TANEJA GROUP
TECHNOLOGY ANALYSTS

P R O D U C T P R O F I L E

Building the Virtual Infrastructure with DataCore SANsymphony-V

February 2011

THE NET NET

There's little doubt that server virtualization delivers a quantum leap forward for IT - it enables multi-generation improvements in processing flexibility and utilization, while simultaneously reducing server costs. But too often, the larger payoff proves to be a pipe dream, held back by remnant ties to the physical world- especially when it comes to storage.

Complex, expensive and inflexible storage hardware proves to be the single biggest constraint holding back the realization of mobility, flexibility and device independence that server and desktop virtualization promise.

In this product profile, we'll examine these roadblocks and take a look at one product that circumvents them - DataCore's SANsymphony-V storage virtualization software. Moreover, while SANsymphony-V tackles storage-related challenges found across all of the popular virtual environments from VMware, Microsoft and Citrix, we'll examine specifically how it unleashes unique capabilities for Windows Server Hyper-V scenarios.

Extending Virtualization to Storage

It is quite clear that server virtualization has forever altered the IT infrastructure. Today, the percentage of IT shops that are not touched by virtualization in one way or another are in the single digits, and look like a rounding error. That isn't to say everything is virtual, but clearly virtualization is everywhere, and it is still ramping in adoption; mission-critical applications are the next playing field, and desktop virtualization is yet in its infancy but rapidly storming the data center.

Such rapid adoption of server virtualization technology is taking place because of the promise of more efficient, more scalable, and more flexible computing. The full promise and vision of better IT requires that virtualization be extended beyond the server. True IT flexibility requires no less than a virtual infrastructure.

What is a virtual infrastructure? A virtual infrastructure is an abstraction of all of the hardware resources that make up a working IT environment so that the organization can realize the following benefits:

- Increased utilization. A virtual infrastructure should promote better utilization of all resources, not just servers.
- Easier use. Whereas server virtualization simplifies server provisioning and

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com

Virtual Infrastructure Capability	Delivered by Server Virtualization Alone?	Issues
Utilization Increased	Yes	Improved, but stops with the processor. Storage utilization can scale out of control as demand for IO and an inability to easily migrate storage can drive storage provisioning and consumption sky high. With growth, demand for disk capacity surpasses the physical equivalent, especially with challenging projects like desktop virtualization.
Easier Use	Limited	Shared storage requirements too often force users to turn to expensive, complex external storage systems that force multi-discipline teams to interact, make storage inaccessible to the server admin, and complicate provisioning and management.
Scalability Improved	No	Storage hardware imposes bottlenecks that complicate scaling, and introduce limitations long before today's processors are fully utilized.
Availability Enhanced	No	Availability risk increases with the consolidation of multiple servers on single pieces of storage hardware, no matter how reliable. Complex systems and practices for data sharing are required to make recovery to other hardware possible.
Flexibility Improved	No	The physical connections to physical resources, especially those as physically-oriented as storage, throw much flexibility out the window.

Table 1: Physical limitations in the virtualized infrastructure.

management, it complicates storage management. A virtual infrastructure must streamline provisioning and management of all resources.

- Enhanced scalability. Unlike in many server virtualization environments, users should be able to virtually scale storage and networking capacity and/or performance as they add or expand virtual servers. And do so in small steps without breaking the bank.
- Enhanced availability. Once again, availability must be infrastructure-wide, and not specific to a single resource.
- Enhanced flexibility. Users should be able to rapidly and incrementally expand, contract or re-allocate all types of IT resources in response to changing business needs.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com

Storage blocks, stumbling blocks

With that in mind, it is equally clear that these capabilities are missing from organizations that have confined their virtualization to servers, leaving them tethered to and inseparable from other physical components in the infrastructure that are not yet virtualized. The most critical system that holds back the virtual infrastructure is storage. Storage tends to be more physical and thus less flexible than any other fabrics or systems on the datacenter floor; the end result is that it has become the biggest cost factor.

Tethering virtual servers to an anchor inextricably set in the physical limitations of storage compromises nearly all of the perceived benefits of virtualization. And the incremental shared storage costs to dynamically relocate virtual machines and virtual desktops often cripple any meaningful progress. In reality, virtual servers are insufficient by themselves without extreme investments in storage. Table 1 highlights some of the glaring deficiencies.

Shared disk IO to physical storage effectively creates a constrictive funnel for all interactions. Every interaction connects to a physical resource and crosses a physical bottleneck. For many readers, this is not news. In a recent TG survey of more than 350 end users, well over half indicated that one or more storage shortcomings were compromising the effectiveness of their virtual infrastructures.

The Economic and Business Value Drivers for Virtual Infrastructures, and How Storage Fits In

Virtual infrastructures are being driven by the need to reduce risks, improve economics and deliver business value. Therefore, a virtual infrastructure needs to drive down capital and operational expenses, speed-up ROI, improve utilization, increase productivity, reduce downtime and endure over multiple generations to future-proof and extend the useful life of investments. We believe, based on ongoing conversations with mid-sized and enterprise server virtualization end users, that by far the greatest opportunities for increasing efficiency and reducing TCO lie in their storage infrastructures, which tend to be far from virtualized today. Removing the dependence on – and simultaneously increasing the utilization of – underlying physical storage systems will help significantly reduce storage CAPEX over time. But virtualizing storage can have an even greater and more immediate impact on OPEX reductions – both on hard costs (related to space, power and cooling) and the softer costs of storage provisioning and management. The opportunity for users to boost storage administrative efficiencies through the right virtualization approach is extremely compelling, and will demonstrate tangible and ongoing cost savings from day one of deployment.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com

The Virtual Enabled

To transform a physical infrastructure into a vibrant, virtual infrastructure, we first need to sever the bonds that tie the remaining unvirtualized domains of the enterprise to the physical world. For most domains, this is steadily happening, as adapters, switches, firewalls and more components pop up as virtual devices. Storage is the last bastion of the physical holdouts. What does it take for storage to step up to the plate and enable the same type of flexibility in the storage layer as the infrastructure now has in the server layer? We see several core characteristics that storage solutions must possess to make storage well integrated and aligned with the virtual infrastructure:

A Single Foundation. The virtual infrastructure is built with adaptability and flexibility in mind – an ability to balance out resources and easily scale when faced with unusual demands. That capability drives many virtualization projects because it promises improved business capability and lower cost of ownership by setting servers free from the constraints of many individual devices that each consume space, require deployment time and effort, and add to management overhead. The virtual server in effect is injected into a virtual computing layer that has been completely removed from the physical computers behind it. Storage needs this capability too. Storage is too often constrained and complicated by being scattered across many individual physical devices with their own unique capabilities, management requirements, and physical limits. The storage foundation for a virtual infrastructure should be one layer that remains consistent irrespective of scale or how the business changes over time.

Utilization and Performance. Storage frequently represents a mismatch for the virtual infrastructure. Not only is utilization challenging to predict, with dire consequences stemming from the danger of outgrowing physical storage; but performance is even more limited, more mismatched, and harder to adjust. At the most basic level, it takes more than simple storage blocks to fit the virtual server infrastructure – storage needs to be carefully matched to demands for capacity and performance, and needs to be equipped with the tools and technology to make sure it stays matched to changing and bursty demands. A single foundation is a start, but the tools and technology to deliver fully optimized performance and capacity, keep it optimized, and return it to an optimized state after storage and/or server alterations is a requirement. This is in fact no less than what we already expect from the virtual server layer.

Feature Integrated. Storage is more than purely the retention of digital bits – businesses require storage built for availability and long-term, enterprise-class protection of stored data. While some hypervisor vendors have tried to attack storage features in the hypervisor layer itself, time has demonstrated this strategy can run out of steam, and further complicate managing stored data. There's a reason dedicated storage systems cost big bucks. Turning to a single storage foundation should not require a step back into relying solely on the hypervisor to deliver all storage

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com

functions. Nor should a single layer require the execution of advanced data management features on underlying storage systems. Instead, users should demand that any technology aiming to enhance virtual infrastructure capabilities come fully feature integrated, and execute all of those features itself - this includes the likes of thin provisioning, capacity optimization, delta-based snapshots, replication, and more.

VM-coupled Storage. Even with a full spectrum of storage features, a specific storage solution designed to simplify the storage infrastructure can fall far from being ideal for the virtual environment. An additional layer in the data center does little good if it does no more than create one more layer in which resources must be manipulated outside of the virtual server layer. Stored data needs to be transparently accessible by the virtual compute instances, with complete independence from physical systems outside of the virtual infrastructure. Moreover, any single instance of stored data should be completely wrapped in storage features that can be executed at the virtual server level rather than solely at the physical LUN level where many such technologies operate today. Each virtual machine in turn should have the entirety of the organization's storage capabilities at its disposal, irrespective of where the workload lives or how the physical infrastructure changes.

Management Integrated. The virtual infrastructure today continues to sit almost entirely apart from storage, with limited visibility, and little ability to do anything storage-related without an entirely separate storage management system, and an entirely different set of skills. Managing storage in the virtual infrastructure shouldn't have to deal with any of these challenges today. The virtual server is already abstracted from hardware, even when running 500 different types of servers. Storage should be abstracted too, and the management of storage for virtual servers should be simplified and put in the virtual server administrator's hands, where it belongs.

Transforming the Physical

One vendor suggests such a transformation would divorce capacity allocation, data protection, and performance management from the underlying hardware storage devices, making the virtual storage infrastructure just as flexible as the virtual servers. That vendor is DataCore, with their SANsymphony-V solution in tow. Their software layer of innovative, high-performance, and feature-rich storage virtualization completes the circle of infrastructure reinvention that server virtualization cannot deliver upon alone.

SANsymphony-V encapsulates even modest, unintelligent storage devices with a new level of "beyond the enterprise" capabilities, while keeping costs to a minimum. It frees the business from relying on complex and costly big-iron storage arrays, along with their associated licensing and litany of tools for every imaginable aspect of provisioning, protecting, replicating, and more. The story is compelling, and worth holding up against our storage requirements for the true virtual infrastructure. But we'll go one step further – since every hypervisor technology has its own unique requirements and limitations, we

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com

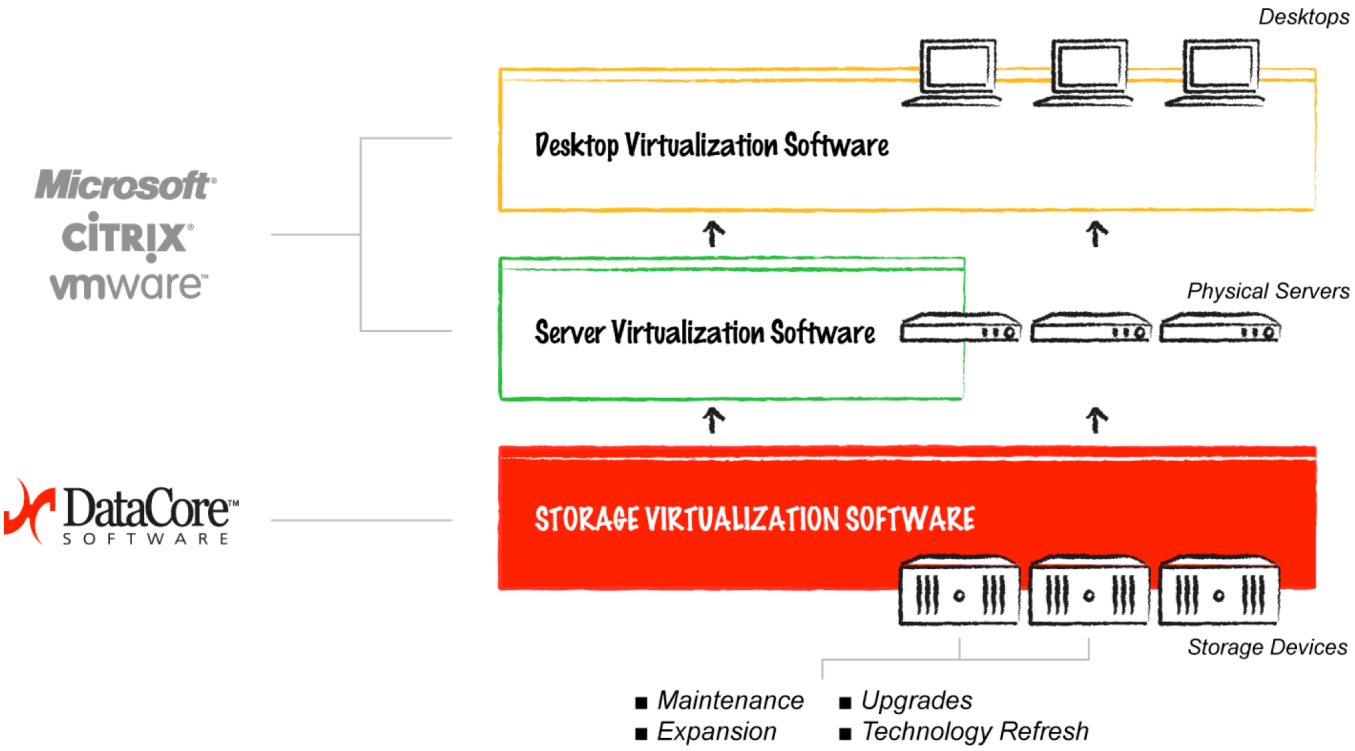

want to shed light on how SANsymphony delivers the virtual shared storage infrastructure for one specific hypervisor, Microsoft's Hyper-V.

Focus on SANsymphony-V in the Hyper-V Virtual Infrastructure

SANsymphony-V runs natively on Hyper-V to virtualize shared storage, speeding up disk access from clustered virtual machines. Moreover, SANsymphony-V integrates an entire portfolio of enhanced capabilities into the virtual storage infrastructure to enhance its availability, utilization and flexibility.

At the heart of SANsymphony-V is a proprietary IO stack that harnesses the underlying processors, memory and I/O channels of the server on which it is installed. In some cases the DataCore software co-resides and runs on the same machine as the VMs, in others it runs outboard on separate servers exclusively dedicated to virtualizing storage. These DataCore nodes can pool anything from locally attached storage to the biggest enterprise arrays, virtualized through nearly any configuration of x86-64 hardware ranging from the latest Nehalem gear to the dusty server in the corner. SANsymphony-V's IO stack uniquely optimizes the caching and pass-through of low latency IO while consuming minimal server resources.

DataCore nodes running on similar or dissimilar hardware provide a scale-out storage foundation for enhanced caching and extended availability on top of a single virtualized storage pool that can simplify the

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com

most complex infrastructures. On top of that optimized platform, DataCore has also turned SANsymphony-V into a uniquely extensive storage services foundation that delivers an alphabet soup list of capabilities: CDP, Asynchronous Replication, NAS, HA, DR, Thin Provisioning, Snapshots, Tiering, and more.

Since SANsymphony-V sits in-band handling all storage interactions, SANsymphony can optionally run on the same hardware as the Windows Server 2008 and Hyper-V Operating Environment, without requiring additional hardware investment, and with the ability to layer advanced services over any storage, internal or external.

While larger customers opt for dedicated DataCore nodes, the software flexibility gives midsized and smaller customers the added cost benefits of making better use of underutilized hardware, or maintaining extreme cost effectiveness where budgets are tight. Moreover, the SANsymphony-V architecture can accommodate any imaginable set of scalability and availability needs. Anytime a storage device or a DataCore node is taken out of service, disk access is automatically rerouted to a redundant node allowing the virtual machines and virtual desktops that depend on those resources to continue uninterrupted.

While DataCore technology has virtualized the storage behind heterogeneous virtual server environments for more than 10 years, SANsymphony-V architecture makes for a distinctly close coupling with the Windows Operating System sitting underneath the Hyper-V operating environment. That blending takes Hyper-V well beyond the realm of pure server virtualization, and into the realm of a true virtual infrastructure. Let's take a look at how SANsymphony-V's pairing with Hyper-V meets our five criteria.

SANsymphony-V Infrastructure

SANsymphony has been built with virtual servers in mind, and is designed to help free them from any lingering ties to the physical storage. Let's see how DataCore delivers on each of the key virtual attributes we expect from a Hyper-V environment:

A Single Foundation. As a separate layer of Windows Server-based infrastructure software that sits in-line between hosts and a set of underlying storage systems, SANsymphony-V provides a single foundation of virtualized storage. SANsymphony-V delivers this virtualization across a heterogeneous set of disks or arrays, providing Hyper-V based servers and desktops with access to a shared pool of scale-out storage. DataCore's virtual disk pooling technology enables users to easily and incrementally scale a storage pool by allocating or de-allocating resources from any type of storage. Existing pools can in turn be split into tiers, with each tier satisfying different price, performance, and availability characteristics. SANsymphony-V removes dependencies on underlying physical storage devices, and by providing a shared and fungible pool of storage, delivers flexibility that fully complements the hardware-independent pool of compute resources enabled by Hyper-V.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com

Utilization and Performance. By pooling storage across the boundaries of heterogeneous physical arrays, SANsymphony-V provides contiguous space for allocation and eliminates stranded pockets of unused capacity. On top of this abstracted pool that is now free from physical restrictions, SANsymphony-V then delivers technologies like thin provisioning, which effectively allows large volumes to be allocated to a Hyper-V virtual server or desktop, while actually only consuming space on disk that contains written data. This capability generally drives up capacity utilization to 75% or more, and by freeing Hyper-V hosts from hard space allocations, relieves administrators from having to constantly provision new storage.

SANsymphony-V provides users with a similar level of flexibility to optimize storage performance. Relying on a virtual layer of pooled storage, SANsymphony-V load balancing allows performance demands to be distributed dynamically across multiple heterogeneous devices and arrays. Taking advantage of its tight integration as a virtual IO stack in Windows Server, SANsymphony-V optimizes performance by intercepting and tuning IO between Hyper-V hosts and storage. Each SANsymphony-V node utilizes underlying system processors and memory as a powerful, inexpensive cache, and shapes IO on both the disk and client side to maximize storage performance. Moreover, multiple DataCore nodes can be aggregated for both availability and performance.

Such capabilities have fed DataCore's notable claims in Virtual Desktop price/performance and even allow DataCore to harness direct attached storage alongside SAN storage with optimal performance and protection for both.

Feature Integrated. With respect to our third facet, SANsymphony-V makes directly available to Hyper-V administrators a rich set of data protection, high availability, and disaster recovery features and capabilities. SANsymphony-V automates the creation and management of incremental copy-on-write snapshots and provides optional Any-Point-in-Time Continuous Data Protection of virtual disks. SANsymphony-V also automates the creation, failover and recovery of replicas and mirrors across heterogeneous storage, thereby insulating virtual servers from device or array-level failures.

The software flexibility advantage of SANsymphony-V becomes evident when we see how it can co-exist with and cost-effectively leverage Windows Server cluster file-sharing and Hyper-V-based capabilities to deliver a high-availability, clustered NAS solution. SANsymphony-V can be easily configured to significantly accelerate performance and add a new level of data protection to Microsoft's Clustered File Shares. This resulting combination of SANsymphony-V and Microsoft is additive; it is simple to set up, requires no additional purchases and best of all it allows organizations to meet both their NAS and SAN requirements from one virtual infrastructure.

Moreover, to facilitate site recovery management and disaster recovery, asynchronous replication enables a Hyper-V administrator to maintain remote, up-to-date data copies without impacting local IO

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com

Two Use Cases in which DataCore SANsymphony-V Extends Microsoft Capabilities

- **Clustered Shared Volumes (CSVs) and Large Scale Metro Clusters:** SANsymphony-V plays a unique role in Microsoft Failover Clusters configurations by providing uninterrupted access to CSVs that are synchronously mirrored between two sites. The active-active shared storage infrastructure in effect stretches metro-wide to provide the highest level of business continuity protection to prevent downtime impacting cluster protected applications (Exchange, SQL, etc.).
- **NAS / Clustered File Shares:** SANsymphony-V leverages built-in Microsoft NAS capabilities and adds needed performance and stretch-HA capabilities to prevent downtime by eliminating the storage-related disruptions that impact Clustered File Shares. SMEs attain a cost-effective way to achieve unprecedented levels of performance and high-availability to meet both their NAS and SAN needs.

performance. These and other SANsymphony-V advanced capabilities may be deployed consistently across a diverse set of underlying physical storage.

VM-coupled Storage. SANsymphony-V's tight integration with Windows Server and Hyper-V enables a set of highly efficient virtual server storage capabilities. For example, virtual disk migration enables Hyper-V data to be transparently migrated across a virtual server infrastructure without disrupting virtual server applications. DataCore also eases data migration in the initial move from physical to virtual via the use of pass-through disks. It simplifies resource load balancing, management, and even opens the door to capabilities like instant recovery on any hardware. The entirety of SANsymphony-V capabilities, including the data protection and business continuity features discussed above, are directly available to Hyper-V administrators, giving them the power to manage the virtual storage pool without needing to call upon storage administrators.

Management Integrated. SANsymphony-V joins storage to virtual infrastructure management making it easy for the Hyper-V administrator to manage all Hyper-V storage from a single centralized wizard-driven management interface, without learning the nuances of the storage equipment. The same GUI and tools are used, regardless of the mix of underlying physical storage systems. The SANsymphony-V management GUI automates a wide range of storage tasks at a macro level. In this way, DataCore enables Hyper-V customers to more rapidly grow their virtual infrastructures without compromising control.

Taneja Group Opinion

Five plus years of server and desktop virtualization efforts – and the repeated hard-learned lessons of countless end users –

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com

have put storage virtualization software in the industry spotlight as the key pre-requisite for a smoothly running, cost-optimized virtual infrastructure. These lessons have been painful for virtualization users, and Hyper-V customers in particular. Physical storage systems alone fall far short of meeting the storage performance, availability and management requirements of virtual machines and virtual desktops.

What's needed is a new approach: a layer of storage virtualization software that is tightly coupled with Hyper-V and that abstracts away the complexities and dependencies of underlying storage and server hardware. For Hyper-V users that would like to build an enterprise-capable virtual infrastructure, DataCore SANsymphony-V is an ideal fit. SANsymphony-V delivers on all five of the fundamental characteristics we consider to be essential attributes of a virtual infrastructure, and does so without requiring significant investments in or changes to other elements of the infrastructure, including virtual servers.

SANsymphony-V takes the pain and cost out of virtualizing Hyper-V infrastructures, and should give IT managers in Microsoft environments the confidence to push forward with virtualization initiatives, and even to begin to virtualize their most critical applications.

We believe that DataCore will not only accelerate server and desktop virtualization efforts, but provide Microsoft, VMware and Citrix customers with the means to build truly virtual infrastructures, and realize long-promised benefits such as increased, infrastructure-wide utilization, scalability and flexibility at a price they can afford.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-2556 Fax: 508-435-2557 www.tanejagroup.com