JOURNEY TOWARDS SOFTWARE DEFINED DATA CENTER (SDDC)

Role of SAN Volume Controller (SVC) and Spectrum Virtualize Software

APRIL 2015

THE NET NET

While it has always been the case that IT must respond to increasing business demands, competitive requirements are forcing IT to do so with less. Less investment in new infrastructure and less staff to manage the increasing complexity of many enterprise solutions. And as the pace of business accelerates those demands include the ability to change services… quickly. Unfortunately, older technologies can require months, not minutes to implement non-trivial changes. Given these polarizing forces, the motivation for the Software Defined Data Center (SDDC) where services can be instantiated as needed, changed as workloads require, and retired when the need is gone, is easy to understand.

The vision of the SDDC promises the benefits needed to succeed: flexibility, efficiency, responsiveness, reliability and simplicity of operation… and does so, seemingly paradoxically, with substantial cost savings. The initial steps to the SDDC clearly come from server virtualization which provides many of the desired benefits. The fact that it is already deployed broadly and hosts between half and two-thirds of all server instances simply means that existing data centers have a strong base to build on. Of the three major pillars within the data center, the *compute* pillar is commonly understood to be furthest along through the benefits of server virtualization.

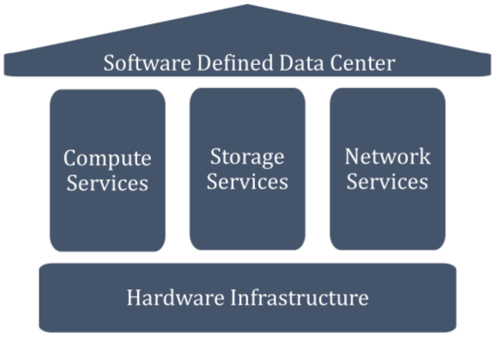

The key to gaining the lion's share of the remaining benefits lies in addressing the *storage* pillar. This is required not only to reap the same advantages through storage virtualization that have become expected in the server world, but also to allow for greater adoption of server virtualization itself. The applications that so far have resisted migration to the hypervisor world have mostly done so because of storage issues. The next major step on the journey to the SDDC has to be to virtualize the entire storage tier and to move the data from isolated, hardware-bound silos where it currently resides into a flexible, modern, software-defined environment.

While the destination is relatively clear, how to move is key as a business cannot exist without its data. There can be no downtime or data loss. Furthermore, just as one doesn't virtualize every server at once (unless one has the luxury of a green-field deployment and no existing infrastructure and workloads to worry about) one must be cognizant of the need for prioritized migration from the old into the new. And finally, the cost required to move into the virtualized storage world is a major, if not the primary, consideration. Despite the business benefits to be derived, if one cannot leverage one's existing infrastructure investments, it would be hard to justify a move to virtualized storage. Just to be sure, we believe virtualized storage is a prerequisite for Software Defined Storage, or SDS.

In this Technology Brief we will first look at the promise of the SDDC, then focus on SDS and the path to get there. We then look at IBM SAN Volume Controller (SVC), the granddaddy of storage virtualization. SVC initially came to market as a heterogeneous virtualization solution then was extended to

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557
www.tanejagroup.com

homogeneous storage virtualization, as in the case of IBM Storwize family. It is now destined to play a much more holistic role for IBM as an important piece of the overall Spectrum Storage program.

THE VISION OF THE SOFTWARE DEFINED DATA CENTER

The Software Defined Data Center (SDDC) seeks to free the enterprise from the limitations and cost associated with hardware-bound solutions. A straightforward way to understand the benefits that come from the insertion of a clutch-plate between the hardware and the services that use the hardware is to consider the pervasiveness of server virtualization.

The lessons from server virtualization were well learned. Many physical servers could be consolidated into a few. Virtual machines could be provisioned with a press of a single key. Additional resources could be added or subtracted non-disruptively. Workloads could be balanced and moved to the right server automatically. Virtual machines could be grouped into consistency groups and protected as a whole. The list goes on and on. And yet all this flexibility came at lower cost to the enterprise.

One of the main reasons for this lower cost was the fact that all existing servers, regardless of their make or model, could be brought into play and the enterprise did not have to throw away any resources. One could argue that server virtualization would have failed if it required all new hardware. And one would be correct in that assessment.

The logic of SDDC is essentially to take these lessons and apply them to the other major layers that make up the overall IT infrastructure: storage and networking. Once hardware is abstracted away from all three layers then the infrastructure could have a major payoff for the enterprise. IT would no longer just be a cost center but rather a full partner of the business it serves. Business would ask for IT resources and get them almost instantaneously. The utilization of resources would go well beyond the 15-20% today to 80% or more. The business entity would itself be more agile and able to respond to competition more nimbly. And new ideas would be converted into new products with alacrity. Such is the vision for the Software Defined Data Center. And given what we have learned from server virtualization such flexibility is not a pipe dream.

Both the storage layer and the networking layers have been undergoing this transformation independently and on their own time scale over the past decade. We will leave the progress related to virtualization of the networking layer to others. Here we will examine the transformation taking place in the storage layer.

At the industry level storage virtualization has had a checkered past. In one sense, storage endorsed virtualization well before any of the other layers did. If one looks at the definition of RAID, it is simply a ganging up of a number of HDDs to look and behave like one from an operating system's perspective. This is the most elemental form of virtualization. The HDDs were virtualized to create a pool of storage and the boundaries between the HDDs were eliminated. In spite of this first step towards virtualization, storage eschewed further advancements until about 2001. This is when a number of new of startups developed storage virtualization software that could sit atop heterogeneous storage arrays and essentially virtualize across them. For a variety of political reasons, however, these solutions did not find a market.

The only exception was the introduction of IBM SAN Volume Controller (SVC) in 2003. If IBM had a grandiose vision for this product at that time one could not tell, for the first implementation of SVC was simple: to sit in-band, in front of a variety of storage arrays, mostly from IBM initially, and abstract away all the idiosyncrasies of the different boxes and 1) deliver a common set of data services and 2) present a common interface to manage those heterogeneous storage resources. Little did IBM know at that time that SVC technology would become the basis of a new line of storage arrays, namely Storwize family, and become the lynchpin for what later became known as Software Defined Storage or SDS (described below). The way to conceptualize this is to recognize that storage virtualization is

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557
www.tanejagroup.com

at the heart of Software Defined Storage, which is at the heart of SDDC. It would not be overstating it to say that since IBM consistently and relentlessly developed SVC during the time that the rest of the industry was unsure about whether or not storage virtualization was detrimental to their existing product lines that today it stands to deliver a complete SDDC quicker than its competitors can.

Role of Software Defined Storage in SDDC

The term Software Defined Storage (SDS) did not come into our vocabulary until only a few years ago. So what is it and why is it important to SDDC? As with any new buzzword SDS is often defined by vendors to suit their own purpose and make whatever solution they have meet the requirements of SDS. But the fundamental concept is not that hard to understand. Let's start with how a typical storage array is built today. Whether it is dual-controller architecture or scale-out the hardware and software is paired in such a way that the two are inseparable. In other words, the software is tied intimately to the hardware. As a result, all the data services built on top of this architecture are also intimately tied to it. This creates a classic proprietary storage product and the industry is laden with such products. In technical parlance, the control plane and the data plane are inseparable. The control plane is responsible for "what action needs to be taken and where is the data located" whereas the data plane is responsible for initiating and controlling the data transfer itself. This is the way storage arrays have been designed for decades.

Enter SDS. The concept of SDS starts with separating the control plane from the data plane and implementing the control plane in software. The immediate effect of this is that the control plane can now interact with a variety of hardware that would act as the data plane. No longer is it beholden to one type of hardware. One common interpretation of SDS is that it works only with commodity hardware. This is a fallacy. It is totally conceivable that the control plane running as software in a server interacts with proprietary hardware. Since all the data services, such as snapshots, cloning, thin provisioning, replication, auto tiering, caching, data deduplication, compression and a myriad of others would be managed by the control plane, these services would therefore work with a variety of hardware underneath.

With this understanding one can begin to see the power of SDS and virtualization together. Think of virtualization as just another service that is managed by the control plane. Not only does that mean that it can be selectively applied to a wide variety of hardware but it also means storage resources can be allocated, moved, transferred and removed under software control. This generates a level of flexibility that has been unseen in the storage industry until now. The sidebar shows the benefits that accrue to any service that is so implemented. In a virtualized environment not only can the server resources now be provisioned, moved and load balanced but the same applies to storage resources. The implications of such a combination are limitless.

Benefits of a Service delivered via SDS

Flexibility – Services are made available in small, discrete pieces in the software layer, and altered when required. One instance of a service can be substituted for another without regard to the physical host on which the data is kept.

Efficiency – Services are created with the capabilities required, and eliminated when no longer needed. There is no need to over-provision. Savings accrue simply through the elimination of waste, and are amplified by higher utilization rates.

Responsiveness – Services are changed dynamically to meet new demands. Manual intervention is minimized through automated, policy driven management. Initial provisioning is greatly simplified and can be done in a self-service portal. Capacity is added incrementally.

Reliability – Replication and automated placement policies protect from failures in a portion of the underlying hardware. Performance is optimized through load balancing ensuring more even utilization of the physical resources.

Simplicity of Operation – A single management interface for the control plane works across disparate physical resources. Further gains come from enhanced automated and policy driven management.

Cost Savings – One gets more from less. Less investment in physical capacity and the staff to maintain it.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557
www.tanejagroup.com

This is the power that can be unleashed by a product like SVC, implemented as a software-defined storage offering. The fact that IBM has nurtured this product for over a decade and seen thousands of installations bear fruit means it is ready to be bonded with server virtualization and software defined networking to enable all the benefits of a software defined data center.

Since SVC is such a crucial piece of the SDDC puzzle we will look in the next section more closely and focus on how the benefits of SDS are realized. Competitively, SVC practically stands alone in the market, which is a true testament to the vision behind it. It will increasingly become the unifying piece for a variety of storage products from IBM and its competitors.

IBM SAN VOLUME CONTROLLER

As a central part of IBM's Spectrum Storage family, what is now known as IBM Spectrum Virtualize software delivers the functionality at the heart of SVC. Its full capabilities require volumes to cover, and we will focus on how SDS capabilities are intrinsic to SVC's fully virtualized implementation.

The primary consideration for any storage solution is *Reliability*. Other necessary attributes cannot make up for downtime or, in the worst case scenario, lost data. The best indicator of a solution's reliability is whether or not it is proven in the field. SVC has tens of thousands of installations with more than a decade of use in production. It has achieved over 5-nines (99.999%) of documented uptime as measured in actual customer environments.

Flexibility is provided by SVC's in-band design that fully virtualizes the physical disks behind its uniform access and management interface. A single pool of all managed physical storage is mapped to logical volumes. These logical volumes provide the basis for the many advanced storage services and automated actions present in a SDS solution. The logical volumes span and allow data to transparently move between different physical arrays and performance tiers within those arrays.

Efficiency is enhanced by in-line "Real-time Compression" that does not rely upon different implementations and efficiency levels in the underlying hardware. The net gain is a 5X increase in storage capacity. Space efficient, point-in-time copies permit the unchanging base (such as a virtual OS) to be shared and only stores the differences that each individual instance requires. Fully integrated thin provisioning ensures all available storage can be used to store data rather than be set aside for future growth. Since it is provided at the software level and not dependent upon each individual array's implantation, it makes provisioning across the pool of storage extremely fast. SDS usage paradigms demand real-time, rapidly provisioned volumes.

Responsiveness is enabled through flexible and nimble provisioning capabilities and delivered through consistent performance of the solution at scale. It doesn't matter how quickly a volume is created if its performance is sub-par. SVC's huge cache, clustered implementation and ability to access the entire storage base as one logical pool combine to deliver impressive performance. SVC combines hardware and Spectrum Virtualize software into a highly scalable, modular solution with up to 4 active-active pairs that can be clustered into a single system with 128 cores, 512 GB of cache and 96 fibre channel ports supporting up to 32 PB of virtual storage.

One particularly valuable capability is "Easy Tier" that optimizes performance and cost-effectiveness by identifying more active data and moving it to faster storage such as flash. Easy Tier can use any faster storage in the virtual pool providing the flexibility SDS demands. Easy Tier is tightly integrated into SVC and can be used with other data-movement, compression and management capabilities.

Simplicity of Operation is necessary to support the constantly changing demands that SDS satisfies. This begins with a powerful and easy to use GUI that provides a single interface to manage multiple storage systems. SVC adds extensive wizards and automated, policy driven management to minimize costly and error prone manual intervention. Built in monitoring supports health and performance visibility over remote, mobile devices. To ease operation tasks in conjunction with server virtualization environments, plug-ins are available to support Microsoft SCOM and VMware vCenter.

To be truly easy to manage one must be able to seamlessly exploit the existing physical environment that SDS is layered upon. Here SVC excels with support for a long list of IBM and competitors' storage solutions. On top of that heterogeneous base, transparent migration between systems supports the incremental management of capacity, upgrades and obsolescence. Easy Tier works on all storage regardless of the different implementations that lie underneath eliminating complex configurations. The end result is that one does not have different islands of capabilities where some are easy and some hard to manage, or the complexity of supporting different management tools and methods.

SVC's modular architecture allows one to start small and grow as needed. The insertion of a single pair of controllers is sufficient to virtualize 8 PB of storage. The first step is to virtualize the underlying arrays of choice, or portions thereof, where unused capacity exists. Data movement capabilities can then be used to migrate data, while it is under use, into the virtual pool. This immediately provides SDS benefits for the migrated volumes. When data on a non-virtualized storage system is evacuated, that system can be re-used by pulling it into the virtual pool. This approach allows one to expand SDS capacity incrementally, without purchasing any additional arrays. New deployments are added to the virtual pool immediately and avoid the stranded HW problem before it begins. This incremental approach allows the administrator to optimize the use of existing and new physical resources for use in the future SDDC world as required, turning a potential liability into an asset.

TANEJA GROUP OPINION

Current IT infrastructures based on hardware-centric architectures of yesterday will not withstand the onslaught of data that the web, social media, sensor-based machines, video surveillance, genomics and other 21st century phenomena have wrought. The data center of tomorrow has to be based on software-defined principles that allow for pooling of resources and abstraction of hardware such that all the resources are brought to bear in a flexible manner and utilized in the most optimized fashion.

As an industry we saw the first signs of this through the server virtualization lens. Server virtualization proved conclusively over the past decade that virtualizing resources at the right level yields major benefits to IT and the user. The industry is in early stages of bringing similar principles to bear on the networking and storage layers. IBM through a brilliant stroke of luck or genius (or both) designed and delivered the SVC product all the way back in 2003. At that time, IBM practically stood alone as most of the major players in the industry eschewed such products. In spite of this, IBM stayed true to its vision and continued to enhance the SVC product not only in its native form but also in new incarnations of Storwize family. And now that the rest of the industry has come to recognize the power of "software-defined everything" and there is a major push towards building the next generation of data centers on these principles IBM stands to finally see a payoff.

We were high on SVC when it first came to market and our position on it has remained unchanged. If anything, we believe the power of this product is now more palpable than ever. It has become the underpinning of SDS, which ultimately is the underpinning of SDDC. What IBM makes of it remains to be seen but it certainly has the ingredients to build a strong foundation based on SVC and Spectrum Virtualize software. We believe customers cannot go wrong in building their SDDC on this foundation.
