led to huge efficiency gains in the data center by managing pools of resources instead of individual hosts and led to operational gains by abstracting the hardware layer. This enabled organizations to provision servers faster, reduce or eliminate vendor lock-in, and increase server utilization. It should be noted that even though the virtualization and sharing of CPU and memory resources are not across the entire data center (they are grouped in clusters of servers, currently a maximum of 64 in VMware environments), there are significant gains to be had and this is the de-facto data center standard today.

Network

CPU and memory have been abstracted, pooled, and automated for over a decade, but we are only beginning to see fully virtualized networking. Software-defined networking is a similar approach to abstraction, pooling, and automation of the underlying physical network resources to enable dynamic and scalable networking infrastructure. Nicira is an example of this trend.

Storage

There have been different approaches and attempts to virtualize storage over the last few years and this topic is our primary focus. Why hasn't storage virtualization occurred at the same pace as server virtualization? This question is explored below in the section titled Technology Enablers for Software-Defined Storage.

WEB-SCALE: SIMPLIFY MANAGEMENT AND OPERATIONS, REDUCE TOTAL COST OF OWNERSHIP (TCO)

Web-scale companies built a new paradigm for achieving high efficiencies based on the following fundamental principles.

Standard components

Web-scale companies tried to use traditional architectures and found them too complex and limiting. They understood the limitations of specialized hardware and needed software to control everything (software-defined). They also understood the value of driving everything in their data centers down to a few standardized components and the need to deploy services quickly and with ruthless efficiency. Therefore, these companies treated hardware as disposable or easily replaceable. Individual component or server failures did not bring down the entire service.

No data migrations, ever

When the capacity or performance of traditional storage arrays is fully utilized, the only way to upgrade is to buy bigger and faster arrays. This typically involves moving massive amounts of data from old arrays to new arrays. Web-scale companies and today's enterprises have data sets that are too large for traditional data migrations. Simply put, they cannot afford to take weeks or months to plan and execute a migration, nor can they spare the personnel required to coordinate the process with application owners. Data migrations are not only complex and time consuming, but also expensive and risky to application availability. But web-scale companies simply add new servers when needed and remove old servers with no data migration. They intelligently move or replicate data non-disruptively off of the servers they intend to decommission and balance in new servers as they are deployed. The entire data footprint is never moved in a single large operation.

Scale-out

Centralized shared storage such as SAN or NAS did not fit the needs of the web-scale companies. These traditional purpose-built storage systems are designed for a specific purpose with limited scalability, resulting in storage silos. They are also more complex and time consuming to deploy and manage. Web-scale companies need something better—rapid deployment of compute and storage on demand. They build scale-out systems that grow as needed and are under full software control. Instead of the traditional scale-up architectures and forklift upgrades that result in long procurement cycle times and massive data migration time and costs, these companies roll in new hardware when needed and decommission old hardware at will—achieving true elasticity.

Predictable performance at scale is another important characteristic of web-scale because it is almost impossible to anticipate growth. Another advantage of the web-scale operating model is that operating or administration costs do not increase as the number of systems do. It is common in web-scale companies to have an administrator-to-server ratio of one administrator to thousands of servers.

Software-defined storage (SDS) applies the above principles (abstraction, pooling, and automation of resources, and a web-scale operating model) to storage. Dell EMC ScaleIO brings the operational efficiencies that server and virtualization teams have enjoyed for years to the storage team—and delivers these efficiencies at *data center scale*. This allows the storage team to operate IT with the efficiency of a web-scale company, regardless of the scale of the particular organization. In other words, even a company with small-scale IT can leverage the web-scale efficiencies for storage.

TECHNOLOGY ENABLERS FOR SOFTWARE-DEFINED STORAGE

Why hasn't software-defined storage become a reality until just recently? There are a few key technologies necessary for SDS to become a reality in the data center. They are outlined below.

Networking

Ethernet is overwhelmingly used to interconnect servers within and across data centers. Ethernet speeds have risen from 1Gbps to 10 Gbps and we are beginning to see adoption of 25, 40, and 100 Gbps infrastructures. Coupled with bandwidth and reliable low-latency links, Ethernet has enabled SDS vendors to deliver high-quality and high-performance storage services to applications.

Compute

Advancements in processor design and manufacturing enabled server vendors to offer powerful multi-core processors and plenty of RAM. Modern servers are capable of running not just the client workloads or applications but also distributed storage services, eliminating the need for special purpose storage arrays. This collapse of compute and storage into a single server is what we generally refer to as hyper-convergence. Because of the collapse of the server and storage layers into one, procurement of hardware is simplified as well. IT organizations traditionally have separate server and storage teams, but with the rise of hyper-convergence, these organizational barriers can be broken down.

Economics

In the last decade, costs for CPU, RAM, and network bandwidth have been drastically reduced. This gave rise to the ubiquity of powerful servers and high-speed network interconnects in the data center.

Hypervisor

Even a few years ago, applications requiring massive performance were deployed on bare metal servers (i.e., servers without a hypervisor). Hypervisors are becoming more mature and, by pooling resources, can now satisfy the performance and I/O requirements of intensive applications.

Scale-out Computing

Traditionally, a server was 'scaled' by adding more CPU or memory. Advancements in scale-out computing enabled additional resources by adding more servers or nodes to an existing cluster. These typically involved standard hardware components. One big advantage of scale-out is that an application won't suffer an outage when one server or a component fails. Coincidentally, web-scale companies such as Google and Amazon made scale-out architectures and self-healing systems popular.

Flash Media

Before the wide-scale availability of flash media, thousands of spindles were needed to satisfy I/O intensive applications. Today, a single SSD can deliver the equivalent performance of hundreds of HDDs with much faster response times. Also, flash media costs have become highly competitive compared to performance-oriented hard drives. These two trends together mean that a large monolithic traditional storage array with thousands of spindles is no longer a necessity. A few SSDs sitting inside a standard x86 server can offer the same or better performance than a large traditional hard drive-based storage array.

WHY SOFTWARE-DEFINED STORAGE?

In this section, we will connect the dots between data center technology trends and the technology enablers discussed above. To recap, VMware demonstrated the value of pooled resources for compute and memory. In fact, the pooling of the compute and memory resources drove the increasing adoption of shared storage arrays, which are really a form of pooled storage resources in which a number of HDDs and SSDs act together as one system and serve a multitude of virtual machines and physical servers.

OPERATIONAL AGILITY

Web-scale companies showed the value of reducing hardware costs and deploying the least amount of variation in the data center. To accomplish this, companies pick their favorite server vendor, standardize on a few server models and configurations, and roll them out until a newer or better model supersedes it. There is no room for external shared storage arrays in this model because they are different hardware, can't be easily and quickly ordered, deployed, or decommissioned, and they almost always involve massive data migrations. To operate with web-scale efficiency, companies need something that is built entirely from the internal storage configured in the servers. Recall also that in order to achieve operational agility, you cannot rely on special hardware. In addition, the standard hardware needs to be driven by software and automation.

Why Not Direct-Attached Storage?

Direct-attached storage (DAS) does not allow you to share storage resources. Recall that the abstract, pool, and automate concept applies for storage as well and mandates that storage resources are shareable across the data center to gain efficiencies. In addition,

managing hundreds or thousands of DAS islands leads to extremely poor resource utilization, a problem that has already been solved by shared storage.

SDS allows for abstraction of the local storage from each server and pools all of it together just as VMware does for compute and memory. You then allocate out of the pool only the resources each application needs. In addition, software control enables you to scale those resources—for performance and capacity—as needed. Storage is completely abstracted from the underlying hardware. You can also have a combination of server vendors, old and new server models, and different media types (HDD or SSD).

In VMware virtual server environments and in web-scale companies, lifecycle management (deploy, manage, and refresh) is critical to operational efficiency and cost savings. Once a pool of resources is running on as few variations of standardized servers as possible, you need a way to roll the infrastructure non-disruptively. Again, VMware makes this possible for compute and memory so long as there is external shared storage underneath—you can move your VMs non-disruptively to a new host and decommission the old host. Since with SDS you are no longer using external shared storage arrays, the SDS itself provides the capability to roll the hardware. The hardware is just standard x86 servers and may even be the same hardware the applications are running on if a hyper-converged deployment is chosen. SDS automates the processes of moving data off of servers that will be evicted and accepting new servers into the SDS pool and keeping it balanced. There is no such thing as a wholesale data migration. Ever.

PREDICTABLE PERFORMANCE AT SCALE

Scale-out distributed architecture allows predictable performance at scale. There is no need to buy a large over-provisioned storage array to meet the future demands of business. Instead, you scale your performance by adding standard x86 servers when needed. When implemented correctly, the result is a near linear performance growth when adding servers.

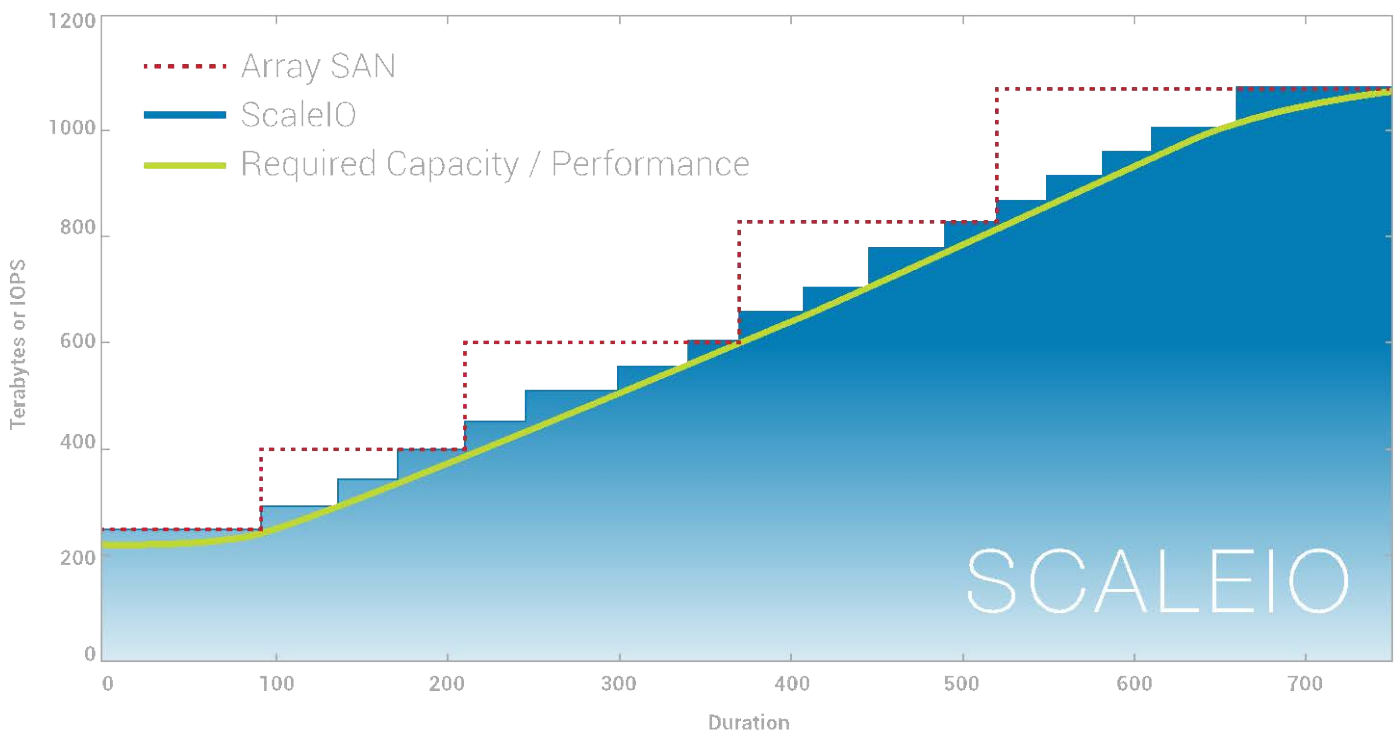


Figure 1. A timeline showing an increase in capacity or performance requirements (green). To meet these requirements using SAN (red), capacity or performance is overprovisioned as each new array or storage shelf is added. To meet these requirements using ScaleIO, overprovisioning is eliminated by managing growth at a fine granularity.

ELIMINATE SILOS

“Abstract, pool, and automate” applies for storage as well and those storage resources must be shareable across the data center to gain efficiencies. The pooling of the compute and memory resources is what drove the increasing adoption of shared storage arrays, which are really a form of pooled storage resources where a number of HDDs and SSDs are pooled together as one system and serve a multitude of virtual machines. Unfortunately, these shared storage arrays (NAS or SAN) are built for a purpose and this often creates silos of storage that cannot be shared at data center scale. Additionally, DAS does not allow you to share the resources across each server. Managing hundreds or thousands of DAS storage islands is not an option either.

REDUCED TCO

There are several ways that software-defined storage at data center scale delivers reduced TCO. First, simplification of the hardware platform to a few standard components across the data center results in lower procurement costs by reducing SKUs and managed vendors. Storage can be deployed and provisioned much faster than the weeks or months it typically takes with traditional arrays. Also, scaling is easier because you simply add servers when you need them and there are no expensive and long data migrations. Since SDS allows for full abstraction of underlying hardware, it is easy to procure and add newer server models when desired and decommission old hardware. In addition, SDS simplifies storage management via automation and API-based software. New application deployments are agile and thus provide revenue growth for businesses.

Further, hyper-convergence can result in higher utilization and thus reduced footprint, power, and cooling costs. On the other hand, some organizations have server and storage organizational silos that cannot be overcome immediately, therefore a flexible SDS solution must not mandate hyper-convergence but offer the flexibility to deploy in either hyper-converged or a traditional two-tier (Server SAN) topology and provide a seamless path from Server SAN to hyper-converged or vice versa.

WHAT DOES SOFTWARE-DEFINED REALLY MEAN?

ARE ALL SDS SOLUTIONS THE SAME?

In a word, no! Software-defined-storage is an over-used term that many vendors have adopted because it is a hot industry topic. We define software-defined storage as more than just delivering an array as software. In fact, there is no room for special purpose hardware to perform specific functions in this definition. Everything is done by software using standard x86 hardware that is readily available.

ScaleIO builds on VMware’s abstract-pool-automate model, allowing a cluster of one thousand physical hosts to share a single striped pool of storage. ScaleIO software-defined storage runs on standard x86 servers. It discovers the local HDDs and SSDs in each server, abstracts them away from each individual server, pools them together as a shared resource, and automates the provisioning of storage performance and capacity out of that pool back to individual applications. Just as VMware made computing more efficient and more manageable at scale, ScaleIO software-defined storage is making storage more efficient and more manageable at scale.

To further help understand how to think about software-defined storage, consider the following taxonomy:

TYPE 1 SDS: SOFTWARE DELIVERY

Type 1 SDS is making a software-only version of a hardware array. These products began life in an array form factor. Now they are offered in “software only” form, typically as a VM. This is really a packaging mechanism as opposed to a new way of doing things, and generally no Type 2 or Type 3 SDS capabilities are offered.

TYPE 2 SDS: STACK INTEGRATED

Type 2 SDS is software that abstracts and pools storage within a single cluster of servers, and generally within the confines of a specific hypervisor or operating system, designed to provide storage services within that system only. Type 2 SDS systems emulate some of the functions of an external array. Think about Type 2 SDS as analogous to having a dedicated physical storage array for each compute cluster.

TYPE 3 SDS: INFRASTRUCTURE SERVICE

Data center-scale SDS solutions were originally conceived and developed to be software-only and to provide storage services across an entire data center. As such, they are not tied to any specific hypervisor, operating system, hardware, or deployment topology. A purpose-built SDS abstracts and pools storage within and between clusters of virtual servers as well as for stand-alone physical servers. It has full flexibility between physical and virtual operating systems, as well as between traditional two-layer (where storage is a

distinct set of resources from compute) and hyper-converged (where storage and compute combine onto the same set of servers) operating models. Data center-scale SDS solutions provide value not just by replacing the traditional functions of an external storage array, but through automation at the data center level that provides additional functionality and flexibility.

ScaleIO is Type 3 SDS. A helpful way to think about the benefits of Type 3 SDS such as ScaleIO is to consider server virtualization. In VMware vSphere, while you abstract and pool all resources in the cluster, any single VM can only be assigned the resources of a single host. In other words, you can't create one large VM that utilizes the compute from a half-dozen ESX hosts. With ScaleIO SDS, there is not a limitation. A single VM can get storage capacity and storage performance that far exceeds what is available on a single ESX host. In this respect, ScaleIO SDS is actually more advanced than the server virtualization it is supporting. In fact, ScaleIO abstracts and pools the resources from multiple vSphere clusters, clusters running other hypervisors, and servers that are not virtualized (Figure 2). Non-homogenous server operating systems and hypervisors are still a reality in today's data centers. This is a completely unique capability that allows automation of storage at data center scale.

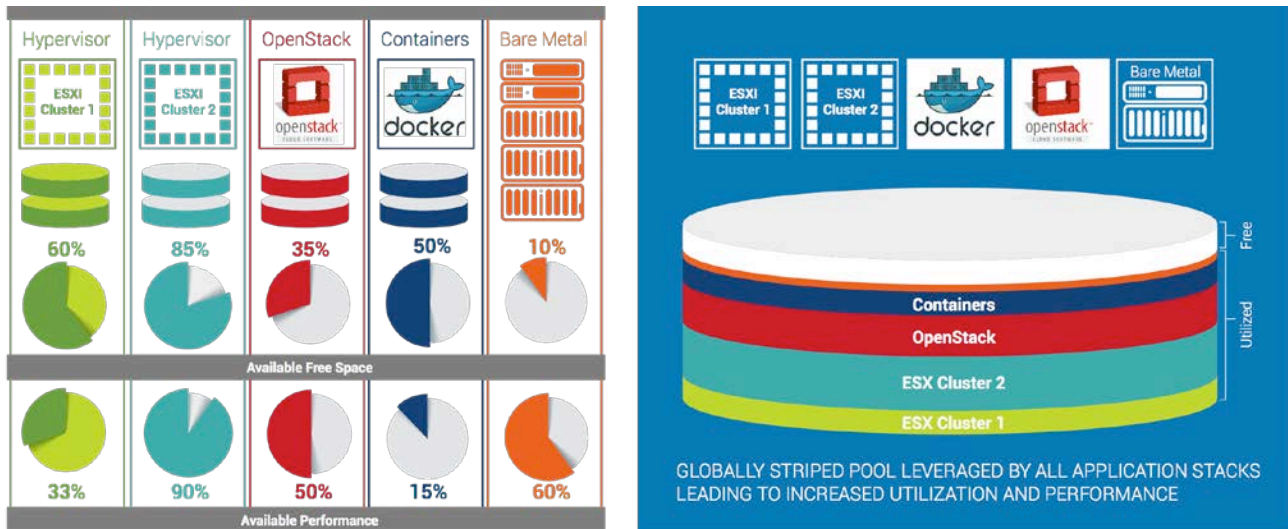


Figure 2. Left: Disparate ESX and Hyper-V clusters, bare metal and container storage each with its own silo of Type 2 SDS storage. Resources may only be provisioned within each silo. Right: A Type 3 shared ScaleIO resource pool. All of the storage capacity and storage performance is available to all hosts across all hypervisors and operating systems.

THE PROMISE OF SOFTWARE-DEFINED STORAGE

In order to understand what SDS can deliver, let us approach this from the perspective of a typical enterprise IT organization. It is helpful to think about this through the storage lifecycle phases—deploy, manage, and refresh.

DEPLOY

In this phase, the IT organization does the procurement, preparation, and initial set-up of new storage resources.

- Minimized hardware costs:** The SDS model minimizes hardware costs because everything is a standard x86 server. The entire data center can run on just a handful of standardized server components, with no proprietary hardware. This makes it very easy to manage the growth of resources because the operational run books for the data center become very streamlined. The model also reduces initial capex costs due to the elimination of the specialty storage hardware.
- Reduced operational complexity:** The SDS model also reduces operational complexity in the data center because management and operations (M&O) teams only deal with one or two standardized servers builds. They become efficient at rolling out new servers and removing old ones because there is little or no variation, just like the web-scale companies.

- **Easy scaling:** The scaling unit is minimal. Purchases are not large or expensive storage arrays out of necessity. Scaling can be done on-demand for only as much resource as is needed, scaled in bite-size chunks based on the server configuration. Servers can be rolled out one server at a time if desired. Gone are the days of over-provisioning which was typically done with traditional arrays to circumvent the long deployment cycles.
- **Minimized networking costs:** Networking costs are also minimized, as there is no Fibre Channel SAN fabric, HBAs, or expertise required. Everything runs on 10GigE (or 25G/40G/100G), which is the standardized data center fabric today.
- **Management automation:** ScaleIO management automation quickly discover new servers, joins them to the cluster (both the ScaleIO storage system itself as well as the vSphere cluster(s), if applicable) and begins using them in minutes. You eliminate the complex and multi-week process of purchasing, receiving, installing, and deploying a storage array. Instead, just receive the server, rack it, and add it to the cluster.
- **Small footprint:** ScaleIO software is lightweight in terms of CPU and memory usage, leaving the majority of the resources for use by applications. In addition, the standard x86 servers are not like the bulky traditional storage arrays. Footprint demands in terms of rack density, power, and cooling are low, making data center planning for accommodating growth easy.
- **Flexible topology:** With ScaleIO, you don't have to decide at deployment time whether you choose to run in hyper-converged or two-layer Server SAN topology. You can start with one and change it on the fly. ScaleIO even allows running in mixed mode with co-existence of Server SAN and hyper-converged in the same cluster. For example, if your organization has separate server and storage administration organizations, ScaleIO offers the flexibility to start with a traditional two-layer Server SAN topology as shown in Figure 3. In this topology, compute and storage services act as services by different sets of nodes. With ScaleIO, you can seamlessly move to a hyper-converged deployment topology as shown in Figure 5. ScaleIO also gives you a choice of running in a mixed mode (Figure 4), where some servers are in hyper-converged and others are in a two-layer deployment configuration. Thus, ScaleIO gives customers the power of choice so they can choose when and how fast the transition to hyper-converged or two-layer can take place.

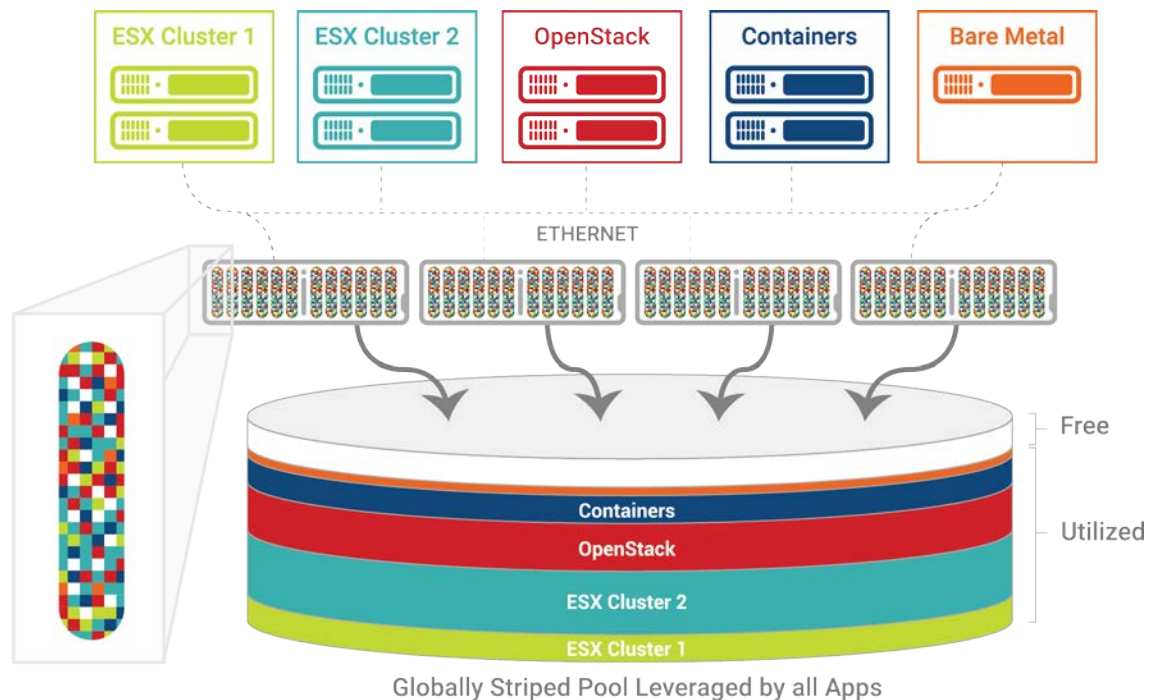


Figure 3. Traditional two-layer deployment with separation between compute servers shown on the top and storage servers shown in the middle. All of the storage devices contribute to form a global storage pool serving all of the compute clusters regardless of the operating system or hypervisor. ScaleIO balances the data from all the applications across all the storage servers leading to a perfectly balanced cluster eliminating performance hot spots.

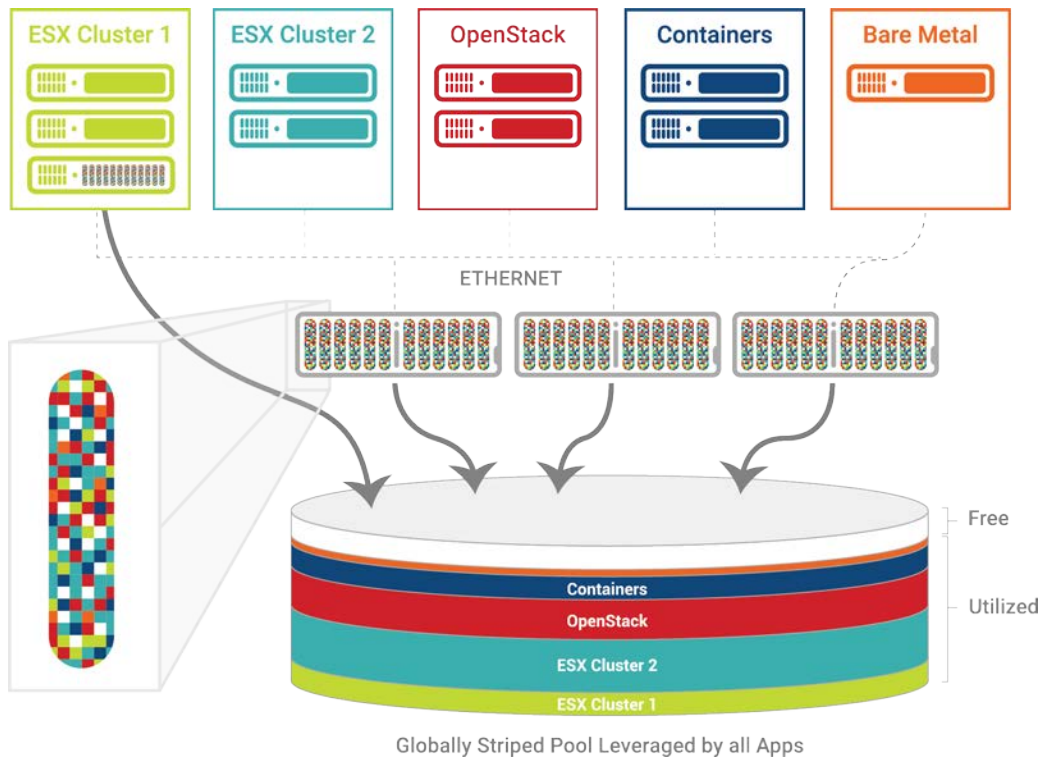


Figure 4. Mix of two-layer and hyper-converged deployment. For example, ESX Cluster 1 now requires more compute power. Instead of buying another server, ScaleIO can configure the first storage server on the left in Figure 3 to also run compute services. This server is moved into the ESX Cluster 1 box to how that it is now serving compute for this cluster. This server is continuing to serve storage services to all the application stacks as before.

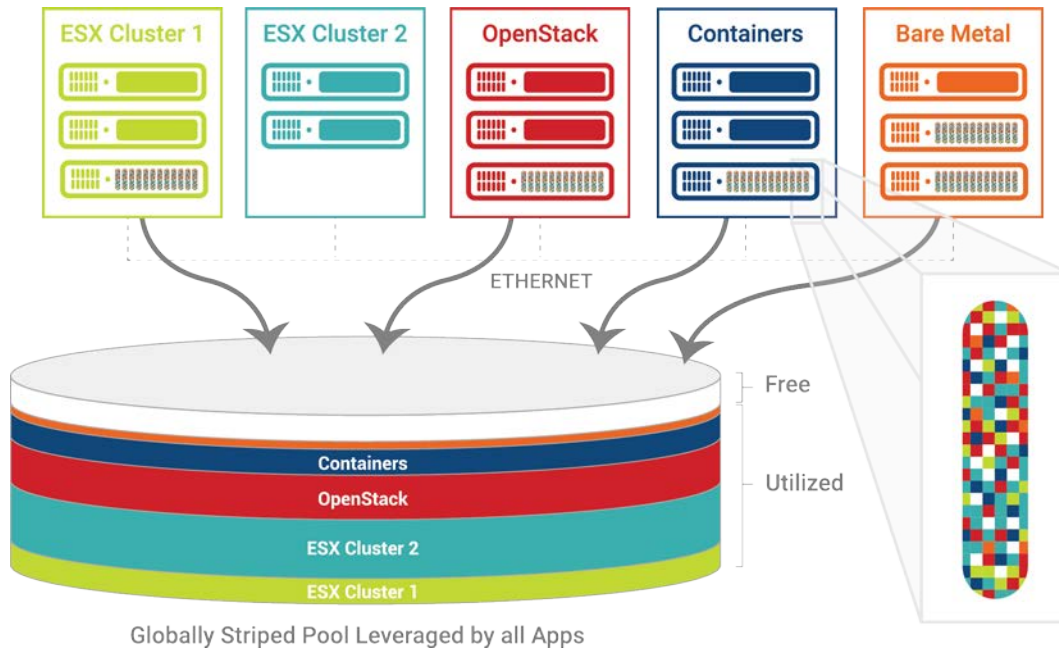


Figure 5. Transition to hyper-converged deployment is now complete. The three storage servers from Figure 4 are now also serving compute to OpenStack, containers and bare metal stacks respectively while serving storage to all the application stacks as before.

MANAGE

In this phase, the IT organization provisions storage to applications, adjusts to changing performance and/or capacity requirements, and deals with inevitable component failures.

- **Grow as you need:** ScaleIO is truly a data center-scale SDS. ScaleIO can start on as few as three servers, but has proven scale to thousands of servers. Storage is managed as a single service for the entire data center, rather than in silos.
- **Performance flexibility:** When ScaleIO abstracts and pools all of the HDDs and SSDs across the servers it is running on, there is a tremendous amount of performance potential available. You can allocate as much or as little performance and capacity as needed to each host—and elastically change it on demand—all under software control. ScaleIO automatically distributes the workloads it is running so that all applications get ideal performance.
- **Seamless tiering:** ScaleIO is designed to seamlessly move data between tiers. For example, an application that sees an increased performance demand can be non-disruptively migrated from an HDD pool to an SSD pool.
- **Workload flexibility:** Workload profiles change over time due to changing business. You can dynamically and non-disruptively allocate more HDD or SSD devices or servers to increase performance for a given workload or application.
- **Data protection:** The data protection scheme in ScaleIO is different from traditional RAID solutions and is designed around the idea that minimizing the window of vulnerability during hardware failures is of the utmost importance when designing at data center-scale. Drive rebuilds can be completed in seconds and entire server failures containing multiple drives can be healed in minutes without disruption.

REFRESH

In this phase, IT organizations manage the lifecycle of the SDS by bringing in new servers for growth or replacement of old servers and decommissioning and evicting old servers that have failed or have exceeded their useful life.

- **Seamlessly add new servers:** It is easy and seamless to add new hardware and ScaleIO does the rest. ScaleIO automatically balances the capacity and performance across all the servers. No manual intervention is needed.

- **Decommission old servers:** If servers are old and need to be retired, you simply indicate in ScaleIO that you would like to remove the server. ScaleIO handles redistributing its data among all remaining servers—all without impact to application availability.
- **Add new media:** Similar to adding new servers into a cluster, it is also easy to add new media (SSDs or HDDs) into the servers and let ScaleIO do the rest.
- **Eliminate complex, time-consuming, and expensive data migrations:** Just as VMware allows compute resources to be removed and replaced without application disruption, since ScaleIO is abstracting, pooling, and automating the storage inside those servers, the process of rolling the storage is similarly seamless.

CONCLUSION

Until today, storage has been trapped in silos in the data center, creating islands of under-utilized and over-provisioned storage pools and leading to vast inefficiencies in deployment, management, and refresh of old hardware. Web-scale companies have taught us that you can achieve operational efficiencies by standardizing on a few hardware components and eliminating data migrations forever.

With ScaleIO software-defined storage, we can now do for storage what VMware did for compute and memory years ago. By abstracting the underlying storage hardware, we are able to pool all the storage resources and allocate to applications as needed. Since the scaling unit is small and is a standard x86 server, we can cut down on long procurement and deployment cycles and grow incrementally as needed. When more performance or capacity is needed, simply add new servers and ScaleIO automatically balances the resources across the entire cluster. When it's time to decommission old hardware, simply evict the hardware from the cluster. Storage lifecycle management and operations (deploy, operate, refresh) become simpler and more efficient.

Not all software-defined storage is the same. To realize the true promise of software-defined storage, you need a combination of abstract, pool, and automate for *any* hypervisor, operating system, and hardware with web-scale management and operational efficiencies and economics—all applied at *data center-scale*. Type 3 SDS is best suited for medium to large enterprise and service providers with large amounts of servers, large amounts of data, stringent performance requirements, and rapid growth.

Software-defined storage is technology to be considering today. If you remember a few years ago, all-flash was that technology and it still is. Dell EMC is leading the SDS market—the technology is already mature and in production at multi-PB scale with some of the world's leading companies and service providers.

Think about SDS like Tesla electric cars. We all know the world is heading to the eventuality of electric cars. It doesn't mean everybody will buy a Tesla today. But some people will. And the world has come to understand that electric cars are the way of the future. Software-defined storage is in this phase today. Some people are "driving" it already, and they are doing so in increasing numbers with great results.

We encourage you to take ScaleIO for a free test drive and experience it yourself—[ScaleIO Free and Frictionless](#).

RESOURCES

ScaleIO: <http://www.emc.com/storage/scaleio/index.htm>

ScaleIO User Guide: <https://community.emc.com/docs/DOC-45035>