VIRTUALLY PERFECT PROTECTION

SIMPANA® software

Commvault® Simpana® software protects hundreds of virtual machines in minutes. This paper describes how Simpana software achieves this and why it is crucial for businesses virtualizing their data centers.

Simpana software features:

- Multiple recovery operations, including granular recovery
- Application-integrated protection
- 360-degree actionable reporting
- Flexible, business-friendly licensing

commvault®
solving forward®

Virtualization is spreading like wildfire— and turning up the heat on the need for a better way to manage and protect virtual data. Companies are finding that legacy-software-and-hardware-based data protection systems, and even some of the newer appliance-based systems, are simply not capable of meeting their needs.

Highly consolidated, rapidly changing virtualized environments require data management systems to efficiently and automatically manage data, and seamlessly connect to physical storage.

Simpana® software, with its single unified codebase and platform is uniquely equipped to bring together the physical and virtual worlds of computing. It can back up hundred of virtual machines in minutes and scale to protect thousands of virtual machines across the enterprise. It can also help you create an available, reliable, robust cloud and protect critical applications with SnapProtect™ snapshot technology.

In short, it's the "virtually" perfect solution you need to make the most of your virtualization investment.

Modern trends in IT are changing how data centers are designed and managed. Data centers continue to evolve from an environment based on physical servers and storage to one based on virtualization platforms that abstract out the server and storage hardware from the user applications. Data management solutions need to evolve and adapt to these shifts in the data center as well.

Data Centers are Changing

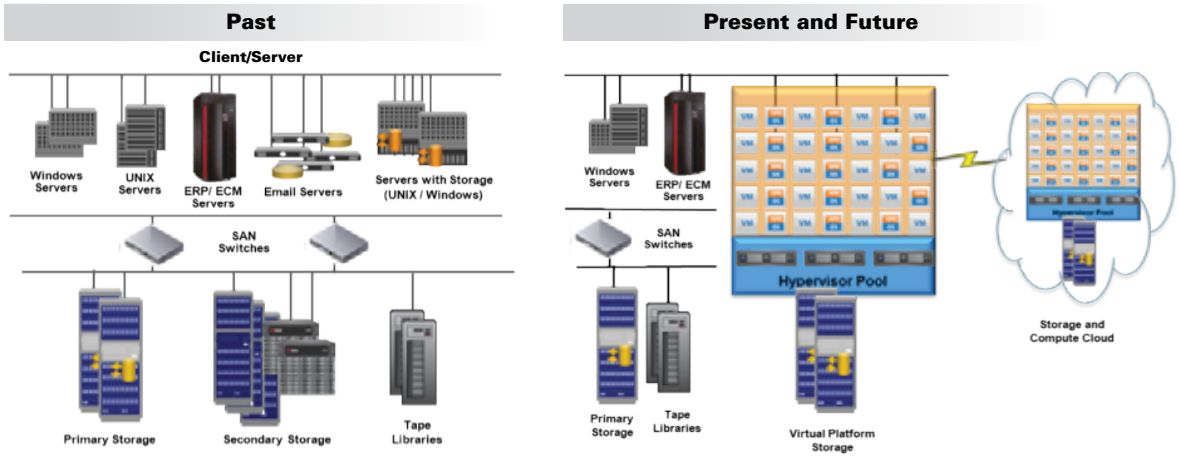

Traditional data centers revolved around physical servers and storage dedicated to these servers. The amount of data per server was limited and typical servers had plenty of unused processing cycles available at any given time. With the growing influence of virtualization, many of these physical servers are converted to virtual servers. The newest version of the VMware platform provides unprecedented scale and server consolidation capabilities. Ever increasing virtual machine density per physical server and higher resource availability per virtual machine allow almost all servers in a data center to be virtualized. However, some mission critical servers will continue to remain on dedicated physical servers and storage. The next step into this evolution is cloud computing. Over time, there will be an increasing move to shift computing and storage resources into public clouds. At the same time, the data center itself will resemble a private cloud where the physical data server complexity is abstracted out to users and application owners. This model provides a lot more **agility and automation** to how data centers are managed.

Data Management in the Virtualized Data Center

Data management solutions need to evolve to keep pace with these trends in the data center. Traditional data management (backup/archive/reporting) solutions deploy resource intensive agents in the physical server which physically collect and move data from production storage to a backend disk or tape target. This worked well for limited data sets within the boundaries of pre-defined operating windows, where no major activity is assumed on the production systems. Over time, not only has the amount of data to be moved grown dramatically, but operating windows have collapsed. The transition to virtual environments has only accelerated this trend, increasing the need to move more data in less time. While deduplication, especially source-side deduplication, can be an effective way to move more data in a shorter window, with less storage footprint. However, the impact of deduplication is increasingly limited as more and more servers are consolidated, resulting in ever fewer physical resources available to move those larger data sets as needed.

With the shift to virtualized data centers, there is a need to rethink the traditional data protection techniques. The legacy method of using brute processing power to move data from one place to another will no longer suffice. Also, the traditional practice of restoring to the previous night's backup is simply no longer sufficient to meet many Service Level Agreements (SLAs). With round the clock operations, data protection, data recovery, data management and reporting operations must have minimal impact on the production systems and require minimal administrative intervention. The winning solution needs to address existing issues in the traditional data centers, but more importantly, ease the transition to a virtualized and eventually a cloud-based data center.

CommVault Simpana Software: New Approach to Data Management

CommVault Simpana software is a revolutionary data management solution that not only addresses the challenges arising out of limitations in legacy data center environment, but more importantly, accelerates the shift into virtualized and cloud-enabled data centers. With Simpana software, businesses can start to realize tangible benefits on day one of deployment as they transition from traditional environments to the modern data center. More critically, using techniques that access data once, but reused for multiple operations, Simpana software ensures that businesses can start reaping the advantages of the modern data center immediately, bypassing many of the pitfalls that arise from trying to forcibly fit legacy techniques in the modern data center.

Key features of CommVault Simpana software:

- Protect hundreds of virtual servers in minutes with no impact on physical production servers.
- Embedded source-side deduplication for rapid creation of secondary DR copies.
- Create 100% application consistent protection copies.
- Granular information mining of application data for granular restores, content search and eDiscovery.
- 360-degree view using SRM that reports on the physical and virtual infrastructure contents for more informed management of the platform.

Simpana software enables customers to rapidly modernize their data centers into private clouds to take full advantage of the advances in virtualization technology while continuing to meet all the data management and data retention needs.

Protect Hundreds of Virtual Machines in Minutes

The Challenge

Every virtual machine is essentially a set of large VMDK files. These large files are stored on LUNs known as datastores. Datastores can be configured on iSCSI or Fiber Channel block storage volumes or on NFS volumes. Traditional data protection techniques for VMware such as vStorage API for Data Protection (VADP), or VMWare Consolidated Backup (VCB) rely on an external agent to protect VMDK files associated with virtual servers. Typical steps are as follows:

- Queisce virtual servers to get a consistent set of VM image files.
- Use the VADP enabled agent to read the VM image files from the datastores.
- Copy the image files to a backup disk target.
- Release the Virtual Servers for normal operations.

While VADP brings much efficiency to this process, it is still a streaming method that moves the image files from the datastore to backup disk for protection. For small to moderate sized environments with large enough backup windows, this method may suffice. For larger environments with ever shrinking backup windows, there is simply not enough time or bandwidth to move all the VM data. Even if the infrastructure is available to copy all this data, it places a tremendous burden on the datastores as the data is read. Moreover, with an estimated average of 40% data growth rate every year, this method is simply not sustainable without frequent and expensive infrastructure upgrades.

Solution: SnapProtect for Virtual Server Agent

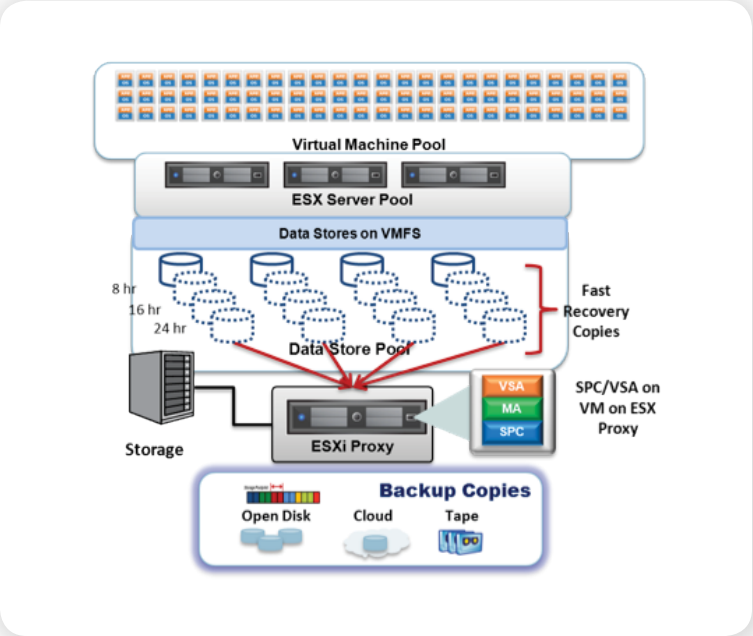

SnapProtect for Virtual Server Agent (VSA) integrates with the snapshot engines built inside the storage devices to create rapid recovery copies of datastores containing virtual machines. These snapshots represent fast VM recovery copies.

The SnapProtect for VSA and the MediaAgent module are configured on a Windows system. It is recommended that this system be a virtual machine. While the VSA virtual machine can run on any ESX server, it is recommended that it be hosted on an ESX server that is designated for backup, i.e. the designated server should not have too much additional workload. This backup ESX proxy can also run on the ESXi version of the hypervisor.

A SnapProtect job follows the same sequence as a regular backup job; however, instead of copying data blocks, it executes a rapid snapshot. The sequence is as follows:

- Discover new VMs based on pre-defined criteria.
- Quiesce VMs to ensure a consistent set of image files.
- Determine datastores associated with the VMs.

- Execute hardware snapshot using storage system's APIs, this takes a few seconds.
- Unqueisce the VMs to resume normal operations.
- Index the snapshots and the VM list inside the snapshot.

Since the SnapProtect job creates a rapid snapshot, the protection operation takes only a few minutes, and the VMs themselves are quiesced for a very short time. The short duration allows the creation of multiple recovery copies per day with minimal impact on virtual systems providing better recoverability and enabling more aggressive recovery SLAS.

In a typical environment, one would create a snapshot-based recovery copy every 4-8 hours and selectively copy contents of one of the recovery copies to a backup to disk or tape target as described below.

Business Benefit:

- Extremely fast, reliable protection for data in virtual machines, with minimal production impact. Allows more and more critical systems to be virtualized for more rapid ROI from the shift to the virtual platform.
- VMware is a strategic choice, whereas the hardware is a tactical decision. SnapProtect for VSA is storage platform independent, providing the same set of data management features across a wide variety of storage systems, allowing flexibility in the choice of hardware platform, without compromising on capabilities.
- Multiple recovery points per days allows users to recover from a point closer to the point of failure, allowing recovery to a more recent point in time than simply last night's backup.
- Evolve to a cloud strategy while the Simpana software data management solution transparently bridges the gap between the legacy and modern data centers.

Create Rapid DR and Long Term Retention Copies

The Challenge

SnapProtect for VSA eliminates the copy process completely by leveraging hardware snapshots for creating fast recovery copies. However, the number of snapshots that can be retained is limited for most storage systems. Moreover, the more snapshots retained, more reserve disk space is consumed on the production storage. To retain and allow recovery of data for more than a few days, it is necessary to create a second copy of the contents outside the production storage and on cheaper disk or tape targets.

Additionally, snapshots are usually dependent on the source LUN. It is rare in the modern storage system for a LUN to fail. However, LUNs do fail and when that happens, all snapshots associated with that LUN are lost. Hence, it is critical to create a second copy of data contained within the snapshot.

The challenge is in creating secondary copies without diverting scarce resources from the production VMs or production ESX hosts and without consuming excessive I/O cycles on the production stores.

Solution: Snapshot Backup Copy with Embedded Deduplication

SnapProtect for VSA includes the ability to copy the contents of selective snapshot based recovery copy to a secondary disk or tape target for DR and for data retention beyond a few days. Users can create multiple recovery copies per day. One of these snapshots can then be copied to a secondary disk or tape target. This allows users to store more recovery points than the storage system would allow.

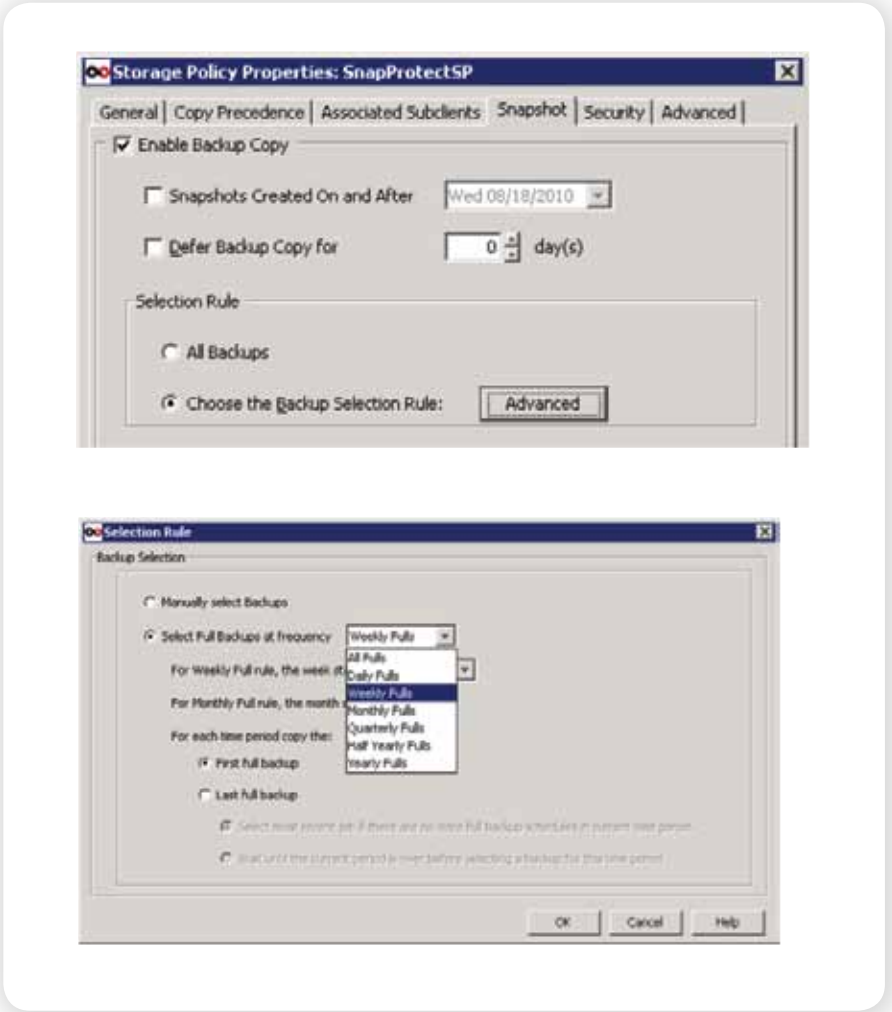

When creating the Snapshot backup copy, VSA mounts the selective snapshots as temporary datastores on a designated ESX server. VSA then leverages VADP against the VM images inside these temporary datastores to read data blocks and write them to secondary storage device. However, since these are distinct datastores within VMFS and the backup ESX server has exclusive access to them, there is no IO contention for these datastores. As a result, we do not need to lock out production datastores from the production ESX servers. All issues arising from SCSI locks required during regular VADP backups are completely eliminated. Moreover, since the VM images are already consistent, there is no need to quiesce the production VMs again. With this approach, VSA completely eliminates the impact of VADP copy on the production ESX servers and production VMs.

VADP does impose an IO burden on the storage LUNs. To minimize the workload on a single LUN and to distribute the workload of copying across all data stores, VSA has the option to create backup policies by data store and limit the number of streams per store. This allows many VM images to be copied simultaneously but limits the impact on any individual LUN to one or two streams, while involving all LUNs in the copy process, leading to much faster data transfer.

Since the backup copy job has no impact on production servers, this could be run any time thus providing a larger operating window in which to create the second copy. Conversely, since you have more time to move data, you do not need to provision many ESX servers in backup role. For example with regular VADP backup, assuming a throughput of 1TB/hr per backup ESX, it would require 2 backup servers to backup 15-20 TB of data in a given 8 hour backup window. Because of the additional time available with VSA you can use a single backup ESX server for that role.

The backup role ESX server should be sized for high IO, not for high or dense computing. In other words, while the other ESX servers may require a large number of CPU cores with hundreds of GB of RAM to host as many VMs as possible, the backup ESX server need not have more than 8 cores and 24-32 GB of RAM. Also worth noting is that having these dedicated backup ESX servers offloads the burden of backup from the larger ESX servers. If the production ESX servers also need to bear the burden of data movement and iSCSCI locking for backups, they will probably host fewer VMs than they are capable of. The backup ESX servers eliminate that burden, allowing the other expensive ESX servers to have higher VM density for more optimal utilization of server hardware and the processing and IO capabilities to serve business applications.

When creating backup copies on disk target, VSA leverages embedded data deduplication to reduce the amount of data that is transferred over the network and written to disk by 90-95%. As a result, copies can be created almost 50% faster as compared with regular VADP without deduplication.

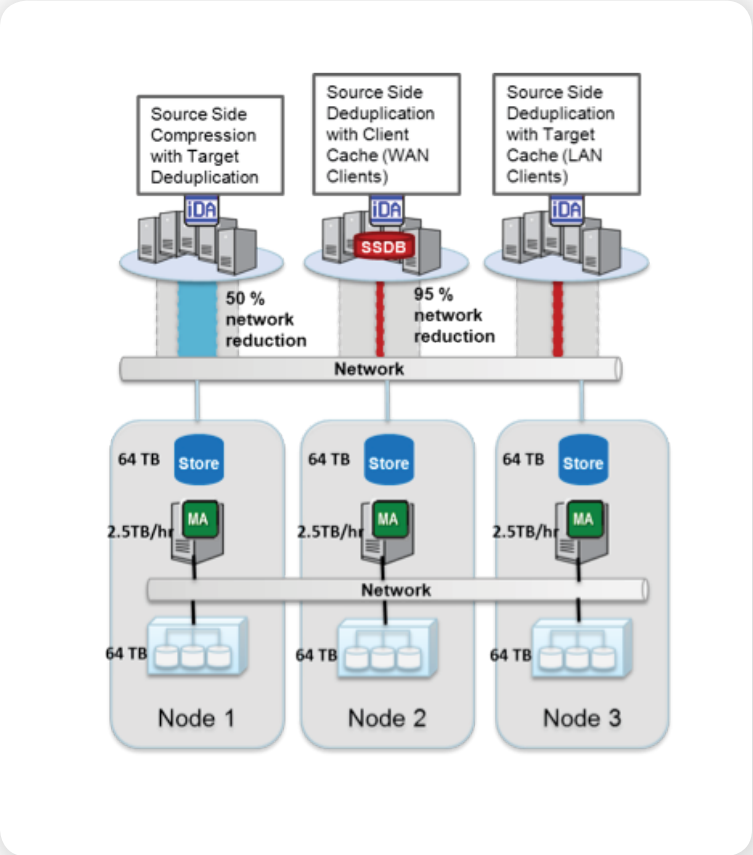

The Simpana software deduplication feature offers a highly scalable and global, node based, deduplication option that can scale to handle datasets even in the largest of enterprises. At the same time, it is easy to incrementally add more capacity and increase throughput simply by adding disk incrementally or adding additional MediaAgent nodes. There is no need to replace the older dedupe store with a completely new one simply because you have run out of space on the old store. This embedded deduplication, including source-side deduplication, is available for all types of backups, including virtual machine backup. More critically, the source-side deduplication has no limitation on data size on a single client and can perform at high speed even with multi-terabyte data sets on individual clients.

The example in the image illustrates a 3 node deduplication configuration that supports up to 196TB[1] of disk capacity and a throughput of 7.5 to 9 TB per hour. Assuming a conservative 10:1[2] deduplication ratio, that represents a logical capacity of almost 2 Petabytes, sufficient to satisfy the needs of very large environments.

[1] Assumes 128K deduplication block size, set on the Storage Policy or Global Dedupe Policy.

[2] 10:1 deduplication ratio typically requires around 60 day data retention on dedupe disk with weekly full and daily incremental backup schedule. More frequent full will lead to higher deduplication ratio sooner.

Business Benefit:

- Ensure rapid creation of DR copies to recover from any type of data or hardware loss.
- Scalable deduplication capability that allows store to increase both throughput and capacity incrementally as the business needs grow, without having to throw out older investments.
- Retain data longer on cheaper secondary disk in deduplicated form allowing access to a lot more data than can be retained on snapshots.
- Create long-term retention copy by moving data transparently to tape in dedupe form.
- Divert the backup workload from production-expensive ESX servers and virtual machines to cheaper backup ESX servers, allowing more optimal and appropriate use of the expensive, multi-core, high memory, high computing servers, realizing a lot more value for money spent.

Automatic VM Discovery and Automated Data Protection

The Challenge

Given the ease of creating virtual machines, new ones are created all the time. This is almost always done without the knowledge of the storage or backup admins who are tasked with ensuring that all data in the environment is protected. As a result, these admins have to spend an inordinate amount of time every day, sometimes as much as 75% of their day, hunting down new virtual machines and their owners, determining their purpose and deducing their data protection and retention policy. These manual tasks lead to tremendous loss of productivity and wastage of time that could otherwise have been spent on projects that help drive business growth.

Solution: Automatic Discovery for VSA

The virtual server agent includes the ability to automatically discover new virtual machines based on pre-defined criteria and transparently add them to data protection policies. This ensures virtual machines are protected even though administrators have no knowledge about them. A wide variety of auto-discovery rules are available that allow administrators to fine tune discovery policies that best meet their needs. There is even a catch-all policy to ensure VMs that do not meet any of the pre-defined criteria are protected nevertheless.

With this option, administrators can set the discovery rules once and never have to bother spending a minute looking for new virtual machines to protect. This can save hours per person, time that can be utilized for more fruitful activities.

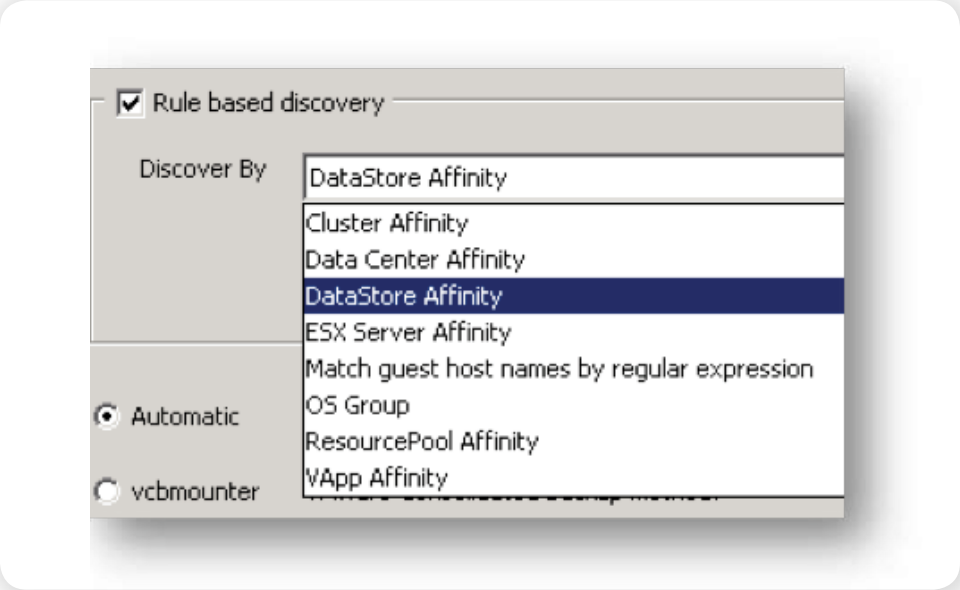

The most commonly used rule for SnapProtect is the Data Store Affinity rule that groups all VMs in the datastore in a single data protection policy, or subclient. This ensures that when a datastore LUN is snapped, all the virtual machines hosted in that data store are in a consistent state, eliminating the possibility of "dirty" files within the recovery copy. The Data Store Affinity rules also allows the backup copy process to distribute the workload of copying data blocks equally across all LUNs without excessively over-burdening any particular LUN.

Business Benefit:

- Increased IT staff productivity, freeing up more time for administrators to drive business critical projects that drive revenue.
- Effective workload distribution ensures best possible use of storage LUNs without overwhelming any single one, allowing production processes to run with minimal impact.

360-Degree View and Actionable Reporting

The Challenge

Storage and backup administrators and even VMware admins, often lack insight into the contents of virtual machines. As a result it is extremely challenging to make smart optimal decisions on the amount of and quality of storage resources that need to be made available for virtual machines. Traditional host-based Storage Resource Management tools require agents inside every server, in this case, virtual server, to collect that information which can be quite burdensome. Even if agents are not required in every machine, these tools need to run data collection operations on a frequent basis that has a significant direct impact on the VMs and the ESX servers. To avoid the disruption they cause, most users avoid using these tools and lose a critical technique in ensuring optimal storage allocation and utilization for these virtual servers. This often leads to a waste of critical storage resources.

Solution: Storage Resource Management for VSA

The Simpana software Storage Resource Management capability integrates with the Virtual Server Agent to provide detailed contextual and historical reports of the different elements in a virtual infrastructure. SRM includes detailed reports and dashboard views of physical resource consumption like CPU, memory, network and consumption per ESX server and per VM. It also provides historical reports along these dimensions.

However, what distinguishes SRM is the ability to provide detailed file level analytics of content inside virtual machines, without the need to install agents in each virtual server. The SRM capability within VSA does not require any data collection operation. Rather, it reuses the data collected by the VSA backups and generates reports based on information collected inside the index. Thus the SRM feature introduces no additional burden on VMs or production systems and presents detailed information of not just the virtual infrastructure, but also what is precisely contained inside each VM.

All the reports in the world are worth nothing if they do not contain some automation that takes decisive action based on information found in the reports. The Simpana software File System Archiver agent integrates with SRM reports to automatically archive off files indicated by the SRM report that have met pre-defined archiving criteria.

For example, an SRM report could return files inside virtual machines that are older than six months. When this report is fed to the File System Archiver inside the virtual machines, these old files will automatically be archived. Note that even though this requires an Archiver agent inside the VM, it is not as burdensome as backup agents for two reasons. First, Archiver is always incremental and it only moves files that have met the criteria since the last time the Archiver job ran, which will be a small collection. Second, more than the data transfer, it is the scanning of the file system that consumes a lot of VM resources. Since the SRM report hands over a ready list of files, the Archiver agent need not scan the file system, thus eliminating the biggest workload.

Business Benefit:

- Integrated 360-degree view of the physical and virtual environment.
- File level analytics allow storage administrators to make smart decisions on storage allocations to "important" VMs, ensuring the best storage resources are available to only those VMs that require high-end capabilities, ensuring better utilization of more expensive storage.
- Integrated and automated file level archiving of data inside virtual machines prevents rampant storage growth inside VMs, helping keep production storage costs under control, with very little administrator intervention.

Application-Integrated Backups and Information Mining

The Challenge

Applications running inside virtual machines often do not get the same type of protection as they would when they are running on physical servers. Specifically, tasks like log truncation and information mining of object data are not possible. As a result many application owners are still reluctant to virtualize application servers.

While some tools are available to perform granular object level restores from application files inside virtual machines, these usually require that all of the application data be available on disk in native format. However, most object restore requests are for objects that are quite old. For example, it is common for restore requests for Exchange e-mail messages that are 30 days or older. However, this will require application files from 30 days ago to be stored on disk. If the application files total 1TB, and you are running application backup, this implies you need 30 TBs of disk to keep application files around for 20 days. If older messages need to be recovered, more disk is needed.

Because of these excessive disk requirements, tools that allow granular object level recovery from applications are prohibitive and rarely used in practical environments.

Solution: Application-Integrated VSA

The SnapProtect agent for VSA provides application-integrated recovery copies that ensure that the recovery copy contains a point in time consistent image of the application running inside a virtual machine. In addition, SnapProtect for VSA allows the application to treat this recovery copy as an application-aware protection copy, allowing the application to perform post backup tasks such as log truncation.

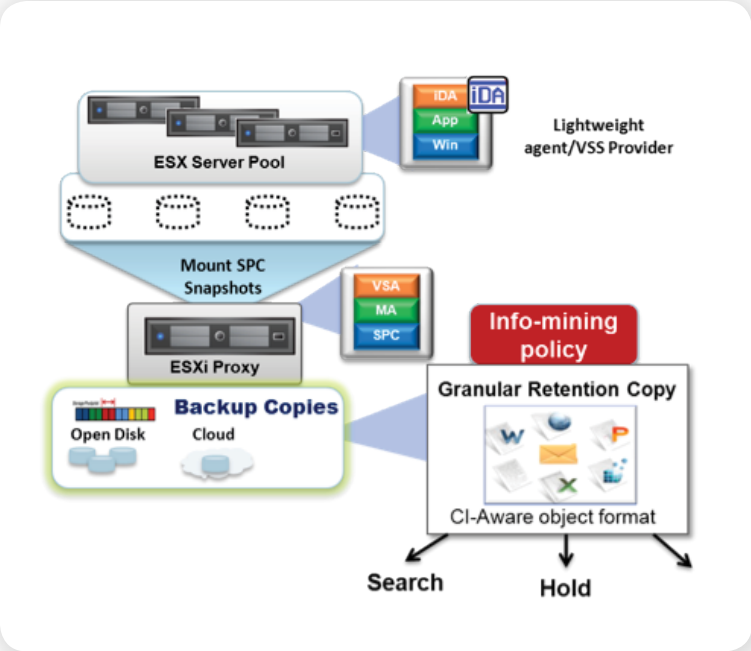

The image illustrates how SnapProtect for VSA integrates with applications running inside a virtual machine. It utilizes a lightweight application-aware agent inside the virtual machine hosting the application. Before the SnapProtect job queisces the VM, it interacts with the lightweight agent to prepare the application for backup. Once the application is ready and all other VMs in the subclient are quiesced, SPE job creates a hardware snapshot. Once the snapshot creation is complete, the SnapProtect job asks the agent in the app VM to truncate any logs. Once the logs are truncated, the application is released for normal operations and the VM is unqueisced. This whole operation lasts a few minutes.

Conventional wisdom dictates agents inside the virtual machines are bad. In fact, agents by themselves are not bad; what is bad is the amount of resources they consume, the amount of data they move and the cost of licensing and maintaining each one separately. The agent inside the VM does not actually move any data and consumes some resources only at the time of backup for a few minutes. Hence it adds very little burden to the production system. This agent can be remotely installed and updated without having to login manually in the VM. Most importantly, because the agent is not moving any data, it does not consume any license. It functions in restore only mode that does not cost a license. In addition, it acts as a restore target, so that it is possible to restore data directly inside the virtual machine, without copying it to a temporary location and then manually copying it back to the VM.

Simpana software also includes the option to run information mining operations against recovery copies created by VSA. The information mining policy extracts out selective object level information from application files and provides the option to store the objects individually in separate Storage Policies with different retention. This allows for the ability to perform granular object level restore, for example individual message level restores. More importantly, granular objects can be retained for long periods of time without retaining all the application files. As a result, much Simpana software requires much smaller disk space to be provisioned for secondary copies.

Initially Simpana software will support Exchange and SharePoint applications for application-integrated backups and information mining, with more to follow.

Business Benefit:

- Integrated application protection and information mining allows the confidence to virtualize more applications, helping realize benefits and returns on the virtual platform sooner.
- Long-term retention of select individual objects does not require retention of entire application data, thus utilizing a lot less secondary storage.

Test Environment Setup

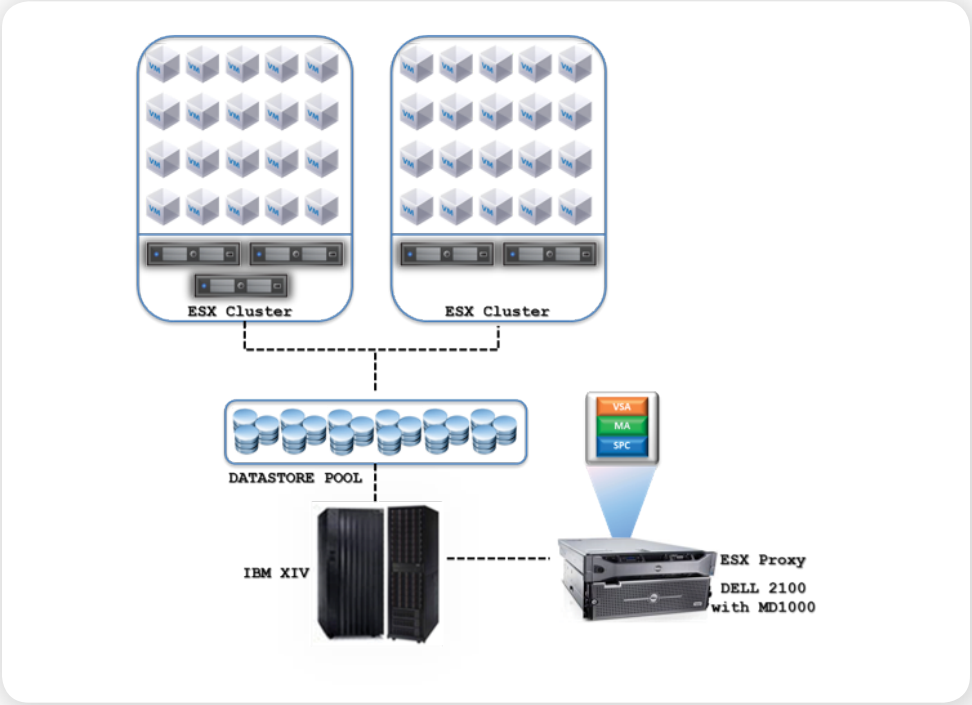

The diagram above depicts the environment used for testing SnapProtect for VSA.

VMware Setup

All of the tests were performed on vSphere 4.1.

ESX Clusters Configuration

Cluster Name	# of ESX Servers	ESX Server Configuration	Approximate # of VMs
DevCluster	2	8 CPU cores, 16GB RAM, 6 NIC ports	250
ProdcertCluster	3	4 CPU cores, 8 GB RAM, 2 NIC ports	250

Virtual Machine Configuration

# of VMs	Avg size of VM	Avg. Space Provisioned Per VM	OS	Total Size	Total Space Occupied
500	20 GB	17 GB	Windows 2008 Enterprise	10 TB	11 TB

Note: Due to the limited number of ESX servers, not all virtual machines were powered on during these tests. This does not have any impact on the test results, since the tests were set up to quesice virtual machines even when the VMs are powered off.

Just enough virtual machines were powered on to reach CPU and memory utilization of approximately 80% on each ESX server.

Data Store Configuration

Storage System	Number of LUNs	Size of Each LUN	# of VMs per LUN	Total Space Provisioned
IBM XIV	10	1 TB	50	11 TB

Simpana Software Configuration

CommServe®: Windows 2003 R2 x64 server with 8 GB of RAM

Backup ESX Proxy: Dell DL 2000 with MD1000

Processor	RAM	Network	Storage
2 CPU quad core for 8 cores total	16 GB	6 GigE ports	6 15K RPM internal SAS drives for hosting guest VM pagefile and dedupe database. 6 TB usable MD1000 with 7.2K RPM SATA disk used as backup to disk target.

Hosts a single virtual machine that contains the virtual server agent, MediaAgent and dedupe database.

Virtual Server Agent/MediaAgent: Configure on the following VM

Install IBM XCLI in the virtual machine for control access to the data store LUNs from the VSA agent.

OS	vRAM	vCPUs	Network
Windows 2008 Enterprise	16 GB	8	2 GigE ports

Test Execution and Results

SnapProtect™ for VSA

Test Case Configuration

For this test, we created a single backup set and configured 10 backup policies (subclients). Each subclient was configured to auto-discover with Datastore Affinity. This ensures that each subclient contains VMs from a single datastore, preventing any "dirty" or "hanging" VM images in the recovery copy.

Test Execution

All subclients were scheduled to start backup at the same time. This started to 10 different jobs. In the first phase of each job, the VSA discovers new virtual machines defined meeting the subclient discovery rules criteria. Once the VMs are discovered, each job interacts with vCenter separately to queisce virtual machines in each data store. Once the virtual machines in the datastores are all queisced, each job created a hardware snapshot for the datastore LUN. Once the hardware snapshot is executed, all the VMs in the data store are unqueisced.

The entire SnapProtect™ operation took approximately 17 minutes to complete. It took roughly 2 minutes to complete the discovery phase, then just over 14 minutes to run the SnapProtect jobs. Most of this time was spent in quiescing and unquiescing the VMs.

Test Outcome

Approximately 500 virtual machines protected in 17 minutes.

Conclusion

CommVault Simpana software offers a completely fresh approach to data management that is tuned for the needs of the emerging data center. Features like source-side deduplication are designed to ease the challenges facing traditional data centers. However, the new virtualized data centers need a lateral shift in thinking in terms of data management. Simpana software offers SnapProtect for VSA that leverages the snapshot capability to create fast recovery points with minimal impact on virtual servers, allowing users to protect hundreds of VMs in minutes.

Simpana software includes application-integrated backups and the mining of object level information for applications running inside virtual machines. Integrated reporting with SRM and automated archiving help users manage storage utilization in VMs more effectively. Finally, the ability to use cloud storage for long-term retention or archive targets ensures users can be well on their way in evolving a cloud-based data center that provides the automation, flexibility and scale needed to meet increasingly challenging business needs.

Simpana software's exclusive single-platform architecture gives companies unprecedented control over data growth, costs and risk. CommVault Simpana software was designed to work together seamlessly from the ground up, sharing a single code and common function set, to deliver superlative backup and recovery, archive, replication, search and resource management capabilities. More companies every day join those who have discovered the unparalleled efficiency, performance, reliability and control only CommVault can offer. Information about CommVault is available at www.commvault.com. CommVault's corporate headquarters is located in Oceanport, New Jersey, in the United States.

www.commvault.com ▪ 888.746.3849 ▪ E-mail: info@commvault.com

CommVault Worldwide Headquarters ▪ 2 Crescent Place ▪ Oceanport, NJ 07757 ▪ 888-746-3849 ▪ Fax: 732-870-4525

CommVault Regional Offices: United States ▪ Europe ▪ Middle East & Africa ▪ Asia-Pacific ▪ Latin America & Caribbean ▪ Canada ▪ India ▪ Oceania