Taxonomy

IDC's Worldwide SBS, SDS, and FOBS Storage Solutions Taxonomy, 2018

Eric Burgener, Laura DuBois, Ashish Nadkarni, Amita Potnis, Eric Sheppard, Andrew Smith

IDC'S WORLDWIDE SBS, SDS, AND FOBS STORAGE SOLUTIONS TAXONOMY

FIGURE 1

SBS, SDS, and FOBS Storage Solutions Primary Segments

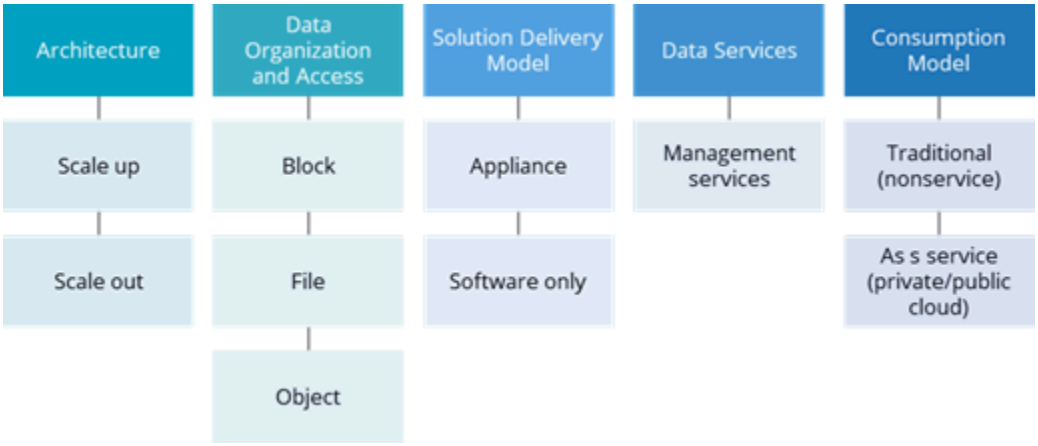

Source: IDC, 2018

SBS, SDS, AND FOBS STORAGE SOLUTIONS TAXONOMY CHANGES FOR 2018

Starting in 2018, IDC is consolidating the server-based storage (SBS), software-defined storage (SDS), and file- and object-based storage (FOBS) taxonomies into a single storage solutions taxonomy. IDC believes that this new approach will provide a simplified understanding of software-defined storage solutions and help customers understand how these three segments of the market correlate. Separate taxonomies that inform how IDC defines these markets will cease to exist. This taxonomy replaces three former IDC documents:

- *IDC's Worldwide Server-Based Storage Taxonomy, 2017* (IDC #US42308717, February 2017)
- *IDC's Worldwide Software-Defined Storage Taxonomy, 2017* (IDC #US42258917, January 2017)
- *IDC's Worldwide File- and Object-Based Storage Taxonomy, 2017* (IDC #US42376717, March 2017)

March 2018, IDC #US43579118

There are no substantive changes to the taxonomies as part of the consolidation, although in integrating the SDS and FOBS taxonomies we have moved to a consistent set of primary segment descriptors, which now include architecture, data organization and access, solution delivery model, data services, and consumption model. We have also slightly evolved some of the definitions to ensure that they can now apply to all three consolidated taxonomies (SBS, SDS, and FOBS).

TAXONOMY OVERVIEW

In combining the taxonomies, IDC aims to simplify the common and distinguishing features of SDS and FOBS solutions, both of which run predominantly on SBS hardware. The following sections outline in detail the many segments IDC includes as a part of both SDS and FOBS, and how they interrelate with SBS:

- The Definitions of SBS, SDS, and FOBS
- Two Key SDS and FOBS Solution Building Blocks
- Primary Segment Descriptors
- Market Segments

DEFINITIONS

The Definitions of SBS, SDS, and FOBS

Server-Based Storage

SBS hardware is defined as a commodity off-the-shelf (COTS) server that consists of a motherboard that hosts the CPU, memory, buses, and other devices that support storage, networking, and management interfaces. Additional components and/or devices can be attached to the base hardware to provide a functionally differentiated platform. Internal storage devices are used to persist data. A low-level software program placed on the read-only memory allows the system to boot and access the operating environment (e.g., hypervisor and operating system) residing on persistent media. The software-defined storage controller software (SDS-CS) is then installed on top of an operating system (which may be running on bare metal or in a virtual machine or container) to create an SDS or FOBS solution.

IDC identifies six different classes of SBS hardware: commercial systems (tier 0); configured systems (tiers 1–3), with increasing levels of limited customization; and customized systems (tiers 4 and 5). Differentiation between the tiers is based on the level of customization in the hardware, how the servers are procured, the assembly channels, and the types of internal components. To meet the definition of SDS for IDC revenue tracking purposes, a software-defined storage controller software product must reside on one of the first four tiers (0–3) (see the SDS Controller Software section). FOBS solutions can run on SBS hardware in any of the six tiers. Products like Microsoft Storage Spaces and VMware vSAN are typical hyperconverged infrastructure (HCI) solutions that can run on top of a variety of different COTS-based SBS hardware platforms. Products like SimpliVity (now HPE), an HCI offering that requires each node to have an internal and highly customized PCIe-based flash card to meet marketed performance levels, does not meet the SDS requirement even though it is considered to be an SBS solution. It does, however, meet the FOBS solution definition and is tracked under the FOBS tracker and forecast.

Table 1 illustrates the types of SBS hardware, indicating the types (by tier) of those considered as SBS building blocks for SDS and FOBS solutions.

In Table 1, "class" refers to whether the system is considered industry standard (though it may vary by configuration) or is considered a fully custom build. The next six areas — base hardware, compute, storage (composed of server-side flash, solid state disks [SSDs], and hard disk drives [HDDs]), additional hardware components, ROM profile, and core operating environment — indicate where customization may exist within that particular tier of SBS platform. "SDS building block" identifies whether or not this tier of SBS platform contributes to overall SDS systems revenue, while the "procurement model" identifies how SBS platforms of this tier are assembled and procured.

Increasingly, storage systems that IDC tracks under the external storage systems market are built using either a commercial or a configured systems approach. For example, vendors like Nutanix (a hyperconverged systems vendor) or Kaminario (a vendor of software-based all-flash arrays) are using commodity CPU and motherboard components, installing their SDS-CS on these systems before they are shipped.

TABLE 1

Server-Based Storage Classification

Tier	Class	Base Hardware	Compute	Server-Side Flash and Nonvolatile Memory (PCIe/NVMe)	SSDs and HDDs	Additional Hardware Components	ROM Profile	Core Open Embedded (OE)	SDS Building Block?	FOBS Building Block?	Procurement Model
0	Commercial	Standard	Standard	NA	Standard hardware, no custom firmware	None	Standard	No modifications	Yes	Yes	OEM Server vendor
1	Configured	Standard	Standard	NA	Standard hardware, no custom firmware	None	Low-impact configuration or customization to standard ROM/firmware	Functional modifications only	Yes	Yes	OEM Server vendor or ODM Direct
2	Configured	Standard	Standard	Standard hardware, custom firmware	Standard hardware, custom firmware	Contains additional COTS hardware components	Low-impact configuration or customization to standard ROM/firmware	Functional modifications only	Yes	Yes	OEM Server vendor or ODM Direct
3	Configured	Custom	Standard	Optimized hardware, custom firmware	Optimized hardware, custom firmware	May contain additional COTS or specialty hardware components	Custom	Functional modifications only	Yes	Yes	OEM Server vendor or ODM Direct

TABLE 1

Server-Based Storage Classification

Tier	Class	Base Hardware	Compute	Server-Side Flash and Nonvolatile Memory (PCIe/NVMe)	SSDs and HDDs	Additional Hardware Components	ROM Profile	Core Open Embedded (OE)	SDS Building Block?	FOBS Building Block?	Procurement Model
4	Major customization	Custom	Optimized	Optimized hardware, custom firmware	Optimized hardware, custom firmware	May contain additional COTS or specialty hardware components	Custom	Functional modifications only (minor software modifications to take advantage of optimized silicon)	No	Yes	ODM Direct
5	Fully customized	Custom	Custom	Optimized hardware, custom firmware	Custom hardware, custom firmware	May contain additional COTS or specialty hardware components	Custom	Programmed specifically for custom silicon (software is tuned to take advantage of custom silicon)	No	Yes	ODM Direct

Source: IDC, 2018

Software-Defined Storage and File- and Object-Based Storage

SDS and FOBS solutions are storage solutions (hardware plus software) that deliver a full suite of persistent storage services via an autonomous software stack that can run on specific tiers of SBS hardware as outlined in Table 1. SDS/FOBS solutions can be sold either as appliances or as software-only products that are paired with SBS hardware either in the channel or by end-user customers to create a production-ready storage system.

Software-Defined Storage

For a solution to be classified as SDS, it must satisfy the following requirements:

- An SDS solution will be deployed on SBS tiers 0–3 (refer back to Table 1) including (but not limited to) x86-based servers running Windows, Linux, or other off-the-shelf operating system distributions like AIX or Solaris. The solution should not contain any proprietary hardware components like custom ASICs, chipsets, memory components, and/or CPUs, and the software code should not make any assumption of such components being present to offer any essential data services such as inline data reduction, snapshots, and replication.
- SDS solutions can support block-, file-, or object-based data organization. While HCI solutions can be built on any of the three data organization methods, IDC's SDS tracking segregates HCI to provide a granular view of the SDS market that differentiates between disaggregated and hyperconverged storage architectures.
- In general, the solution should be able to run on multiple (physical or virtual) hardware instances that are not factory configured by the supplier. IDC notes that there are some cases where the vendor chooses to sell the SDS product only as an appliance despite the fact that it otherwise meets SDS requirements and could very well run (unsupported) on multiple types of SBS hardware. In general, buyers should be able to procure the platform as software and deploy it in a virtual environment or directly on any physical hardware of their choice (as long as this hardware can be categorized as being in tiers 0–3 in Table 1) without loss of functionality.
- The solution is a standalone system. In other words, it provides all essential northbound storage services and handles all southbound data services and persistence functions without requiring additional hardware or software. IDC therefore considers file systems and logical volume managers to be components of an SDS solution rather than complete systems.

File- and Object-Based Storage

For a solution to be classified as FOBS, it must satisfy the following requirements:

- A FOBS solution can be deployed on SBS tiers 0–5 (refer back to Table 1) including (but not limited to) x86-based servers or proprietary hardware platforms running commodity (Windows, Linux, AIX, Solaris, etc.) or proprietary operating systems. The solution may contain any proprietary hardware components like custom ASICs, chipsets, memory components, and/or CPUs, and the software code may make assumptions of such components being present to offer essential data services such as inline data reduction, snapshots, and replication.
- FOBS solutions can support file- and object-based data organization. HCI products built on file- or object-based data organization schema are included in the relevant file and object segments of the FOBS market.
- The solution should be able to run on multiple (physical or virtual) hardware instances that may or may not be factory configured by the vendor. The software may or may not be

"decoupleable" from the underlying hardware. In certain cases, decoupling the software from the hardware when not supported by the vendor may result in loss of functionality.

Two Key SDS and FOBS Solution Building Blocks

The two key SDS and FOBS solution building blocks are SDS-CS and SBS hardware. Note that these two building blocks are required to create the SDS or FOBS solution, regardless of whether storage is consumed as traditional on-premises infrastructure through access protocols like iSCSI, Fibre Channel (FC), FC over Ethernet (FCoE), NFS, SMB, or CIFS or as a cloud-based service.

SBS Hardware

The SBS tiers are defined in Table 1, and SBS hardware in tiers 0–3 are used in SDS solutions, while FOBS solutions may also include hardware in SBS tiers 4 and 5. Regardless of whether an SDS or FOBS solution is sold as an appliance or a software-only product or offered through both solution delivery models, the solution will typically be running on SBS hardware. In the FOBS market, there are a few examples of older products where this may not be true, but going forward, the trend in both these markets will be to deploy software on SBS hardware that is available from multiple sources.

SDS Controller Software

SDS-CS provides for data persistence, a set of data services (snapshots, replication, etc.), and a method of data organization (block, file, and/or object) along with one or more defined access methods (block, file, and/or object). IDC's SDS-CS trackers and forecasts include all the software value in SDS appliances (but not the SBS value) as well as all SDS software-only sales revenue. The SDS-CS forecast is smaller than the SDS forecast because the latter includes some hardware revenue for appliances. Many SDS solutions are built around a scale-out architecture that leverages a cluster of SBS nodes, all connected via a high-speed switched fabric, although there are single-node SDS solutions. SDS solutions can be deployed on dedicated physical hardware that is not running a hypervisor like Nexenta's NexentaStor or IBM's Spectrum Scale or can be deployed on a virtual machine, in which case it is designed to always require a hypervisor like Dell EMC's ScaleIO or VMware's vSAN.

Table 2 provides detail around SDS-CS offerings, introducing the four flavors of software-only products (see the Solution Delivery Model section) and indicating which of these can be used as SDS and/or FOBS building blocks. Some vendor examples are also provided where they exist.

TABLE 2

Software-Defined Storage Controller Software

SDS-CS Software Categories by Data Organization	SDS Building Block?	FOBS Building Block?	Consumption Model	Examples
Open source software				
Block	x		Software ISV	
File	x	x	Software ISV	
Object	x	x	Software ISV	
Commercial open source software				
Block	x		Software ISV	Red Hat Ceph Storage
File	x	x	Software ISV	
Object	x	x	Software ISV	SwiftStack, Red Hat Ceph Storage
Commercial software				
Block	x		Software ISV	
File	x	x	Software ISV	Quantum StorNext, Dell EMC IsilonFS
Object	x	x	Software ISV	Scality, IBM Cloud Object Storage, Cloudian, Caringo, Hitachi Vantara HCP, NetApp StorageGRID Webscale, DDN WOS, and so forth.
Self-built software				
Block			Storage services only (as a service)	
File		x	Storage services only (as a service)	
Object		x	Storage services only (as a service)	

Source: IDC, 2018

Software-only products come in several flavors: open source software, commercial open source software, commercial software, and self-built software. In detail:

- **Open source software.** Open source software is developed by a community of interested market participants and made available for free download to developers that can then modify and enhance the source code themselves prior to deployment. Support is offered through the community (rather than a commercial vendor) that develops and/or uses the software. Open source initiatives include OpenStack, Ceph, and Sheepdog.
- **Commercial open source software.** Commercial open source software is a commercial distribution based on open source software that is packaged by a vendor as more of a commercial offering with support, upgrades and, quite often, vendor customization. This approach appeals to customers that specifically want to work with open source software stacks but do not want to handle product support and maintenance issues themselves.
- **Commercial software.** Commercial software is proprietary software developed by vendors using their own intellectual property (IP). Commercial software offers support, upgrades, and features that are evolved and maintained by the vendor (rather than the end-user customer).
- **Self-built software.** Self-built software is self-developed by a hyperscaler (e.g., Amazon, Google, Facebook) or very large service provider that delivers subscription-based storage services to end users. These customers generally have a very sophisticated in-house development capability and want to build the software themselves for either functionality or economics reasons.

Primary Segment Descriptors

The sections that follow detail the primary segment descriptors of the SDS and FOBS markets. In essence, these descriptors help define and differentiate the different architectures and attributes of the various SDS and FOBS solutions. For instance, SDS and FOBS solutions fall into either a scale-up or scale-out architectural design. Within each of these architectures, solutions offer data organization and access methods that support block, file, and/or object as well as various combinations. Altogether, there are five primary segment descriptors.

Table 3 shows the primary segment descriptors along with their definitions that IDC uses in this consolidated taxonomy to differentiate solutions offerings.

TABLE 3

SDS and FOBS Market Technology Segments

Segments	SDS Coverage	FOBS Coverage	Additional Notes: SDS Coverage	Additional Notes: FOBS Coverage
Architecture				
Scale up	x	x	While SDS predominantly employs a scale-out architecture, a few products based on the scale-up architecture may fall into the SDS market segment (e.g., OSNEXUS, Kaminario, NexentaStor).	The general FOBS market is moving toward a scale-out architecture. However, many legacy product-based scale-up architectures continue to support FOBS workloads and therefore are tracked in the FOBS forecast.
Scale out	x	x		
Data organization				
Block	x			
Block HCI	x		SDS reports block HCI under a separate HCI market segment in the annual SDS forecast. Examples include ScaleIO and Microsoft Storage Spaces.	
File	x	x		
File HCI	x	x	SDS reports file HCI under a separate HCI market segment in the annual SDS forecast. Examples include Cisco HyperFlex and Nutanix.	FOBS reports file HCI under the FBS market segment in annual FOBS forecast.
Object	x	x		
Object HCI	x	x	SDS reports object HCI under a separate HCI market segment in the annual SDS forecast. Examples include Dell EMC VxRack/VxRail and HPE SimpliVity.	FOBS reports object HCI under the OBS market segment in annual FOBS forecast.

TABLE 3

SDS and FOBS Market Technology Segments

Segments	SDS Coverage	FOBS Coverage	Additional Notes: SDS Coverage	Additional Notes: FOBS Coverage
Solution delivery model				
Appliance only	x	x	Any hardware (appliances or servers) must qualify the SBS tier 0–3 category to be reported under SDS.	Any hardware (appliances or servers) must qualify the SBS tier 0–5 category to be reported under FOBS.
Software only (+SBS hardware)	x	x		
Data services				
Management services	x	x	Relevant data services types vary by workload requirements for the data organization method.	Relevant data services types vary by workload requirements for the data organization method.
Consumption model				
Traditional (on-premises/ off-premises)	x	x	As long as the previously mentioned categories are met, any SDS solution can be consumed in the form of traditional, private cloud, or public cloud.	As long as the previously mentioned categories are met, any FOBS solution can be consumed in the form of traditional, private cloud, or public cloud.
Private cloud (on-premises/ off-premises)	x	x		
Public cloud (off-premises)	x	x		

Note: See Table 2 for the SBS tier 0-3 category.

Source: IDC, 2018

Architecture

A key trend in SDS and FOBS in general is the migration from scale-up to scale-out architectures. While there is clearly still industry revenue being generated by products based on scale-up, dual-controller designs, over time, more and more revenue is being generated from products based on multinode scale-out designs that use SBS hardware. In detail:

- **Scale up.** Scale-up SDS/FOBS architectures are typically based on a server with one or two controllers and one or more internal storage devices (usually HDDs or SSDs — or some mix of the two). In systems with two controllers, cache coherency across controllers is maintained. The design can implement either an active/active or an active/passive controller architecture, each of which has implications for system performance in degraded mode (e.g., a controller failure scenario) as well as for overall performance. Controller performance is relatively fixed, although firmware and software upgrades over the life of the product may produce some incremental performance improvements. Each node runs a standalone operating system. Storage capacity can be scaled by adding more storage devices, which can be either internally or externally attached. In these designs, the performance that can be brought to bear against a single volume, file or object is limited to the performance of a single, standalone node.
- **Federated scale up.** This approach allows multiple scale-up SDS/FOBS platforms to be clustered together under a single namespace. This model enables much larger data stores (regardless of whether they are block, file, or object based) to be centrally managed as a single cluster, allows the overall throughput of the cluster to be scaled as nodes are added, and provides additional resiliency through redundancy (node failover, etc.). It still does not, however, allow performance to be scaled for a single volume, file or object beyond the capabilities of a single node.
- **Scale out.** Scale-out SDS/FOBS architectures are typically based on multinode server configurations, each of which can have one or two controllers and some number of internal storage devices (HDDs or SSDs). Nodes are typically connected over a dedicated, high-speed network. These systems feature a truly distributed software operating environment that enables all the resources spread across multiple nodes to operate as a single storage system. Performance against a single volume, file or object can be scaled as nodes are added to the cluster since each can effectively be spread across multiple nodes using a shared-nothing, data sharding architecture.

  The distributed nature of the software also offers other advantages, such as more predictable performance in environments with widely varying workloads, increased resiliency (because of the inherent redundancy of clustered configurations), easier technology refresh (newer nodes can be transparently added and older nodes can be removed to nondisruptively evolve a system to next-generation technology), and better recovery (since they can more highly leverage parallelism to speed recovery times in the event of component or node failures). A key advantage to scale-out architectures is the ability to independently expand the performance and capacity of the system as a whole just by adding more nodes.

Data Organization and Access

Data Organization

SDS typically uses one of the three data organization schemes for managing persistence (block, file, or object based), while FOBS uses one of the two data organization schemes (file or object based). Products based on block- and file-based data organization are available in both scale-up and scale-out implementations (depending on vendor), but object-based products all use a scale-out design. Definitions of the three data organization schemes and hyperconverged infrastructure are as follows:

- **Block-based storage.** This is a category of data storage used mostly in storage area network (SAN) environments where data is saved in volumes known as blocks. Each block in block storage is configured by a storage administrator and acts like an individual hard drive. The blocks are controlled with the help of server-based operating systems and can be accessed by protocols such as iSCSI, FC, or FCoE. Block storage includes a lightweight abstraction layer, and therefore, it can provide better performance for many workloads than file- and object-

based storage (both of which employ higher levels of abstraction for easier management and higher-end scalability).

- **File-based storage (FBS).** This is a category of storage used mostly in network-attached storage (NAS) environments where unstructured data (files) is stored hierarchically in directories and subdirectories (or folders). File systems maintain limited metadata such as file size, file location, file name, date of creation, modification time stamps, access control, and directory hierarchies. File-based storage typically provides a larger range of scalability and makes storage easier to manage (since it is managed at the file level rather than the volume level). A global namespace federates multiple file systems across storage devices, allowing simplified management and access to file data regardless of physical location. File-based storage can be managed by the storage system, whereas block-based storage requires the assistance of server-based operating systems. File access protocols (across scale-up and scale-out FBS architectures) include NFS, CIFS, and SMB.
- **Object-based storage (OBS).** This is a category of storage that uses a flat tenant/account/container/object namespace that is designed for large, multipetabyte data sets. Objects are referenced directly (rather than hierarchically) through a metadata repository that stores and manages the attributes not only for individual objects but also for groups of objects or buckets (a semantic that was created to simplify policy-based management). Policies can be applied at the object or bucket level. In contrast to the limited metadata capabilities in file-based storage, object-based storage supports custom metadata such as application or user-specific information. Custom metadata assists with data management policies, better insights on existing data sets, management, and so forth. There are also object-based storage solutions that leverage NoSQL databases as metadata repositories and persistent data stores (instead of storing data in a file system). Object-access protocols include HTTP/REST, CDMI, S3, OpenStack Swift, and other object-specific interfaces.
- **Hyperconverged infrastructure.** HCI systems collapse core storage and compute functionality into a single, highly virtualized solution. A key characteristic of HCI systems that differentiates them from other storage systems is that they colocate compute and storage resources in the same server. This colocation shortens the data path between "server" and "storage" resources since they both reside within the same server, providing lower access latencies than disaggregated, networked storage can for many environments. These systems are usually deployed as scale-out clusters with three or more nodes per cluster to provide for high availability and recovery upon node or device failure. Each node within a cluster contributes all of its resources to an abstracted pool of compute and storage capacity resources. Compute resources are used to service application workload requirements, while storage capacity resources are used to service storage workload requirements.

  In addition to integrating storage and compute functions into a single node or cluster of nodes, all HCI systems employ:

  - A hypervisor that provides workload adjacency, management, and containerization in addition to providing the hardware abstraction layer
  - A distributed block or file system or an object store that serves as the data organization, management, and access platform
  - An (optional) Ethernet switch that provides scale-out and/or high-availability capabilities

HCI systems can be purchased as a complete appliance or rack-scale system with core features included or as a software-only solution.

Access

Access defines how applications access data stores, and there are again three schemes: block, file, or object based. Block access interfaces include iSCSI, FC, and FCoE. File access interfaces include NFS, CIFS, and SMB. Object-based access interfaces include HTTP/REST, CDMI, S3, and OpenStack Swift. How data is accessed can be independent from how it is organized, and there are getting to be more layered offerings in the market that offer multiple access types against a data store that uses a single data organization type (e.g., the Ceph open source object-based data store simultaneously supports block-, file-, and object-based access). As big data/analytics deployments proliferate, more customers want to store a data set once but access it multiple times with different applications. Some of these applications may, for example, use file-based access, while others may use object-based access. The ability for a single data store to support multiple simultaneous access methods can be a significant simplification that saves time and money for customers.

Traditionally, there have been performance differences between block-, file-, and object-based data stores. Block has historically been very popular for more latency-sensitive workloads, while file- and object-based storage have typically supported better scalability for larger data set sizes. While there still are performance differences between the three today, the widespread use of flash storage devices in enterprise environments has enabled file- and object-based storage to be used in more latency-sensitive environments without undue performance impacts, a fact that enables some additional consolidation of workloads for improved efficiencies.

Solution Delivery Model

SDS and FOBS solutions are delivered as either appliances, software-only products, or cloud-based storage services. An appliance is an integrated storage system that includes both SBS hardware and SDS-CS software. The SBS hardware includes controllers, one or more internal storage devices, and internal cabling and may also include external network connection hardware. The software includes SDS-CS, which IDC tracks and forecasts as a standalone market (see *Worldwide Software-Defined Storage Controller Software Forecast, 2016–2020,* IDC #US41998415, December 2016). The SDS-CS forecast is an input to the SDS forecast (because SDS also includes some appliance-based hardware revenue). Software-only products are designed to be deployed on SBS hardware but put the responsibility on the buyer to select the hardware and install the software. Vendors of software-only products typically certify their products on specific hardware and make a hardware compatibility list available. Appliances are easier to buy and deploy and offer a simpler, one-stop shop for support services, while software-only products provide additional flexibility in the choice of hardware platform (a factor that can lead to lower cost for some customers) but can impose more complexity to purchase, deploy, and support the storage solution.

Data Services

Regardless of whether it is block, file, or object based, storage must support a standalone suite of data services via the software stack that is not dependent upon any particular hardware resource for anything other than data persistence. Workload requirements determine data services requirements, and block-, file-, and object-based storage can support different types of data services. Data services include storage management functions like inline data reduction and thin provisioning, data protection approaches like RAID or erasure coding, snapshots, encryption, quality of service, replication, versioning, audit trails, access control, secure data shredding, orchestration and automation services, and other functions necessary to meet customer requirements.

Products that just provide orchestration and automation services are not considered SDS solutions (they are not a standalone storage solution). In the current taxonomy, networking features such as "software-defined networking" are considered beyond the purview of the SDS definition.

Consumption Models

There are two basic consumption models that describe how an application accesses storage and are in common use in the SDS and FOBS markets — traditional (noncloud) and "as a service" (private/public cloud):

- **Traditional (noncloud).** This model refers to any storage system/solution deployed in a noncloud space. It includes four of the five enterprise storage consumption models that have been discussed elsewhere as applying to enterprise storage solutions in general (appliance, software only, converged infrastructure, and HCI). Under the traditional model, an SDS solution is purchased outright (i.e., software is licensed on a perpetual basis), deployed on-premises, and accessed using standard storage protocols such as iSCSI, FC, FCoE, NFS, SMB, CIFS, or object protocols (regardless of whether it is purchased as an appliance or a software-only product).
- **As a service.** In the "as a service" model, customers can obtain storage on a more granular basis from a cloud. Instead of being accessed directly through protocols like iSCSI or NFS, storage services are accessed through cloud APIs, which can enable more of a "self-service" mindset for consumers of cloud storage services:
  - **Private cloud.** Private cloud services can be deployed onsite or offsite, managed by in-house staff or a third party, and typically provide storage services for a single enterprise (or extended enterprise). In a private cloud managed in-house, the IT organization effectively plays the role of a cloud service provider in making storage available for use. This deployment type has restrictions, defined and controlled by the enterprise, on access and level of resource dedication.
  - **Public cloud.** Public cloud services are shared among unrelated enterprises and consumers, open to a largely unrestricted universe of potential users, and designed for a market rather than a single enterprise. The public cloud market includes a variety of services designed to extend or, in some cases, replace IT infrastructure deployed in corporate datacenters. Examples include Amazon Web Services S3, Rackspace Cloud Files, and DreamHost DreamObjects. It also includes content services delivered by a group of suppliers IDC calls "value-added content providers" ("VACPs") such as Thomson Reuters or Dow Jones. VACPs gather, organize, and provide access to large repositories of digital content (such as news, stock quotations, financial and technical analysis, scholarly works, advice columns, comparative information, resumes, job listings, financial data, videos, and music).

Note that there is a hybrid model where a vendor agrees to deploy a storage appliance in an end-user customer datacenter but charges the customer for the use of "storage services" based on a storage utility pricing model that provides some of the same service elasticity benefits as the cloud. While many vendors offer this hybrid model, few customers use it as a way to acquire and use storage services over a long period of time. Many vendors use this approach as more of a "low risk, try and buy" program that is then moved to the more traditional model within six to nine months. IDC does not track this segment separately, but these products are viewed as using a traditional consumption model.

Market Segments

IDC measures the size of the SDS and FOBS markets in terms of worldwide infrastructure revenue on the two key building blocks (SDS-CS and SBS hardware) and storage capacity shipped. The FOBS market also includes revenue from sales of gateways and general-purpose file servers.

The SDS market is categorized mainly into four segments based on data organization (block, file, object, and HCI):

- **Block.** Block-based SDS vendors include Dell EMC (Unity, a unified storage platform), NetApp (with its ONTAP Select product that actually supports unified storage), and Veritas (HyperScale).
- **File.** File-based storage vendors include (but are not limited to) DDN (EXAScaler, GRIDScaler), Dell EMC (Isilon), IBM (Spectrum Scale), NetApp (ONTAP Select), Nexenta (NexentaStor), Red Hat Gluster Storage, Veritas Access, and Qumulo File Fabric. This segment does not include HCI built on top of FBS and some appliances that employ SBS hardware from tiers 4 and 5 (e.g., Panasas).
- **Object.** Object-based storage vendors include (but are not limited to) DDN (Web Object Scalar), Dell EMC (Elastic Cloud Storage), Hitachi Vantara (Hitachi Content Platform), IBM (Cloud Object Storage), NetApp (StorageGRID Webscale), SwiftStack, Caringo, Cloudian, Nexenta (NexentaEdge), Scality (RING), and Western Digital (Active Archive System/Active Scale). This segment does not include self-built object-based storage (OBS) software deployed for cloud-based storage services, HCI built on top of OBS, and some appliances that employ SBS hardware from tiers 4 and 5 (e.g., Dell EMC Atmos and Centera).
- **HCI.** HCI vendors include Cisco (HyperFlex), Dell EMC (ScaleIO and derivative products [VxRack and VxRail]), HPE (SimpliVity), Microsoft (Storage Spaces), Nutanix (NX-Series), and VMware (vSAN).

The FOBS market is characterized mainly into two segments based on architecture: scale-up or scale-out and file- or object-based storage. Examples of products tracked in these segments are as follows:

- **File servers.** Examples include Microsoft Windows Server 2012, Red Hat Enterprise Linux, and SUSE Linux Enterprise as well as solutions based on proprietary Unix distributions such as Oracle Solaris, HPE HP-UX, and IBM AIX.
- **Scale-up file-based storage solutions.** Examples of scale-up appliances and gateways include Windows Storage Server file servers from Dell and HPE, ZFS-based appliances from Oracle and other suppliers, and NetApp All Flash FAS (AFF) running ONTAP in cluster mode.
- **Scale-out FBS solutions.** Examples of scale-out FBS appliances include (but are not limited to) IBM SONAS (based on IBM Spectrum Scale), Quantum StorNext 5 Appliances, Dell EMC Isilon, NetApp FAS and other cluster ONTAP-based products, Dell FS Series, Qumulo File Fabric, Veritas Access, and Panasas ActiveStor. This subsegment also includes a small percentage of disk-based scale-out file virtualization solutions such as those from Avere Systems (acquired by Microsoft). Examples of scale-out FBS software include Red Hat Gluster Storage, Dell EMC IsilonSD Edge, Lustre-based variants, IBM Spectrum Scale (GPFS), Quantum StorNext, and open source FBS software. HCI solutions such as those from Nutanix, Pivot3, and Scale Computing are also included in this category.
- **Object-based storage solutions.** Examples of OBS appliances include (but are not limited to) the Dell EMC ECS (Elastic Cloud Storage), Hitachi Vantara HCP (Hitachi Content Platform), NetApp StorageGRID Webscale, Western Digital (HGST) Active Archive System and ActiveScale, and Cloudian HyperStore. Examples of OBS software-only products include

Caringo Swarm, Cloudian HyperStore, DDN WOS, Exablox, Red Hat Ceph Storage, Scality RING, and SwiftStack. Some products like IBM's Cloud Object Storage and DDN Web Object Scaler are available as both appliances and software-only products. Examples of self-built OBS platform vendors include Amazon S3, Google, Microsoft Azure, and Rackspace. HCI solutions built on object data organization are included in this category.

Aligning SDS and FOBS to Other Related IDC Taxonomies

The SDS and FOBS taxonomies provide a platform view that includes revenue from the sales of appliances, software only, and SBS hardware to various buyers. The SBS and FOBS markets include some subsegments tracked and reported in IDC's Quarterly Enterprise Storage Systems (ESS) Tracker and IDC's Storage Software QView.

IDC's Quarterly ESS Tracker tracks revenue generated from sales of SBS platforms (tiers 0–5) sold in the following categories:

- Internal storage (OEM server vendors such as Cisco, Dell, HPE, and Huawei)
- External RAID that includes turnkey bundled hardware and software appliances (vendors here include traditional enterprise storage suppliers such as Dell EMC, HPE, and Hitachi Vantara)
- ODM Direct (vendors such as Foxconn, Inventec, Quanta, and SYNNEX)

ESS revenue and capacity from sales of products across the aforementioned categories that classify as SDS and/or FOBS (those created from SDS-CS and SBS hardware) are included in the SDS and/or FOBS taxonomy and forecast. Essentially, SBS revenue tracked in SDS and FOBS is a subset of that tracked in the ESS tracker and forecast. Note that both the SDS and FOBS markets track the full value of the systems (hardware plus software) and therefore account for a markup in categories whereas ESS tracks only the value of associated storage (compute-intensive servers in internal storage and ODM Direct).

IDC's Quarterly Storage Systems QView tracks revenue generated from sales of storage software in six functional markets:

- Archiving
- Data protection and recovery
- Storage/device management
- Storage infrastructure
- Storage replication
- Software-defined storage controller software

Both SDS and FOBS draw revenue from sales of storage software classified under SDS-CS in IDC's Storage Software QView. The other five categories are not tracked in the SDS and FOBS markets.

Figures 2 and 3 show how SDS fits into and contributes to these overall markets. The SDS-CS forecast includes just the SDS-CS revenue in the yellow box under the storage software market. The yellow boxes indicate tracked SDS revenue (which includes all SDS-CS revenue plus some appliance-based hardware revenue), while the green-tinted boxes indicate where those revenue contributions come from in enterprise storage systems and software. The orange boxes show the additional hardware revenue that are in the FOBS forecast that account for most of the significant difference in overall market revenue between the two. Note that SDS appliance revenue (made up of both hardware and software revenue) contributes to the enterprise storage systems market, while the revenue

generated by software-only purchases is split between the storage software market (for the SDS-CS revenue) and the enterprise storage systems market (for the necessary SBS hardware revenue).

FIGURE 2

How SDS Fits into IDC's Enterprise Storage Taxonomy

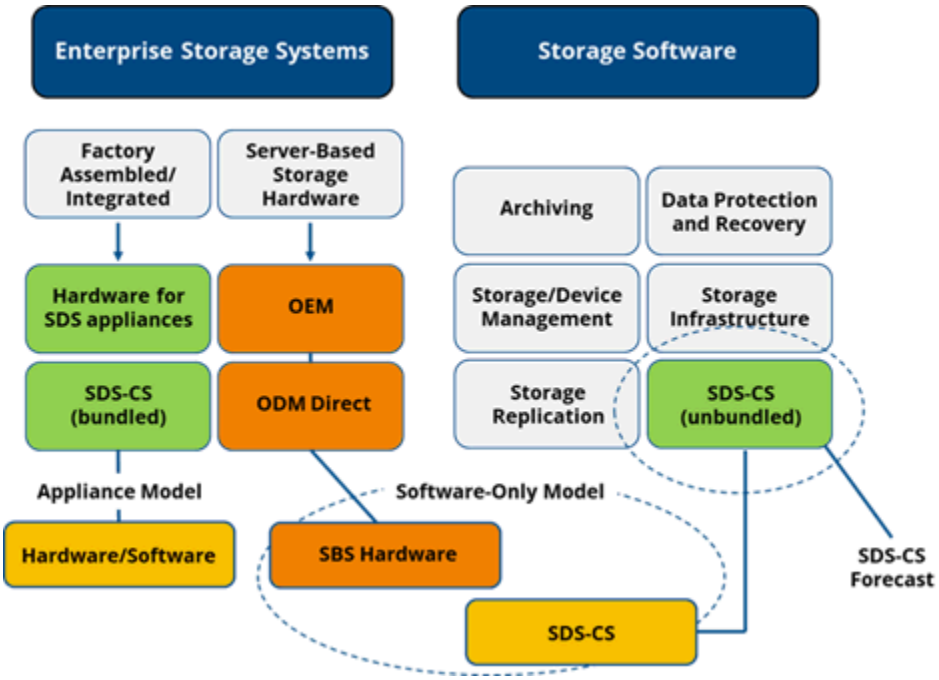

Source: IDC, 2018

FIGURE 3

How FOBS Fits into IDC's Enterprise Storage Taxonomy

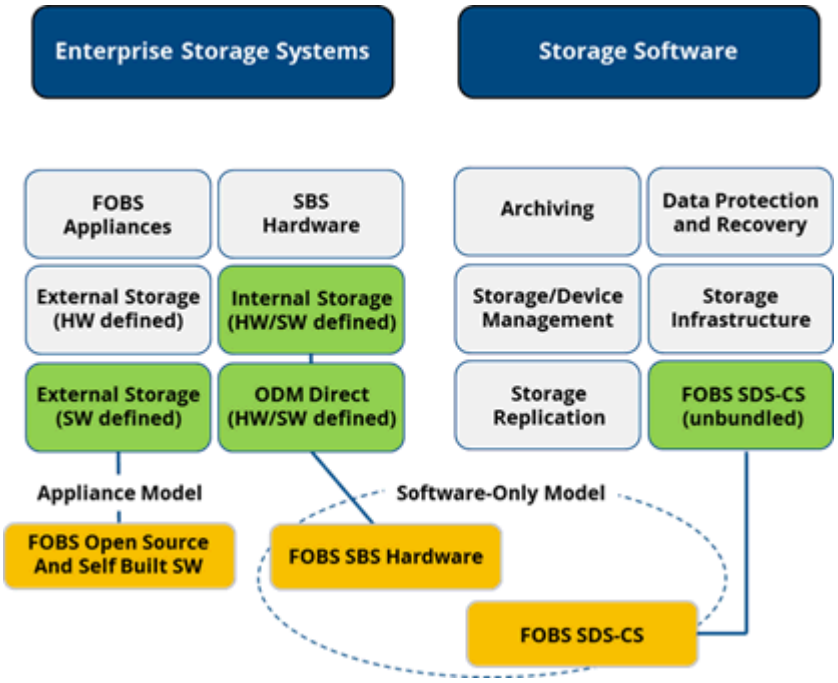

Source: IDC, 2018

LEARN MORE

Related Research

- *Business and Content Applications Will Dominate Evolving Production SDS Deployments* (IDC #US43249017, December 2017)
- *Veritas Vision 2017 – Renewed Focus on Multicloud Data Management* (IDC #IcUS43114717, September 2017)
- *Worldwide Software-Defined Storage Forecast, 2017-2021: SDS Market Growth Significantly Outpaces Enterprise Storage Growth, Led by HCI* (IDC #US43062517, September 2017)
- *Worldwide File- and Object-Based Storage Forecast, 2017-2021* (IDC #US42280717, September 2017)
- *Hedvig Brings Hyperscale Agility and Economics to Enterprises* (IDC #US42983117, August 2017)
- *Successful SDS Deployments Demand Requisite Expertise on the Part of Customers* (IDC #US42919517, August 2017)
- *IDC's Worldwide Storage Software Taxonomy, 2017* (IDC #US42834017, August 2017)
- *Veritas HyperScale for OpenStack: Consistent Performance and Data Protection for OpenStack-Based Cloud* (IDC #US42580517, June 2017)

Synopsis

This IDC study forms the consolidated taxonomy for several markets, including software-defined storage controller software, software-defined storage, file- and object-based storage, and server-based storage. This study specifically covers market background, taxonomy, and definitions for these markets.

"Software-defined storage products are driving an ever-increasing percentage of overall enterprise storage revenue," said Eric Burgener, research director, Storage at IDC. "Enterprises like the agility, flexibility, and economics of the software-defined model, and IDC expects that this segment will continue to drive double-digit revenue growth over the next five years."

"The growing unstructured data sets continue to push the boundaries of the file- and object-based storage solutions market, demanding better scalability, reliability, and performance at low cost," said Amita Potnis, research manager, File- and Object-Based Storage Solutions. "As enterprises grow beyond tens of petabytes, object storage will be a solution of choice, but the need for file-based storage will continue. IDC expects that the file-based storage market will transform itself to add new features and functionality to remain viable, while flash will play a bigger part in object-based storage solutions."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com