New Features in SANsymphony™-V10 PSP1 Software-defined Storage Platform

Updated: November 7, 2014

Contents

Note: Some features are available as separately-priced options. Please consult with your DataCore-authorized solution provider for more details on how to put the new capabilities to work for you.

Introduction

The SANsymphony™-V10 PSP1 release brings major enhancements in the areas of performance / scalability, private clouds and hybrid clouds. The most significant new capabilities are discussed below.

Performance and Scalability Enhancements

Random Write Acceleration for Transactional Workloads (e.g., OLTP databases)

Greater than 30X performance acceleration for workloads characterized by many random writes, such as frequently updated databases, ERP and OLTP systems, are now possible by selecting the new sequential storage option for virtual disks. The remarkable improvements in writes (see graphs below) are most pronounced when using spinning disks (~33X faster), but also result in up to 3.6x faster writes from solid state disks (SSDs). The optimizations effectively overcome undesirable delays incurred when updating widely dispersed data blocks across these storage media.

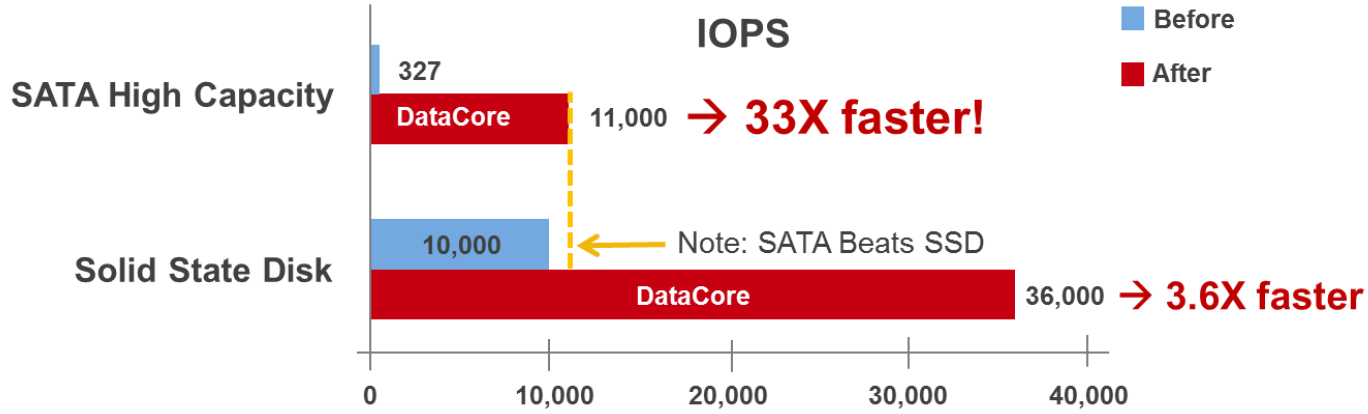

Specifically, random writes are stored in sequential order in new locations to sidestep the slowdowns resulting from:

- Physical head movement and rotational delays across the surface of spinning disks
- Write amplification in SSDs
- RAID-5 reads to recalculate parity data, especially with small writes

For random writes to RAID-5 pools, performance boosts as high as 46X have been observed.

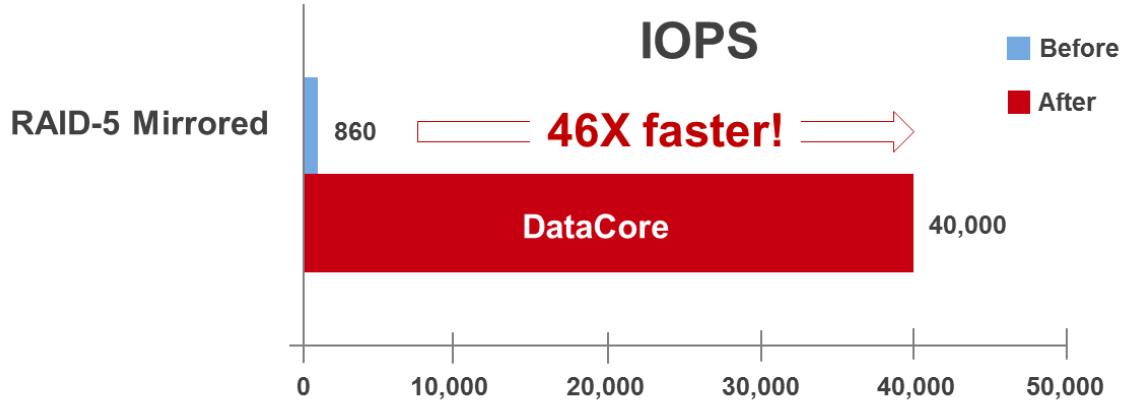

Considerations

The actual performance benefits you experience for any one application will vary depending on the percentage of random writes that make up that application's I/O profile and the types of storage devices you employ. There are also some tradeoffs discussed below. For these reasons, the optional feature can be enabled for individual virtual disks, allowing you to select when to apply the optimization.

It's worth noting that Random Write Acceleration is not expected to improve the performance of workloads that are already largely sequential, although it will still work. Also note that optimizing the I/O pattern to the back-end devices doesn't change the characteristics of the underlying storage. Any limitations in latency, throughput or IOPS will still apply.

Tradeoffs

Random write acceleration is designed maximize performance and therefore uses more backend storage than the amount of active data in the virtual disk. This extra space reduces the copying required and therefore yields greater performance by mitigating the non-sequential write penalty.

Double the Scalability

The maximum number of SANsymphony-V nodes in a server group grows from 32 to 64 with PSP1. In large VMware and Microsoft clusters, you can install a copy of SANsymphony-V10 on each application server to accelerate application response and achieve continuous availability. Doubling the possible number of nodes in a DataCore™ Virtual SAN is particularly valuable for latency-sensitive applications distributed over large-scale clusters with up to 64 nodes, as well as expanded virtual desktop (VDI) deployments. Total aggregate capacity across the 64 nodes rises to 64 Petabytes, and their combined performance grows to 100 Million IOPS.

The higher number of nodes per group also allows organizations to deploy SANsymphony-V in a broader set of distributed branch offices, and centrally manage them as a group from the same console. If instead, one chooses to treat each site as a separate group, PSP1 also introduces the ability to redirect the management console from one site to the other.

Greater scale is important when running SANsymphony-V on centralized SANs for very large environments. You can expand your infrastructure to fulfill larger I/O requirements by harnessing the power and/or port connectivity of up to 64 nodes. You can grow into these larger configurations over time as the need to scale out arises; yet plan for them at the outset. Scaling out the configuration across more nodes enables you to better distribute the workloads and segment the physical storage capacity. Internode mirrors can be stretched over metro-wide distances up to 100 kilometers (~60 miles) to improve overall availability and business continuity. Scaling out to more nodes also enhances the overall resiliency of the infrastructure using an N+1 redundant grid.

Essential Enhancements for Private Clouds and Hosted Cloud Services

Quality of Service (QoS) Controls

New in SANsymphony-V 10 PSP1 are QoS settings to ensure high-priority workloads competing for access to storage can meet their service level agreements (SLAs) with predictable I/O performance. QoS Controls regulate the resources consumed by workloads of lower priority.

Without QoS Controls, I/O traffic generated by less important applications could monopolize I/O ports and bandwidth, adversely affecting the response and throughput experienced by more critical applications. To minimize contention in multi-tenant environments, the data transfer rate and IOPS for less important applications are capped to limits set by the system administrator.

QoS Controls enable IT organizations to efficiently manage their shared storage infrastructure using a private cloud model. Storage resources can be logically segregated, tracked and regulated on a departmental basis as illustrated below.

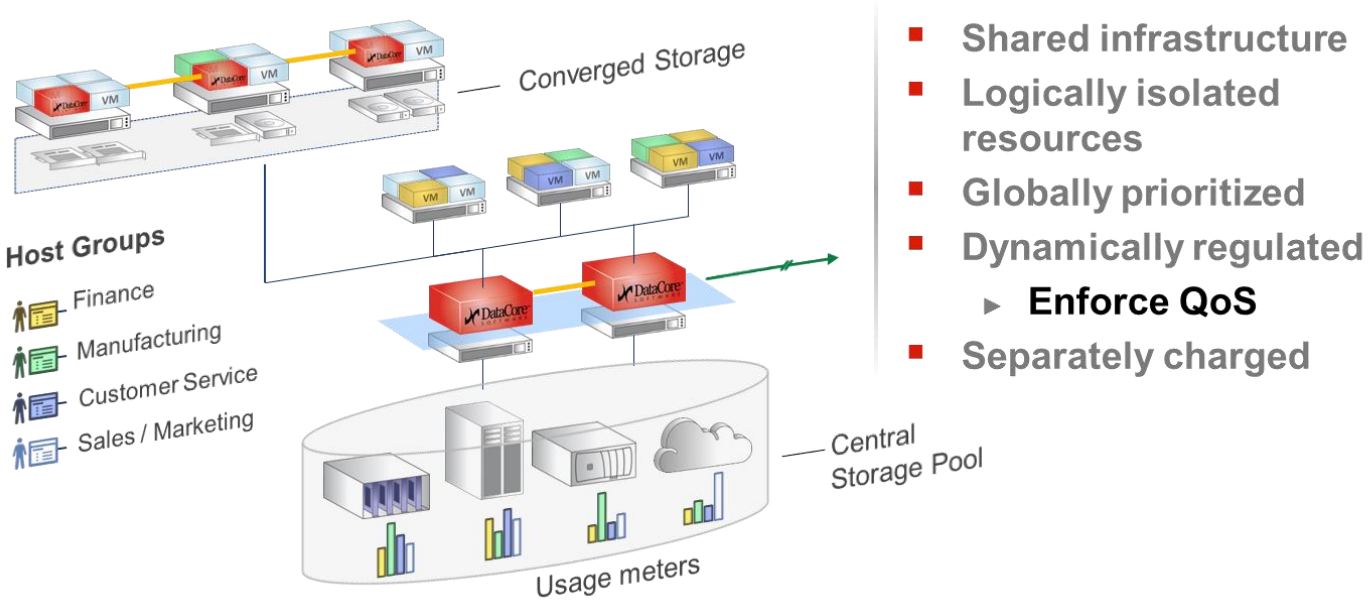

QoS parameters can be set for individual hosts or groups of hosts. For streaming applications which burst data, it's best to regulate the data transfer rate (MB/sec) to minimize their impact. For transaction-oriented applications (OLTP), limiting the IOPS makes most sense. Both parameters may be used simultaneously.

The effects are illustrated below. First, the behavior of competing workloads is shown in an unregulated environment. Next, the effects of limiting IOPS are highlighted. In the last figure, the system throttles throughput for a specified workload.

Without QoS: Unregulated IOs

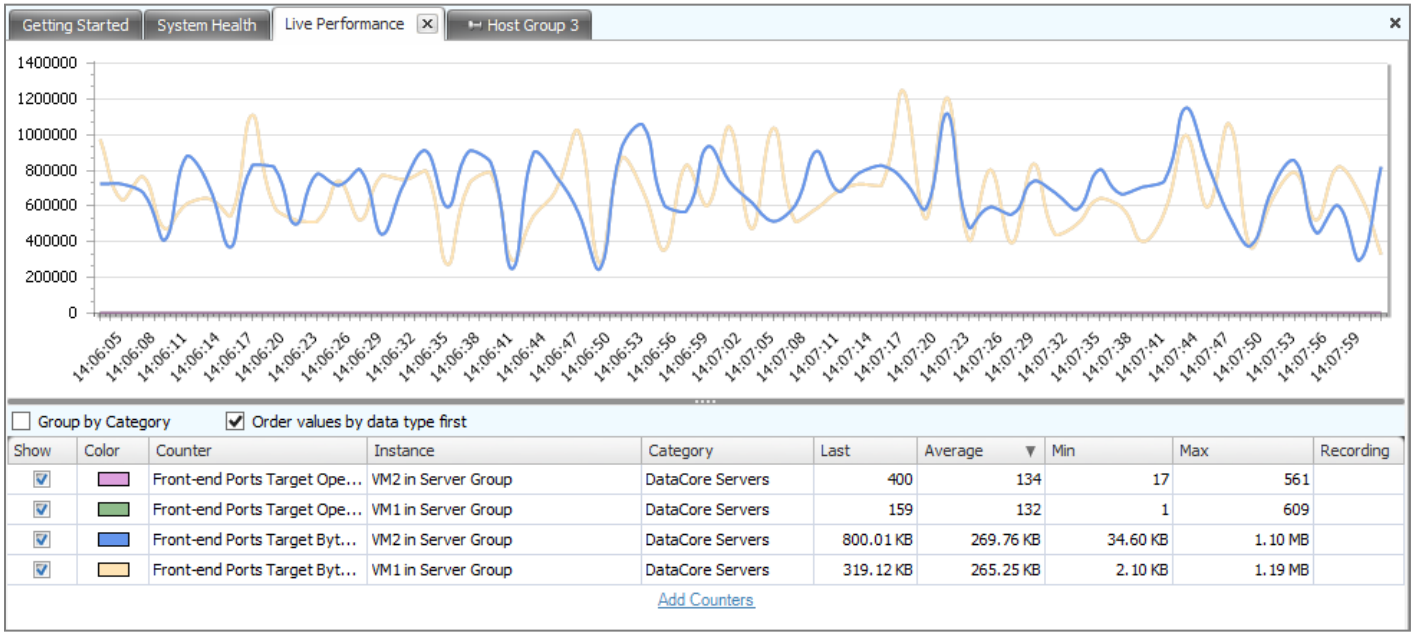

With QoS: Set Max IOPS

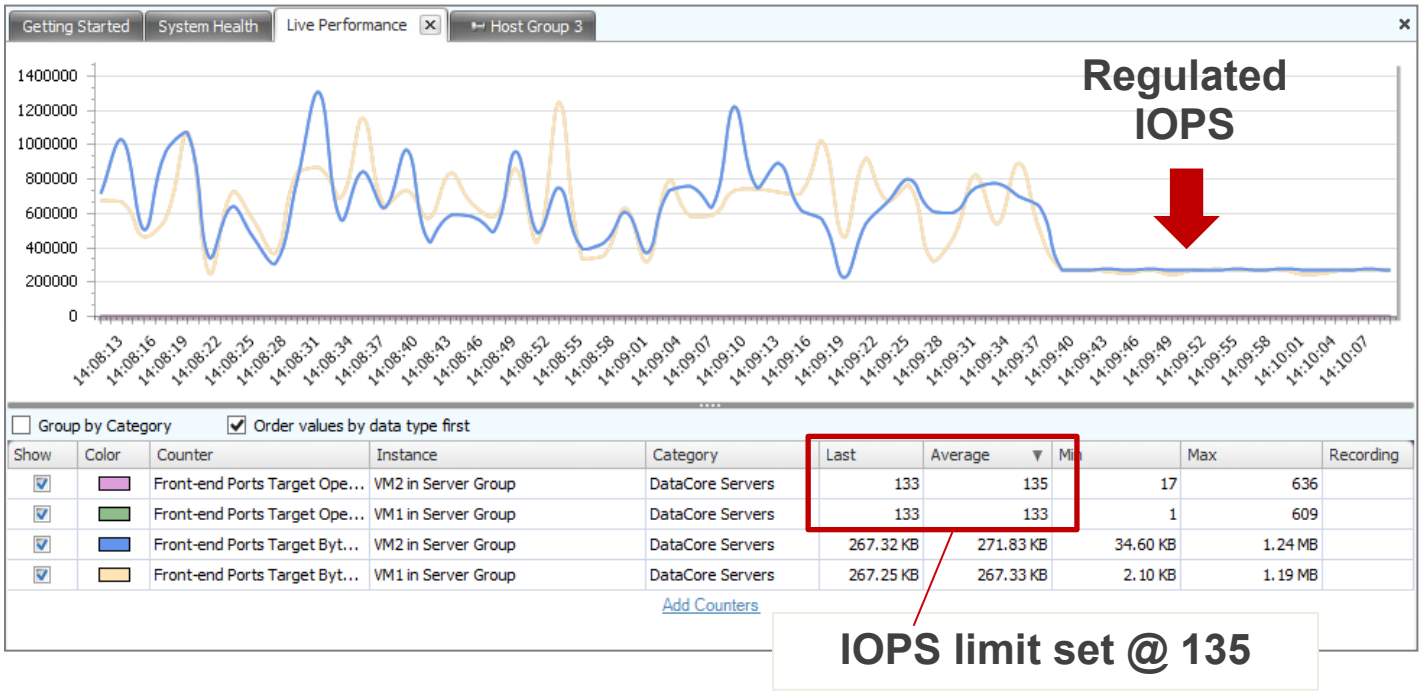

With QoS: Set Max Throughput

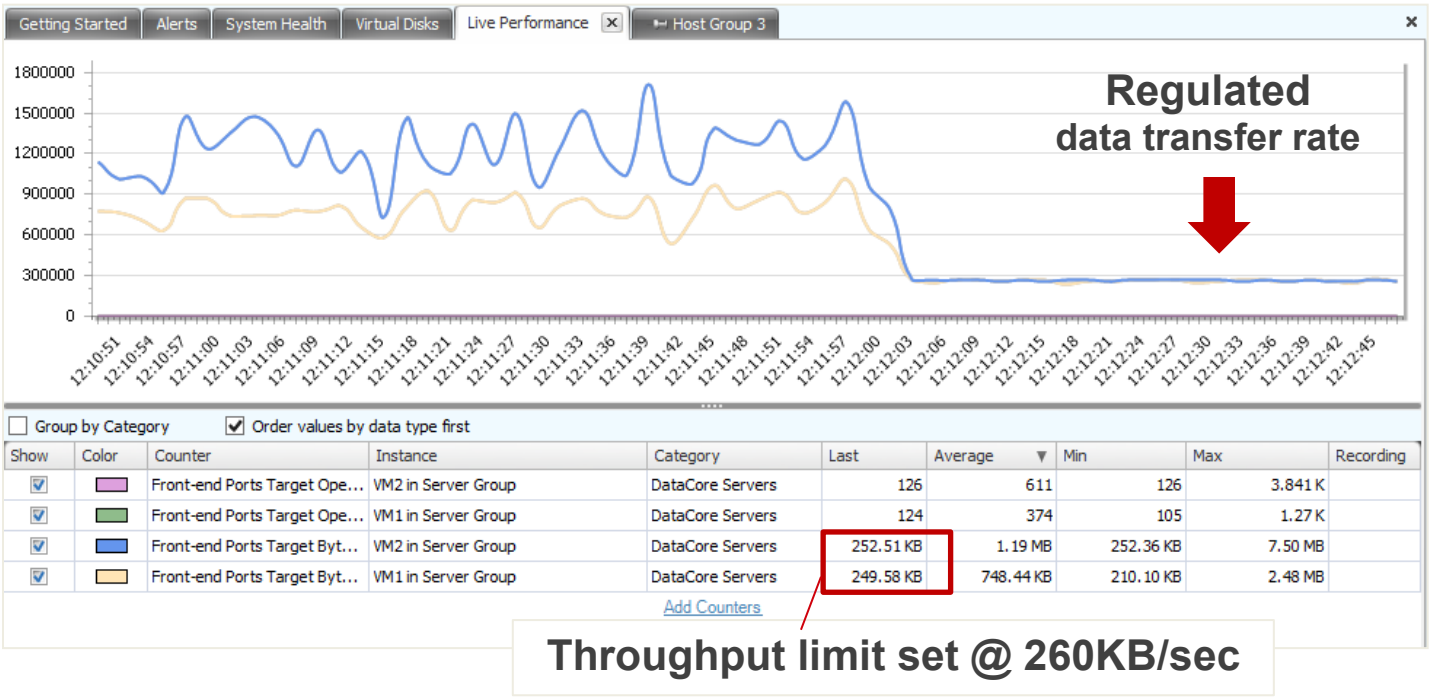

Chargeback and Reporting; Keeping Track of Who Uses What

To help allocate costs, especially in private cloud and hosted cloud services, SANsymphony-V generates reports quantifying the storage resources consumed by specific hosts or groups of hosts.

The reports tally several parameters including:

- Total Bytes Transferred
- Total Bytes Written
- Total Bytes Read
- Total IO Operations
- Total Read operations
- Total Write operations
- Total Bytes Provisioned

Usage information can be exported in several tabular formats to equitably charge consumers for the services they've received over a specified period.

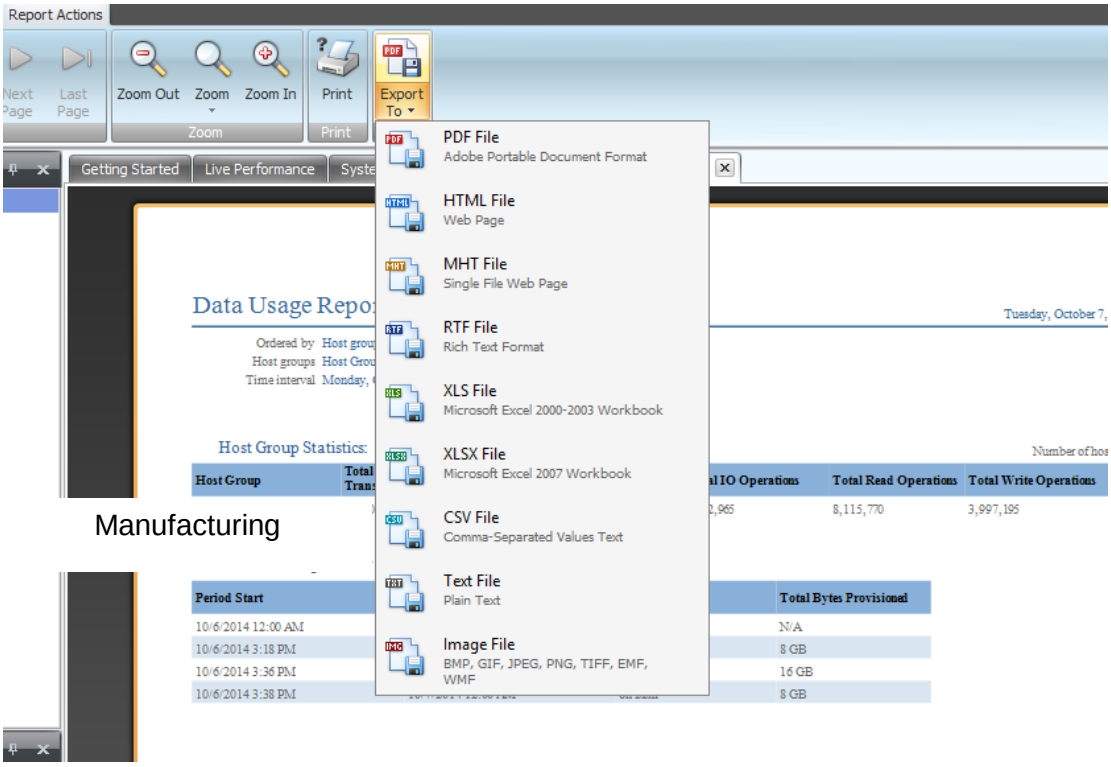

Sustain Redundancy and Performance during Major Maintenance Operations, Upgrades and Expansion

Another important new capability in PSP1 comes into play when taking a SANsymphony-V node out of service, whether temporarily or permanently. Control of that node's shareable storage resources can be transferred to a standby node without disrupting applications. The standby node takes the place of the one undergoing maintenance to maintain two active mirrored copies of data for high availability. In addition, the standby node ensures that the overall throughput and responsiveness of the system remains intact.

In metro or stretch clusters, transferring control to the standby node has an added advantage. Hosts at the affected site no longer need to redirect their requests over the slower metro area network (MAN) to the mirrored copy on the remote node. Instead, they continue to access local storage resources at the fastest possible speed without traversing the MAN.

Some Conditions Apply

The standby node can only take control of storage devices which can be accessed by both nodes simultaneously. Single ported storage devices, whether internal or direct-attached externally, are excluded. Shareable devices would include multi-ported storage devices as well as SAN-accessible arrays. Essentially, the storage devices connected to the node that will undergo maintenance must also be connected to the standby node. Then SANsymphony-V non-disruptively transfers controls of the virtual disks dependent on those storage devices to the standby node.

Offloading Nodes by Redistributing Workloads

The capability to introduce a standby node into an active SANsymphony-V group may be used to offload some of the responsibilities of an overtaxed node. For example, some of the virtual disks and corresponding storage devices from a node near its capacity/performance limits can be redistributed to a 3rd node; essentially increasing the power of the storage infrastructure by 50% or more. This should enhance the responsiveness of the overall system and set the stage to absorb additional workloads.

Additional Centralized Management Capabilities

Integration with Windows Perfmon

In addition to the native real-time charts and performance metrics accessible from the SANsymphony-V console, PSP1 can export those metrics to the Windows Server Performance Monitor (Perfmon) utility executing locally on the DataCore node. In this way, system administrators can get a complete picture of the performance aspects for the Windows Server where SANsymphony-V is running. One can choose which counters should be visible in Perfmon. They can also use Perfmon-aware applications and scripts to incorporate SANsymphony-V performance metrics in the broader understanding of their infrastructure.

Tag Known-Good Restore Points for CDP

You can now undo the effects of malware, accidental user deletions and logic errors by restoring to known-good points in time within the Continuous Data Protection (CDP) log to achieve the best recovery point objective (RPO).

Similar to snapshot requests, one can generate a CDP Rollback Marker by scripting a call to a PowerShell cmdlet when an application has been quiesced and the caches have been flushed to storage. Several of these markers may be present throughout the 14-day rolling log. When rolling back a virtual disk image, one simply picks the marker signaling a restore point just before the incident occurred.

Switch between Multiple Server Groups from the Same SANsymphony-V Console

For organizations running two or more SANsymphony-V server groups, PSP1 offers the ability to redirect the management console from one site to the other using a pull down list. This capability is particularly valuable when managing multiple remote sites from a central location. The software retains server address and login credentials from each initial session to expedite the switch between groups being controlled and monitored.

Automated SANsymphony-V Upgrades

Upgrading to a new release of SANsymphony-V is made simple by a friendly new wizard that automates much of the process. This reduces the chance for errors. The wizard performs the updates across nodes in a group in a systematic way. Of course, one can choose to perform the update manually.

Hybrid Cloud Storage Infrastructure for Cost Savings and Simplicity

The explosive growth of data and increasing requirements for business continuity / disaster recovery make it impractical to keep all data within the confines of privately-maintained data centers. Through new integration with the Microsoft Azure cloud platform, SANsymphony-V enables a hybrid cloud storage infrastructure where active data remains on premises, rapidly accessible, while cold, inactive data resides in the cloud. The same technology enables backup copies to be stored in Azure for disaster recovery purposes and long term archive, securely accessible across the world.

While SANsymphony-V can be configured to utilize many Azure Cloud services, a more efficient and simple solution that is certified DataCore Ready is now available. Key to the solution is the Microsoft StorSimple Hybrid Cloud Storage Array which provides local iSCSI storage as well as a gateway to geo-redundant Azure cloud storage.

SANsymphony-V can auto-tier on-premise storage to StorSimple Arrays as well as pool these arrays just like other iSCSI storage. As new data arrives and older data becomes less frequently accessed, the StorSimple arrays can then migrate the inactive data to Azure over an Internet connection. The process regularly vacates space to make more room for the current working set of active data. The array deduplicates and compresses the cold data before moving it to the cloud, further shrinking the capacity required to store it.

From a disaster recovery perspective, the StorSimple Arrays can also create backup copies (snapshots) in Azure of the current working set. In case of emergency, the backup copies can be retrieved from other locations equipped with StorSimple Arrays, or by Azure-hosted applications. Apps running in Azure access the backup copies without the need for a separate StorSimple Array through a cloud-resident StorSimple Virtual Appliance. The appliance behaves like an array.

Summary

DataCore customers upgrading to SANsymphony-V10 PSP1 can enjoy some significant enhancements that benefit not only conventional data center operations, but those IT organizations seeking the benefits of private and hybrid cloud infrastructures.