SANsymphony™-V
Software-defined Storage Platform

Technical descriptions of the device-independent functions provided by DataCore's advanced storage virtualization software

Federates dispersed storage resources to:

- Maximize combined value
- Increase availability
- Speed up performance
- Improve utilization
- Overcome device-specific differences
- Unifomly manage storage infrastructure
- Provide higher levels of service than devices can offer individually

Runs either externally in the SAN, or on the application servers to form a virtual SAN. Can also span both.

SANsymphony-V software solves many of the difficult storage-related challenges raised by server and desktop virtualization in contemporary data centers and cloud environments. The software forms an active, transparent virtualization layer across diverse storage devices to maximize the availability, performance and utilization of disk resources in IT organizations large and small.

The integrated set of centrally-managed data protection, provisioning, caching, replication and migration functions operates uniformly over different models and brands of storage equipment, assimilating current and future technologies non-disruptively. You'll find that SANsymphony-V cost-effectively delivers uninterrupted data access, speeds up applications, and extends the life of your tiered storage investments, while giving you peace of mind.

Figure 1 - Infrastructure-wide functions work across unlike & incompatible storage devices

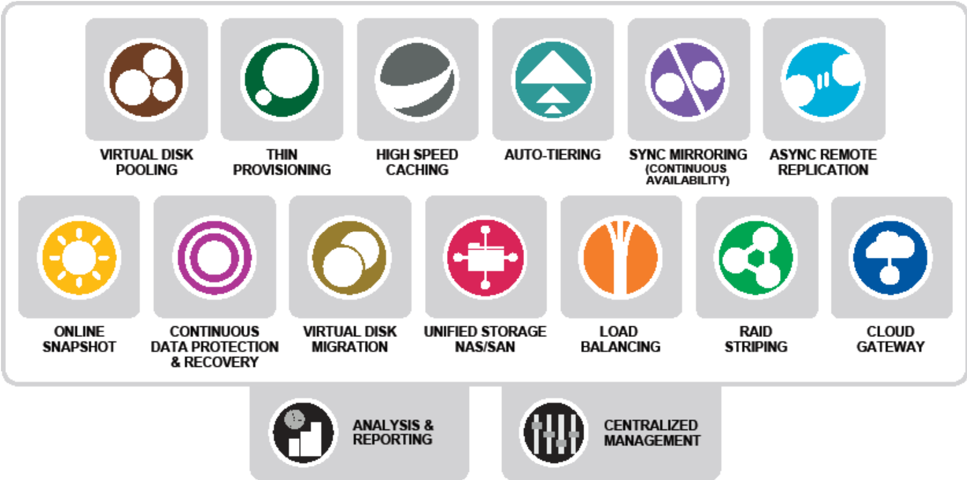

Contents

How it works

First we'll explain how the software fits into your IT environment, and then describe the most salient features in more detail.

Terminology

We'll use the following naming convention throughout the document: "Hosts" refer to servers hosting applications, including database servers, web servers, and file servers. "DataCore nodes" refer to servers running the DataCore storage virtualization software. "Storage" refers to the disk devices, be they internal hard drives, direct-attached disk arrays, solid state drives, flash cards or SAN-connected storage systems. "Virtual disks" refer to the well behaved logical drives presented by SANsymphony-V to the hosts.

Virtualization Technology

SANsymphony-V forms a transparent, scalable virtualization layer across your storage infrastructure in order to enhance its capabilities and centralize its management. The many nuances that distinguish one model or brand of storage from another and render them mutually incompatible no longer stand in the way of using them together. Essentially, SANsymphony-V federates dispersed resources to:

a) maximize their combined value, b) to overcome device-specific differences and c) provide higher levels of service than the devices can offer individually.

You have complete freedom to select the storage devices that make up your physical storage pool; mixing and matching as you see fit, even migrating between old and new in the background. SANsymphony-V works with all the popular models and brands of disk devices supported on Windows Servers allowing you to strike a balance between high-speed, premium-priced resources and lower cost, higher capacity gear.

The same flexibility applies to packaging. Larger environments requiring many Terabytes or Petabytes of disk space generally use external disk arrays to house hundreds of drives in very high-density packages. On the other hand, compact applications in cramped rooms or hardened mobile enclosures can take advantage of drives housed inside the DataCore nodes. Still others start small with internal drives and expand externally. It's common to split the overall disk capacity evenly between a pair of redundant nodes for high-availability and load balancing. The software accepts direct-attached storage (DAS), SAN arrays and solid state disks / flash cards with equal ease, connecting to them via any of the standard disk interfaces (SATA, SAS, Fibre Channel, iSCSI, etc.,).

Figure 2- Active, transparent virtualization layer maximizes availability, performance & utilization

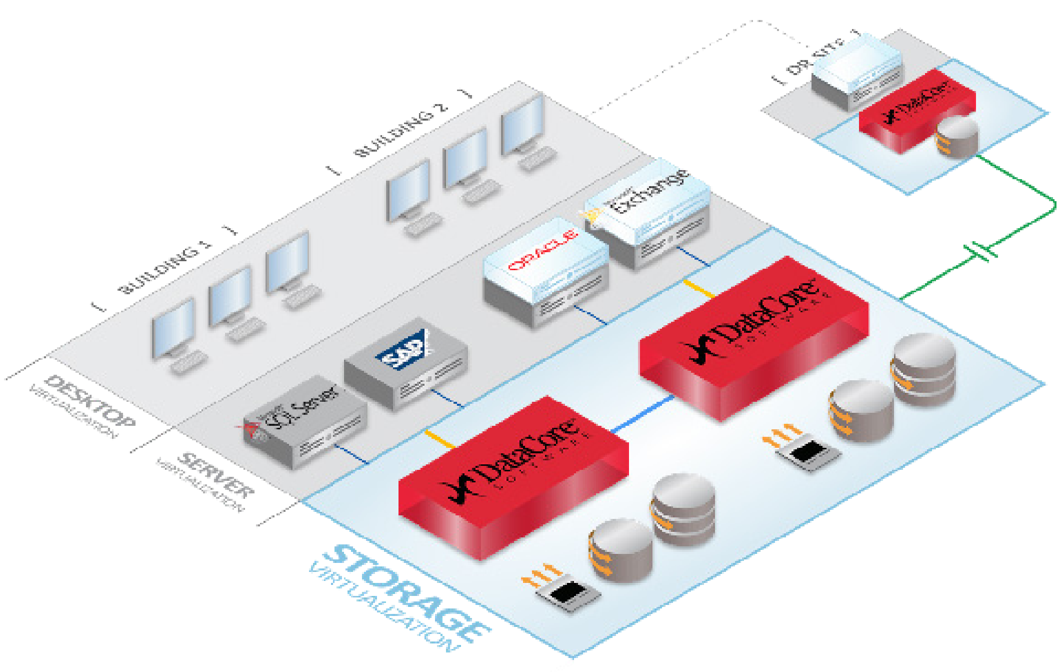

Gain uninterrupted access by stretching mirrored nodes apart

Many customers select DataCore software in order to attain business continuity in the face of much storage-related upheaval. SANsymphony-V shields applications from planned or unplanned outages in the underlying components by providing uninterrupted access to the virtual disks. The software will synchronously mirror virtual disk updates between nodes in separate rooms connected to segregated storage so they won't be exposed to the same facility-related hazards.

Stretching inter-node distances up to 100 kilometers apart via metropolitan area networks (MANs) decreases the chances that an ordinary misfortune such as a roof leak, fire, air conditioning failure, construction crew mishap or flood will affect both sites. A virtual disk in these high-availability configurations logically appears to hosts as a single, well-behaved, multi-ported shared drive, although it is really made up of two widely-separated mirror images updated in lock step.

Entire sites, nodes, disk subsystems, channels and other components of the environment may be taken out-of-service, upgraded, expanded and replaced without disturbing applications or scheduling downtime.

Provisioning disk space

The SANsymphony-V system administrator carves out "virtual disks" on demand from the physical disk pool according to the capacity, availability and performance needs of each workload. For instance, some groups of virtual disks may be defined to be cached, synchronously mirrored, remotely replicated and thinly provisioned. In the background, the DataCore software will draw on multiple real devices and the necessary computing and network connections to meet those requirements. Auto-tiering functions dynamically determine which storage class will best meet the selected storage profile (Critical, High, Normal, Low, or Archive.)

Hosts connect to the DataCore nodes over iSCSI and/or Fibre Channel just as they would connect to an external SAN disk array. Network ports and I/O channels may be configured from a wide range of host bus adapters (HBAs) and network interface cards (NICs) including iSCSI virtual SANs inside virtualized servers. The protocol between the hosts and the DataCore nodes may be different from what the nodes use to connect to the physical storage. For example, SANsymphony-V can essentially bridge iSCSI connections from hosts to trays of SATA and SAS drives.

Hosts only sees those "virtual disks" explicitly shaped and assigned to them over designated ports. In clusters where several hosts must share certain disks, the same virtual disks may be assigned to multiple hosts even if the physical drives in the pool are not multi-ported (shareable).

To maximize disk utilization and eliminate wasted space, DataCore implements very granular, thin provisioning techniques and space reclamation features.

SANsymphony-V can serve virtual disks to physical hosts running any of the popular open operating systems. It also serves virtual machines hosted by the mainstream server hypervisors from companies like Microsoft, VMware and Citrix. Integration with host-based workflows and other infrastructure-wide system management tools is simplified through an extensive scripting library of PowerShell commandlets and SNMP traps, as well as task scheduling wizards. Real-time / historical charts, e-mail event notifications and configuration reports provide additional insights into the health and performance of the virtualized storage infrastructure.

Cache to overcome I/O bottlenecks and help applications go faster

DataCore taps each node's processing, memory and I/O resources to carry out advanced functions across all of the storage devices under its command. SANsymphony-V reserves up to 1 Terabyte (TB) of random access memory (RAM) per node for SAN-wide "mega caches".

Thanks to the sophisticated multi-threaded caching algorithms, data written to or read from disks move swiftly into and out of the caches, harnessing the full potential of each node's high-speed multi-processors.

Caching makes application run faster than they would had they accessed the disks directly, whether connected to top-of-the-line storage systems or lower end gear.

Scale up and Scale out

The intrinsic scalability designed into the SANsymphony-V software gives you the confidence that the DataCore architecture can grow gracefully to meet your future needs. This is accomplished by scaling both up and out.

"Scale up" is accomplished simply by upgrading the underlying server platforms acting as storage virtualization nodes to more powerful systems. The nodal improvements yield higher IOPS and more throughput.

It's very similar to the performance boost you get when moving your web or mail services to a faster machine.

You may also scale out your SANsymphony-V infrastructure beyond two nodes to keep pace with the natural capacity expansion and rising workloads encountered in larger data centers. This may be done one node at a time. Up to 16 nodes may be managed from a central console.

Each additional SANsymphony-V node brings more ports, cache and processing power to address the larger demand. More nodes provide the necessary bandwidth and connectivity to service more hosts and attach more disks to the environment.

Additional nodes also increase the overall redundancy offered by the infrastructure. For example, you may configure 3 nodes which share a common fourth node for failover (more on this in the Continuous Availability section).

Automatically optimize disk access across tiers for best utilization

Disks with different price/performance attributes can be organized into separate storage tiers. For example, the fastest tier may be composed of Solid State Disks (SSDs), with lower tiers made up of SAS and SATA drives, respectively. The same tier may include similar units from different manufacturers acquired over time.

SANsymphony-V automatically promotes frequently accessed disk blocks to faster tiers and demotes less active blocks to slower, more cost-effective drives based on the IT team's criteria. You may override auto-tiering policies by explicitly defining which disks should be used for special workloads or confine the selection of tiers to a narrower set of resources.

Groups of disks may also be segregated for specific requirements, such as test / development or to isolate multiple subscribers (multi-tenancy) in public, hybrid and private clouds.

Advanced features such as caching, synchronous mirroring, asynchronous replication, snapshots, thin provisioning and CDP operate across mutually independent devices within the same tier or across tiers. For example, we find customers prefer to take online snapshots of virtual disks in tier 1 and place them in a tier 2 or tier 3 pool to avoid consuming premium resources for the backup copies. Similarly, tier 1 devices may be remotely replicated to a disaster recovery site that only has tier 2 devices, possibly from a different supplier.

Designed to tolerate failures

Despite our best efforts to choose the ultimate hardware/software platform, we must recognize that no single hardware or software component in the storage infrastructure can possibly be trusted to be bulletproof.

SANsymphony-V has been designed with this reality in mind. In order to maintain continuous access to your storage, we simply ask that you institute a short list of common sense best practices which have been fine-tuned over the past decade. Some of these practices are built into our configuration wizards, whereas others must be adhered to in setting up the broader environment.

It's important to note that these architectural recommendations not only prevent disruptions due to equipment failures and facility outages, but they also allow you to take components out of service for routine hardware, firmware and software updates without suffering planned downtime.

Proven Architecture

Perfected over tens of thousands of installations, SANsymphony-V Release 9 represents the latest enhancements in DataCore's proven architecture.

Its strict adherence to established interfaces for communicating with disks, networks, operating systems and server hypervisors ensures that you can easily adapt to take advantage of the many innovations that the industry will introduce in the next decade, and to do so without disrupting the virtualized environment on which your business depends.

In the following pages, you'll learn more about each of the major features that make up the SANsymphony-V storage hypervisor. Note that some features may be licensed as separately-priced options.

Replicate offsite for Disaster Recovery and Migration

For additional offsite disaster recovery needs, or simply to migrate data between sites, SANsymphony-V replicates disk updates asynchronously over unlimited distances to other DataCore nodes using conventional IP lines. You may also reverse the direction of replication to restore the original site after the danger has been averted. For additional confidence, remote restoration procedures may be tested regularly while normal production processing goes on at the primary data center.

Choosing the Hardware / Software Platform

The core SANsymphony-V feature stack runs on dedicated server-class machines arranged in high-availability configurations to eliminate single points of failure. Most organizations purchase new servers for this role, although some initially repurpose perfectly good application servers vacated by consolidation projects. The SANsymphony-V software may also co-reside with virtual machines on servers where spare processor, memory and I/O resources are available.

Several factors help determine the specific x86-64 hardware chosen for each scenario; among them, I/O Operations per Second (IOPS) and throughput requirements (Megabytes / second), level of redundancy desired, budget constraints and brand preferences. You may choose from any of the popular server models from manufacturers like Dell, Fujitsu, HP, IBM, Intel, Oracle, and SuperMicro, as well as purpose-built appliances from private-label system builders across the world.

In any case, you always have the option to upgrade the underlying hardware with more powerful systems when newer, faster machines are deemed appropriate. There is no software throw-away. The DataCore code and its licenses are fully portable between server platforms for the utmost flexibility and investment protection.

The core SANsymphony-V code executes on top of a locked-down image of the Windows Server 2012 or 2008 R2 operating system (OS).

The OS provides a robust and deterministic platform for memory management, networking services, driver integration, fault-isolation, workflow orchestration and other basic services.

As importantly, the OS provides the portability needed to move across generations of servers and the compatibility to interface with the most popular storage devices.

Pooling tiered storage

Virtual Disk Pooling

- Federate widely distributed storage assets
- Split pool into tiers with different price/performance/ capacity characteristics
- Create and assign virtual disks of desired sizes
- Define access rights (read/write/shared)
- Use templates to rapidly create many virtual disks with common properties
- Explicitly assign virtual disks to specific hosts or groups of hosts
- Expand capacity without downtime
- Decommission disks non-disruptively
- Automatically rebuild virtual disks on a healthy disk pool when the source pool is out-of-service
- Reclaim stranded disk space

Consolidate like or unlike disk resources

Virtual disk pooling is DataCore's overarching feature responsible for consolidating storage capacity from like or unlike disk resources. As noted earlier, the pool may encompass a variety of brands and models of disks, effectively creating storage tiers with different price/performance/capacity characteristics.

Pooling is fundamental to storage virtualization, enabling virtual (or logical) disks to be rapidly created from blocks of space on the physical devices. Using a central administrative interface, these virtual disks can then be assigned to storage consumers throughout the physical or virtual Cloud with specific access permissions; possibly shared among different hosts, virtual machines or clustered applications.

The upper limit on a SANsymphony-V storage pool is well into the petabytes, depending on the product level chosen.

DataCore storage pools can be composed of different brands and models of disks and managed as a central resource to meet the varied needs of applications. Before you exhaust the capacity of devices in a specific tier, you can add more, possibly using the next generation product from the same supplier or a different manufacturer that offers similar technology with more favorable price/ performance. And as new storage consumers arrive, they can immediately tap into the storage pool via their iSCSI or Fibre Channel SAN connection.

Unified Storage

- Highly available storage for NFS/ SMB (CIFS) file shares
- Simultaneously handle disk block and file level services

Converge NAS/SAN

Clustered file systems may be layered on top of redundant DataCore nodes to achieve highly available Network Attached Storage (NAS) with fast, uninterrupted access to NFS and/or SMB (CIFS) file shares. Both the file sharing role and the failover clustering features are available directly from the underlying Windows Server OS. You can configure DataCore nodes as "unified" storage by providing hosts with concurrent access to the file shares over LAN ports while the SAN ports fulfill direct requests for "raw" disk blocks.

Thin Provisioning

- Appears to computers as very large drives (e.g. >2 TB disks)
- Takes up only space actually being written to
- Dynamically allocates more disk space when required
- Notifies you when it's time to add capacity
- Reduces need to resize LUNs
- Reclaims zeroed out disk space

Allocate just enough space, just-in-time

In these tougher economic and eco-friendly times, thin provisioning gets a lot of press and rightly so. DataCore's implementation spans all physical devices in the pool, and like its other features, is not dependent on any one host, operating system or server hypervisor.

There are two parts to thin provisioning. On the one hand, the storage consumer may be assigned very large logical drives to satisfy their long range wishes, while in reality, the space is only consumed when it is actually written to. This solves the problem of guessing wrong and wasting space when an application first asks for a lot of capacity even though they may never use it.

Rather than tie up all that space, the software allocates only small chunks of disk blocks as needed, just-in-time. This works out well with virtual machines that like to mount volumes that are much bigger than they actually need. The technique is sometimes referred to as oversubscription since the virtual capacity assigned to the hosts may exceed the physical capacity on the data center floor.
You may also reserve a portion of the virtual disk capacity without explicitly allocating it, to set it aside for later use, so no one else can unintentionally use it.

For example, we can logically assign 6 TBs to the virtual disks, but the applications have only written to less than ½ TB, so very little capacity has been consumed. When real consumption starts to exhaust the physical space, the software alerts the administrator to add more. You may set different alarms (information, warning, critical) as space starts to dwindle.

Once more drives arrive, you instruct the software to add the new capacity to the existing pool, giving you more room for expansion. The additions take place non-disruptively in the background. Moreover, the additional space could come from completely different equipment given your new choices at that time.

DataCore also includes utilities to reclaim space that was temporarily consumed on thinly provisioned virtual disks and later zeroed out by the host. This happens with scratch disks and other large temporary files.

Auto-Tiering

- Select tiers based on price/performance/capacity
- (Flash, SSD, FC, iSCSI, SAS, SATA, ...)
- Dynamically migrate blocks among classes of storage
- Best location at the time determined by:
  - » Access frequency
  - » User preferences

Best tradeoff between performance & cost

The science of automated storage tiering distills down to monitoring I/O behavior, determining frequency of use, then dynamically moving blocks of information to the most suitable class or tier of storage device. SANsymphony-V software automatically "promotes" most frequently used blocks to the fastest tier, whereas least frequently used blocks get "demoted" to the slowest tier. Everything else floats to the middle.

Of course, there will be exceptions, especially when you need to assign high performance storage to an infrequently used volume, as in special end-of-quarter processing.

In these cases, you can pin specific volumes (virtual disks) to a tier of your choosing, or define an "affinity" to a particular tier. Only if that tier is completely exhausted, will a slower tier be chosen.

In order to appreciate the financial motivation for auto-tiering, consider the price/performance differences among the three most popular classes of disks in use today. Leading with the highest performance at the steepest price are Solid State Disks (SSDs). Next are Serial Attached SCSI (SAS) drives, followed by the most economical Serial ATA (SATA) disks.

For random disk read patterns, SSDs are said to be 25 to 100 times faster than SAS hard disk drives (HDDs) at roughly 15 to 20 times higher cost per Gigabyte. In practice, SSDs substantially reduce the number of HDDs required for heavy random I/O pattern. Just one SSD PCI I/O card may yield the equivalent of 320 hard disk drives. That's 300 times less hardware to house, maintain, cool, and watch over.

In general, you wouldn't want to waste the premium-priced SSD capacity on infrequently used blocks, so the DataCore software will keep the least active information on the lowest cost, capacity-optimized SATA disks corresponding to Tier 3. On the other hand, when the auto-tiering algorithms detect sustained heavy use of certain other blocks, they will naturally bubble up to the Tier 1 SSDs.

Perhaps your diverse storage pool consists of top-shelf, premium disk arrays from one vendor, mid-range RAID subsystems from another, and relatively inexpensive racks of unintelligent disk drives (JBODs).

You define what constitutes each tier, relative to the other members of the pool. You may set up as many as 15 tiers, although 3 or 4 are usually enough to help you make good trade-offs.

As new storage products arrive on the market, the top rung could shift down to make room for an even higher performing disk technology – again, without disrupting applications.

Performance

High-speed Caching

- Accelerates disk I/O response from existing storage
- Uses x86-64 CPUs and memory from DataCore nodes as powerful, inexpensive "mega caches"
- Anticipates next blocks to be read, and coalesces writes to avoid waiting on disks

Speeds up performance

High-speed caching has long been a potent differentiator for DataCore's products. If you have been conditioned by other suppliers to believe that any form of virtualization brings a performance penalty, think again. In the process of virtualizing disks, DataCore software accelerates reads and writes by leveraging the powerful processors and large memories of the x86-64 servers on which it runs.

Up to 1 Terabyte of relatively inexpensive cache may be configured on each DataCore node enabling it to turn around disk requests at electronic memory speeds. Caching essentially recognizes I/O patterns helping it anticipate which blocks to read next into RAM from the back-end disks. That way the next request can be fulfilled quickly from memory absent mechanical disk delays.

When hosts write to a virtual disk, the information goes quickly into cache memory and mirrored to another DataCore node to ensure multi-cast stable storage before being acknowledged. It is later de-staged to disk without impacting apps, often grouped (coalesced) with other writes to minimize delays when storing the information to disk. You may also define certain virtual disks to be in write through mode so the blocks are destaged to disks before being acknowledged.

The simplest way to understand caching on a DataCore node is to view it as a level 1 cache that can respond in less than 20 microseconds, whereas the caches on the disk array may be slower and pricier, taking somewhere in the hundreds of microseconds. Both caches aim to hide the much longer delay of the physical disk I/O which is in the order of 4000 to 6000 microseconds (or 4 to 6 milliseconds). The DataCore caches play an important role when using very fast flash / SSD technology, extending their useful life by reducing the write traffic.

The software uses the CPUs on the DataCore nodes to rapidly poll for inputs and immediately services input or output requests out of SAN-wide caches.

Polling replaces slower (higher latency) interrupt servicing techniques that take much longer to recognize I/O events. The quicker the CPU, the faster the software can turn around an I/O. The same polling method is used whether fielding requests from the host computers or responding to the back-end disks. Write coalescing reduces the delay in writing to disk by grouping inputs, whereas pre-fetching into cache anticipates blocks to be read given earlier reads from that section of the disk. The adaptive algorithms have been perfected over the past decade.

RAID Striping

- Circumvents drive failures
- Spreads I/O across multiple spindles
- Offloads RAID 0 & 1
- Supports popular RAID devices

Better protection and performance

RAID (or Redundant Arrays of Independent Disks) is a common way to gain better performance and protection by spreading I/O across multiple disk spindles. Virtual disks may be intentionally striped across several physical drives or simply mapped to a logical RAID drive supplied by the underlying disk subsystems.

Thin provisioning takes advantage of striping to dynamically allocate more disk space when the initial set of drives runs out of room.

Performance Analysis

- Assists with tuning and scaling decisions
- Reveals potential bottlenecks
- Heat maps indicate unusually high disk activity on specific drives

Tracks and charts storage infrastructure behavior

The built-in administrative interface provides real-time status and historical charts on I/O ports, virtual disks, and other parameters across the virtual storage pool.

This information may be used to reassign overloaded channels or back-end drives. It also helps determine when additional resources are necessary to scale the solution to match growth. A multitude of performance counters can be charted.

Load Balancing

- Overcomes typical storage-related bottlenecks
- Spreads load on physical devices using different channels for different virtual disks
- Detects disk "hot spots" and transparently redistributes blocks across the pool
- Automatically bypasses failed or offline channels

Improve response and throughput

Load balancing across the back-end channels into the physical storage pool complements caching to improve response and throughput. The host computers may also be doing host-based load balancing across their ports. Load balancing helps to overcome short-term bottlenecks that may develop when the queue to a given disk channel is overly taxed, or when one channel fails or is taken offline. SANsymphony-V regularly fine tunes itself by redistributing disk blocks which may be overloading a specific disk drive.

Continuous availability

Synchronous Mirroring & Auto Failover

- Architect N+1 redundant grids for continuous availability
- Eliminate SAN or storage as a single point of failure when combined with host MPIO or ALUA drivers
- Enhance survivability using physically separate nodes in different locations
- Mirrored virtual disks behave like one, multi-ported shared drive, while automatically updating the two copies simultaneously
- Establish a common hot site for multiple data centers in the same metropolitan area
- May be combined with clustered file shares to achieve high-availability NAS

Real-time I/O replication for High-Availability

When it comes to non-stop storage access, synchronous mirroring deserves most of the credit. It handles the real-time replication of I/Os for the ultimate in continuous availability. Having two nodes store the data simultaneously in conjunction with the host's multi-path I/O (MPIO) or Asymmetric Logical Unit Access (ALUA) drivers eliminates single points of failure or disruption.

SANsymphony-V allows you to configure redundant storage pools by synchronously mirroring between DataCore nodes. For any mirrored virtual disk, one DataCore node owns the primary copy and another holds the secondary copy. Those are maintained in lock step.

Under normal operation, all read and write requests issued to that virtual disk will be serviced by the primary copy. The secondary copy need only keep up with new updates arriving from the mirroring function. Generally, nodes are configured to control primary copies for some virtual disks and secondary for others, thereby evenly balancing their read workloads. Alternatively, N+1 configurations consisting of 3 or more nodes, may rely on a common node to keep the mirrored backup copy.

Should any errors be encountered on the preferred path, the host's MPIO or ALUA drivers automatically fail over to the alternate path without disrupting applications. The same is true if a node needs to be taken out-of-service for maintenance or upgrades. If the node encounters any problems trying to reach the disks where the primary copy is stored, it will redirect the request to the node holding the mirrored copy.

From a physical standpoint, best practices call for the DataCore nodes to be maintained in separate chasses at different locations with their respective portion of the disk pool so that each can benefit from separate power, cooling and uninterruptible power supplies (UPS).

The physical separation reduces the possibility that a single mishap or facility problem will affect both members of the mirrored set. Round-trip network latencies govern the maximum distance between mirrored nodes. Current technologies support inter-node distances up to 100 kilometers. The actual limits depend on an application's sensitivity to delays and the network latency experienced between locations.

Where cross-site redundancy is not practical, two or more nodes may be configured to provide multiple paths to the same shared, multi-port disk array. This configuration ensures high-availability across the DataCore storage virtualization layer, but relies on the storage system to provide internally redundant components (disks, fans, power supplies ...) to safeguard its contents. Keep in mind, however, that despite the array's component-level redundancy, a facility problem could render the entire unit inoperable and preclude hosts from getting to their disks.

Migrations & Upgrades

Virtual Disk Migration

- Perform non-disruptive hardware disk upgrades
- Clear and reclaim space occupied on the original drive(s)
- Decommission active physical disks non-disruptively from pools and redistribute their contents among the remaining disks in the pool
- Move virtual disks from one pool to another without interrupting applications
- Provide pass-through access to drives previously used on other systems

Transparently move contents from one disk to another

Another by-product of making storage interchangeable is the ease of relocating data from one storage system to another, non-disruptively. We refer to it as virtual disk migration or volume replacement. It comes in handy during hardware upgrades as you seamlessly roll in new equipment in place of older hardware. Having copied the contents to the new drives in the background, the software will zero out the original disk and reclaim its space into the free pool. You can replace storage devices in the background while applications continue to run uninterrupted. Sometimes the older equipment can be allocated as a slower tier device for less demanding requirements, or it can be decommissioned altogether. The choice is yours.

Occasionally the need arises to shrink disk pools – especially in Cloud environments with subscribers that come and go. This too is supported in the background.

SANsymphony-V also helps you to take previously formatted disks connected to other systems and place them under the DataCore virtualized storage environment temporarily or permanently. This is accomplished by passing through access to those drives once they are attached to a DataCore node. Generally, disks drives will be in active pass-through mode, while their contents are mirrored in the background to another storage system. At that point, the original drive can be wiped clean and its capacity absorbed into the larger disk space pool to complete the transition. In purely migratory cases, the original drives may be retired.

Backup & Restore

Continuous Data Protection (CDP) & Recovery

- Dial back to restore an arbitrary point-in-time within a 14 day time frame
- Logs and timestamps all I/Os to the selected virtual disks
- No need to quiesce or interrupt applications
- Group virtual disks to synchronize their logs
- No host agents required
- Easy to turn on and revert from

Return to an earlier point-in-time without taking explicit backups

Organizations are frequently asked to undo data modifications that had undesirable consequences on their business. The changes may have been in error or they may be byproducts of viruses or other malware.

Falling back to the latest snapshot or backup could mean losing a significant number of updates that transpired before the problem hit. Continuous Data Protection (or CDP) provides a clever way to restore a point in time between the longer interval covered by your snapshots and backups.

CDP continuously logs and timestamps I/Os written to designated virtual disks allowing you to revert back to a time of your choosing within a 14 day period. It's like an undo button.

Online Snapshots

- Recover quickly at disk speeds to a known good state
- Eliminate back-up window
- Provide "live" copy of environment for analysis, development and testing
- Save snapshots in lower tier, thin-provisioned storage without taking up space on premium storage devices
- Synchronize snapshots across groups of virtual disks
- Trigger from Microsoft VSS-compatible applications and VMware vCenter

Capture point-in-time images quickly

Once you've tried online snapshots, you can't live without them. Snapshots capture a known good point-in-time that may be used for several purposes without scheduling lengthy back-up windows. It may give you a recovery point to undo a patch or file deletion. Or it may be used to feed business intelligence analysis. They are also commonly used to verify new software enhancements in test and development before being put into production.

Snapshots are invaluable in cloning working system images to provision identical new servers or new virtual desktops. Although snapshot utilities are commonplace in operating systems, server hypervisors, backup software, and disk arrays, capturing them at the SAN level affords some major advantages. For one, there is no dependency on host software. Nor does it consume host resources. And you don't need mutually compatible disk arrays.

You can snap the contents of disks on a tier 1 array and place it on a tier 2 or tier 3 device rather than tie up expensive space on the top-of-the-line equipment.

In fact, DataCore snapshots reduce capacity consumption in two ways: They can be set to capture only changes from the original source (differential snapshot) and they may reside on thinly-provisioned volumes.

DataCore snapshots may be triggered from Windows Virtual Shadow Copy (VSS-compliant) applications that temporarily quiesce the application and flush OS caches to ensure a consistent recovery image. Similarly, they may be taken on VMware datastores, whose virtual machines are momentarily quiesced.

There's some impressive magic that goes into making DataCore snapshots fast and lean. Part of it comes from the realization that generally, only a few disk blocks really experience much change. With that in mind, the snapshot can reference the original unchanged blocks most of the time without taking up any extra room. We need only keep a separate, unmodified copy of those blocks that get updated. And that copy is taken on the first change to that block, effectively giving us the frozen point-in-time. Think of the software as a traffic cop. If the block is unchanged, read from the source. If the block has been changed, read from the copy of the original contents kept in a separate bucket.

DataCore gives you several snapshot variations. We just described the "Differential Snapshot". You can also ask it to completely clone the source disk in the background. We refer to it as a "Full Clone."

You may refresh the snapshot to a later point-in-time using the "Update Snapshot" command, or you can completely replace the source with an earlier snapshot using the "Revert" command. That one is reserved for those cases when you have to undo changes.

Offsite disaster recovery & migration

Remote Replication

- Perfect for disaster recovery, business continuity or inter-site migrations
- Only needs a basic IP connection to the secondary sites
- Bidirectional transfers
- Asynchronous protocol to tolerate transmission delays and link outages
- Central site to remote sites (1-to-many)
- Remote sites to central site (Many-to-1)
- Many-to-many
- Compressed, multi-stream transfers for fastest performance and optimum use of bandwidth
- Allocate link bandwidth among replicated disks based on user-defined storage profiles
- Expedite initialization of remote site using transportable media
- Test disaster recovery readiness without impacting production

Maintain distant copies up-to-date without impacting local performance

Earlier, we discussed how synchronous mirroring fits into our solution for continuous availability. We spoke of an upper limit near 100 kilometers between mirrored nodes before round-trip latencies make lockstep replication impractical today. In other words, if you tried to synchronously mirror at longer distances, applications might time out waiting for the acknowledgement to come back from the remote end.

DataCore's remote replication function addresses requirements for secondary copies to be housed beyond the reach of synchronous mirroring, as in distant disaster recovery sites, branch offices and satellite facilities.

It relies on a basic IP connection between locations, and works in both directions. That is, each site can act as the disaster recovery spot for the other.

The software operates asynchronously, meaning that it does not hold up the application waiting on confirmation from the remote end that the update has been stored in both places. Instead, it offers to do its best to keep up with changes at the local site, but makes no guarantees. It's far better than trying to constantly make backup tapes and ship them to a safe house or paying extra for point-products to handle only this task. The advanced protocol handles prolonged transmission delays or link outages allowing you to set the priority of which virtual disks should be allocated the most bandwidth.

You can quickly get a remote site initialized by cloning the primary site's disks onto transportable media and shipping them to the disaster recovery center. Then apply the changes that transpired while in transit. To help you build strong, verifiable disaster recovery practices that you can confide in, SANsymphony-V enables you to test restoration at the remote site while production updates continue to arrive. Any changes made during the simulated recovery are then discarded and the standby copies refreshed with the newest updates.

Advanced Site Recovery (ASR)

- Reverses direction of replication from the disaster recovery (DR) site to the primary datacenter
- Universal coverage for heterogeneous scenarios
- Same automated process for virtual and physical systems
- Integrated with VMware Site Recovery Manager (SRM)

Expedite central site restoration

ASR automates and radically simplifies how a remote IT facility takes over workloads from a central site in the event of a disaster or scheduled outage. As importantly, ASR takes care of updating the main IT site's storage pool with the changes that transpired at the DR location to quickly get the main site back online.

Unlike other approaches DataCore covers both virtual and physical servers without taxing applications, hosts, or server hypervisors. Nor does ASR depend on duplicating expensive equipment offsite, such as top-of-the-line disk arrays and specialized networking gear. DataCore allows organizations to leverage readily available IT assets, often differing between sites, to minimize or eliminate the disruption and data loss attributed to planned and unforeseen site outages. ASR builds on our virtual disk provisioning, asynchronous remote replication and online snapshot features to circumvent downtime and expedite central site restoration.

Integration with VMware Site Recovery Manager (SRM) is achieved through the SANsymphony-V Storage Replication Adapter (SRA).

Centralized management

Unified Storage Management across the Infrastructure

- Intuitive to set up and operate
- Automates repetitive tasks
- Self-guided wizards and troubleshooting tips
- Comprehensive diagnostics
- Configurable views of real-time system behavior and historical performance charts
- Group operations for disks dependent on or related to each other
- Resource allocation reports (exportable for billing / chargebacks)
- Configuration reports (channels assignments, parameters used, etc.)
- Role-based, administrative permissions with audit trail of configuration changes

Control and monitor pooled storage resources from a central console through device-independent tools

SANsymphony-V centralizes provisioning, control and monitoring of distributed storage pools through a remote management console.

The intuitive administrative interface offers self-guided wizards to direct unfamiliar operators through best practice procedures. All of the integrated functions are managed from the single pane of glass without having to deal with model- or vendor-specific variations. The graphical user interface (GUI) is highly configurable to accommodate individual preferences. The same comprehensive management tools operate across like or unlike storage devices.

Workflow integration is automated through rules-based task schedulers. To assist in tuning, the software is well instrumented with real-time charting of performance characteristics. Thorough diagnostics accompanied by e-mail alerts and SNMP traps help sidestep potential problems as well as assist with troubleshooting and root-cause analysis.

Integration with popular systems management and monitoring packages is supported as well. Among them: Microsoft System Center, VMware vCenter, and Hitachi IT Operations Analyzer.

Together, these capabilities let you more easily and quickly accommodate expansion, upgrades and other changes in your environment while building on the IT skills you already have in place.

Microsoft Integration

- Runs natively on Windows Server 2012 OS and may be configured on the root partition of Hyper-V
- Monitoring Pack for System Center Operations Manager is available
- Compatible with System Center Data Protection Manager (DPM)
- ODX support offloads hosts from copy-intensive operations
- SCSI Unmap to reclaim cleared out blocks from deleted files
- RDP-accessible management console
- PowerShell commandlet library
- In the file sharing role, Windows clustered NAS gain uninterrupted access to mirrored storage pools

Since SANsymphony-V runs on top of the Windows Server operating system, integration with the extended Microsoft platform comes naturally. The many points of integration cover a number of areas ranging from systems management to virtual machine provisioning and online backups.

Windows Server 2012 Certified

VMware vSphere Integration

- vCenter Plug-in
- SRM Storage Replication Adapter (SRA)
- VAAI support for full copy, block zeroing and hardware-assisted locking functions
- SCSI Unmap support to reclaim blocks from deleted files
- Native vCenter integration

VMware vSphere system administrators can harness and control the complementary services of SANsymphony-V without leaving their familiar vCenter Server console. Close integration between DataCore and VMware ensures a coordinated approach to realizing fully virtualized environments. You may initiate, schedule and monitor key DataCore storage provisioning and data protection services (snapshot) directly from the vSphere Client.

VMware Site Recovery Manager (SRM) leverages the SANsymphony-V SRA to replicate virtual machines and associated virtual disks between remote locations, making it possible to realize fully automated site recovery and cross-site migrations.

From a performance standpoint, SANsymphony-V responds to vStorage APIs for Array Integration (VAAI) to offload hosts and SAN from low-level storage operations best performed externally in the storage virtualization layer.

Extensive native integration with vCenter is also available directly from the SANsymphony-V management console. It communicates directly with vCenter Servers to automatically register vSphere hosts, clusters, virtual machines (VMs), and datastores in the host pane. From there you can serve and unserve virtual disks, create VMFS datastores and set the desired path selection policies (round robin, most recently used, etc.). The SANsymphony-V GUI automatically updates when new hosts are added to vCenter.

Storage services for clouds

Infrastructure as a Service (IaaS)

- Resource abstraction and control
- Provisioning and configuration
- Portability and interoperability

DataCore delivers the comprehensive set of storage virtualization and centralized storage management features necessary to run contemporary data centers large or small. These capabilities become even more crucial when building robust, scalable cloud storage infrastructures.

In private, hybrid and public cloud scenarios, SANsymphony-V forms the core IaaS functions necessary to abstract, configure, provision, protect and control storage resources with the least effort, the highest availability and best performance.

The independent software layer readily adapts to the frequent forces of change; among them, expansion, obsolescence, hardware refresh, software upgrades, facility moves, and policy updates.

Its rich suite of functions may be invoked through PowerShell scripts to automate actions in response to self-service requests. Taken together with automatic self-tuning mechanisms embedded in the software, DataCore measurably reduces manual intervention and the possibilities for human errors.

Compatibility

Storage manufacturers supported

All of the popular disk manufacturers are supported. These include:

- Astute
- Dell
- Dot Hill
- EMC
- Fujitsu
- Fusion-IO
- Hitachi Data Systems (HDS)
- HP
- IBM
- LSI
- NetApp
- Nexsan
- Oracle (Sun)
- Promise
- Seagate
- STEC
- Texas Memory Systems
- Violin Memory
- Whiptail
- X-IO

Host operating systems supported

- Microsoft Windows Server 2012, 2008 R2, 2003 and 2000
- Microsoft Windows 8, 7 and XP
- UnixHP-UX
- IBM AIX
- Sun Solaris
- RedHat Linux
- SUSE Linux

Server hypervisors supported

- VMware ESX, vSphere
- Microsoft Hyper-V
- Citrix XenServer

Systems management

- Microsoft System Center
- VMware vCenter
- Hitachi IT Operations Analyzer

Disk packaging supported

- Internal disk drives
- External JBODs
- External storage systems

Disk interfaces supported

- Direct-attached and SAN-based connections
- SAS
- SATA
- iSCSI
- Fibre Channel
- Fibre Channel over Ethernet (FCoE) via CNA switches
- SCSI
- IDE

Media supported

- Flash Memory Cards and Solid State Disks (SSDs)
- Standard Magnetic Rotating Disk Drives (HDDs)

(Please see www.datacore.com for more recent updates to the list of supported environments)

Network compatibility

- Synchronous mirroring between nodes over iSCSI and Fibre Channel connections
- Asynchronous remote replication over IP LANs, MANs and WANs
- Inter-node management interface over IP LAN
- Remote console access using Remote Desktop Protocol (RDP) and other standard remote desktop access protocols supported on Windows Server

Maximize value from your storage investments

As you can see, the SANsymphony-V storage hypervisor makes storage largely interchangeable. It reconciles disparities between different models of hardware devices, enabling you to choose across the full range of purpose-built products from your favorite manufacturers and later add attractive new technologies, unconcerned with backward compatibility. Furthermore, the skills you acquire to manage today's gear are the same ones you'll use to administer next year's lineup. This novel, yet transparent virtualization software brings unmatched flexibility and cost reductions in system design, procurement and ongoing operations. In so many ways, DataCore maximizes the performance and value you derive from your strategic and tactical storage investments.

1113

For additional information, please visit **www.datacore.com** or email **info@datacore.com**