

Software defined storage



By Alex McDonald, SNIA Europe Board of Directors,
DMTF OSDDC member, NetApp.

YOU'D HAVE TO BE LIVING in a cave not to have heard the latest IT buzzword; software defined. I first heard this phrase at a seminar about two years ago in the context of work Stanford University was doing on SDN (storage defined networking). Even back then, there were the rumblings of extending it out to include and define SDS (software defined storage) and an encompassing SDDC (the software defined datacenter).

So while SDN and now SDS and SDDC is an evolutionary product of that earlier thinking and the direct result of considering how we can best manage and control virtualization and Cloud technologies, the SDDC term was only coined recently.

Before we can discuss SDS, it's important that we get some kind of idea what we mean by SDDC.

SDDC has been defined in a number of ways, with various definitions competing for attention. But while there are already various management standards that exist for physical, virtual and cloud-based systems, the industry is currently lacking

standard architectures and definitions to describe the abstract, pooled resources of the Software Defined Data Center.

That's not unexpected in the early stages of the development of an idea into working technology. To bring clarity, members of the DMTF (a standards body with a varied and wide membership of the great and the good in the IT industry) set up the

OSDDC (Open Software Defined Data Center Incubator at <http://www.dmtf.org/news/pr/2013/5/dmtf-address-need-open-software-defined-data-center-standards>) to identify standards required, possible architectures and definitions to describe the emerging technologies that are expected to revolutionize IT infrastructure in the coming years.

So how does the DMTF OSDDC define SDDC?

"The SDDC domain is proof of the continuing evolution of virtualization through cloud computing into SDDC technologies.

Cloud computing was the new operational model for IT Services built upon foundational and fundamental virtualization

technologies. While cloud computing uses virtualization technologies and/or converged

Infrastructure as a Service (IaaS) approaches, it is still focused on

the delivery and consumption of IT Services. SDDC, as

the next phase in the evolution of this entire technology domain, promises to deliver more intelligent services, better management solutions and value on top of these commodity and standardized hardware platforms.

"An SDDC is a data center or cloud computing infrastructure in which all elements of the infrastructure



“Software-defined storage (SDS) is virtualized storage with a service management interface. SDS includes pools of storage with data service characteristics that may be applied to meet the requirements specified through the service management interface.”

including networking, storage, compute, and security, are virtualized and delivered as a service to the consumer. An SDDC infrastructure is abstracted from the entire underlying physical infrastructure (and even the virtual infrastructure in some cases). This abstraction enables programmatic and automated provisioning, deployment, configuration, and management of the SDDC.”

While other definitions have been proposed by the various vendors and standards organizations, they all have similar, if not identical, intent or wording. For example, SNIA succinctly puts it this way in their dictionary;

“A Software Defined Data Center (SDDC) is a virtualized data center with a service management interface. Application requirements determine the service levels provided.”

Roots in cloud and virtualization

Cloud Computing has been the new operational model for IT Services built upon virtualization technologies. SDDC, as the next phase in the evolution of this entire technology domain, promises to deliver greater more intelligent services, better management solutions and value on top of commodity and standardized hardware platforms.

SDDC incorporates and is heavily dependent upon the use of topologies that abstract, pool and automates the use of virtualized resources. Virtualization technologies can be thought of as a commodity, or the common resources that when integrated are used by SDDC. Industry standardized management models and APIs provides this level of abstraction.

With these topologies we can describe compute, network, and storage resources, and then, with a suitable set of APIs (application programming interfaces) we can manage, organize and manipulate them. With these abstractions, management of resources becomes standardized and simplified. That’s the attraction of SDDC.

A virtualized X isn’t software defined X

It’s important to note that a “software defined X” is not the same as a “virtualized X”. Storage virtualization enables the pooling of physical storage facilities and devices from various physical networked devices into what appears to be a single (or one to many as well) storage resource pool or pools.

Software defined storage (SDS) on the other hand, is an ecosystem of products that decouples the application from underlying storage network hardware. SDS software makes visible all physical and virtual resources and enables programmability and automated provisioning based on consumption or need.

In essence, software defined separates the control plane (how we manage a resource and what we manage it with) from the data plane (how and what we use to access the resource). SDS does that for storage; the data plane is separate from the control plane. This matches well the definition used in SDN, where exactly (and unsurprisingly) the same facilities are provided.

In many respects, SDS is more about packaging and how IT users think about and design data centers. Storage has been largely software-defined for more than a decade: the vast majority of storage features have been designed and delivered as software components within a specific storage-optimized environment.

Again turning to SNIA’s definition; *“Software-defined storage (SDS) is virtualized storage with a service management interface. SDS includes pools of storage with data service characteristics that may be applied to meet the requirements specified through the service management interface.”*

What does software defined storage look like?

There’s no rigid set of features that differentiates SDS from storage you can buy and deploy today. Storage that meets SDS

requirements needs to include:

- **Automation** – Simplified management that reduces the cost of maintaining the storage infrastructure.
- **Standard Interfaces** – APIs for the management, provisioning and maintenance of storage devices and services.
- **Scalability** – Seamless ability to scale the storage infrastructure without disruption to availability or performance.
- **Virtualized Data Path** – Block, File and Object interfaces (or some combination of them) that support applications written to these interfaces.

Some analysts claim that a needed characteristic is that only block storage is suitable for the SDS treatment. That’s not true; any storage, as long as it clearly separates out the control plane from the data plane and has some form of programmatic interface for provisioning and managing virtual instances of itself is required.

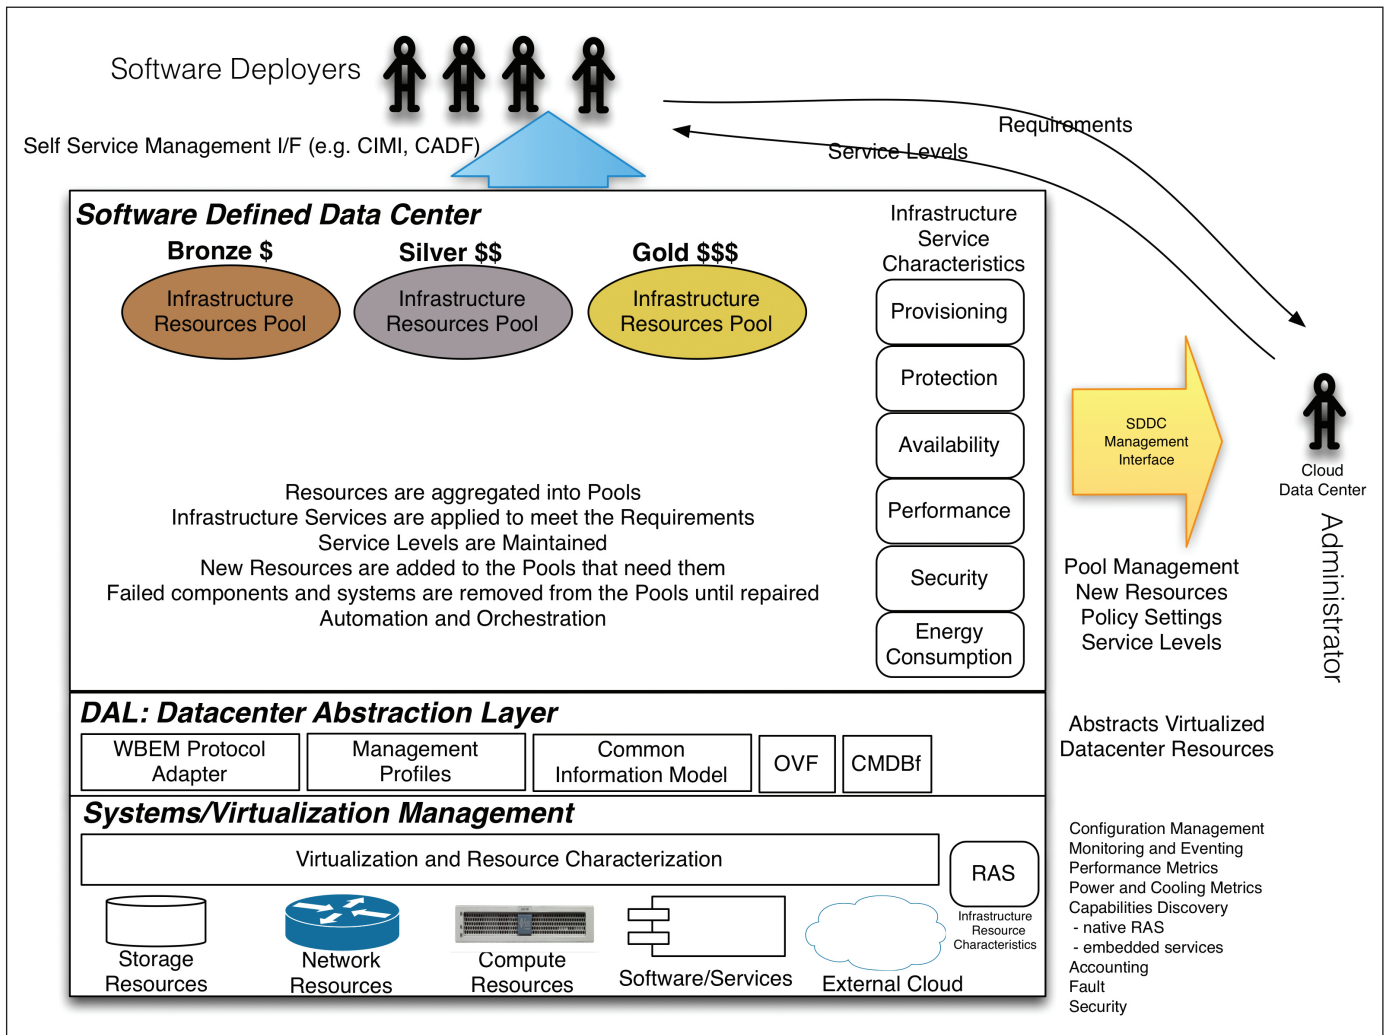
Ideally, SDS offerings allow applications and data producers to manage the treatment of their data by the storage infrastructure without the need for intervention from storage administrators, without explicit provisioning operations, and with automatic service level management.

In essence, SDS is like SDN is like any other resource that we may wish to manage. Taken together, these represent the elements of an SDDC.

If SDx isn’t a virtualized X, then SDDC isn’t virtualization

SDDC differs from Cloud and virtualization in these ways:

- SDDC builds upon the successes of server virtualization, by broadening the individual components of the data center that have been virtualized and envisioning a unified control console & management solution for all the data center’s parts.
- Cloud is a relatively new IT operational model (and marketing model!) focusing on the delivery and consumption of IT



Services. Often, the underlying complexities of the physical and virtualized environments are abstracted from the consumer (as in PaaS and SaaS today).

- SDDC extends this operational model by further refining and expanding upon the three traditional delivery models of Cloud; that is, infrastructure, platform and software as a service (IaaS, PaaS and SaaS respectively).

SDDC isn't a panacea either. It doesn't simplify the complexity or management of the physical data center environment. The physical data center will still be the major underlying component for any virtualized, Cloud or SDDC solution, regardless of vendor.

The provider, carrier or intermediary will still have all of the complexity of managing and operating a physical data center -- as they do today. An SDDC, however, may make the management more efficient.

Realizing an SDDC

To realize an SDDC, data center resources, such as compute, network, and storage, need to be expressed as software. They also need to have certain characteristics, such as multi-tenancy; rapid resource provisioning; elastic scaling; policy-driven resource management; shared infrastructure; instrumentation; and self-service, accounting, and auditing.

This ultimately entails a programmable infrastructure that enables valuable resources to be automatically cataloged, commissioned, decommissioned, repurposed, and repositioned.

Lastly, there needs to be the desire to do this, and to drive vendors, standards bodies and developers of the myriad services that make up a data center to do this.

A recent Forrester Market Overview on Midrange Storage (<http://www.forrester.com/Market+Overview+Midrange+Storage/>)

fulltext/-/E-res110681) noted:

"When asked whether a software defined datacenter was part of their strategic planning, 58% of respondents disagreed or completed disagreed, with 24% somewhat agreeing. Only 17% completely agreed or agreed."

A large number, albeit still a minority, think that SDDC is strategic in the year ahead. For a concept that's only a few years old, it's a significant number. It will be interesting to see, as the DMTF finalizes its report on SDDC, how this changes in the next few years.

For more information about SNIA and the current work and its technical whitepaper on Software Defined Storage that is currently available for public review, please visit: http://www.snia.org/tech_activities/publicreview

Note: This article was based on the work of the DMTF OSDDC Incubator group.