

INSIGHT

V-Ray Gives Symantec NetBackup a Competitive Advantage Today and into the Future

Laura DuBois

Robert Amatruda

IDC OPINION

Over a decade ago, Veritas Software announced NetBackup FlashBackup to address the millions of small files problem, which had been and often remains the nemesis to fast and efficient backup of large file servers. The FlashBackup technology eliminated the random open-read-close process and associated disk I/Os for every file it was backing up. Instead, FlashBackup read file information directly off the raw disk partition holding the files, bypassing the file system all together. FlashBackup was able to achieve this by creating a map of a given file system's content and the blocks associated with any given file, and then writing the entire raw partition to tape. This approach allowed for significantly faster backups and a material reduction in I/O. Now, more than 10 years later, this same technology, although enhanced, forms the architectural underpinning of the Symantec NetBackup V-Ray technology and product vision. Today the FlashBackup technology is used to provide a logical understanding of what (file, application object, etc.) is stored with a VMDK- or VHD-image-level backup, without the necessity to install an agent inside each virtual machine. Further:

- ☒ Symantec's ability to understand what resides within the VM today provides advantages in backup and recovery but in the future can be extended to offer content and context-aware policy-driven IT controls such as data placement, provisioning, load balancing, and security.
- ☒ The visibility into file and application content of a virtual machine image gives Symantec a compelling competitive advantage today and into the future.

Note: This document focuses on the application of the V-Ray technology with Symantec NetBackup. However, the Symantec BackupExec product also leverages the V-Ray functionality for virtual machine backup and restore.

IN THIS INSIGHT

This IDC Insight looks at the V-Ray capabilities in the context of NetBackup virtual machine protection and recovery. Symantec V-Ray is at once a technology and a vision for Symantec's virtualization capabilities and future direction. However, V-Ray is also being leveraged across other Symantec product groups including BackupExec, security, storage foundation, and application availability and clustering.

SITUATION OVERVIEW

There are a plethora of virtual machine protection and recovery solutions available on the market today. These same solutions are all trying to deal with some fundamental challenges with the backup of virtual infrastructure today. Some of these challenges are the same issues that existed in the physical world, the most pronounced being increasing backup volumes and limited backup windows. However, placing multiple systems on a single ESX host has introduced a material network and storage I/O problem. This requires the effective coordination and balancing backups of many VMs residing on the same physical host across shared infrastructure. For example, multiple VMs that share a common LUN are going to face resource contention during streaming backup processes. There is often also CPU, memory, and/or disk I/O overhead, which can impact production applications and virtual machines. To address the I/O challenge with virtual machine streamed backups, VMware recommends a limit on the number of backup streams off the same physical LUN. Another key challenge associated with VM protection is the need for both image- and file-level recovery, without requiring multiple backups. These are all challenges that the Symantec V-Ray solution is expressly suited to address.

Symantec V-Ray

Symantec V-Ray is at once a technology and a vision for Symantec's virtualization capabilities and future direction. In the context of backup, V-Ray technology is embedded in the NetBackup Enterprise client, providing a range of benefits including single file and application object recovery from image-level backups and granular visibility into a given virtual machine's logical contents. For NetBackup operations, the V-Ray technology works hand in hand with the VMware vStorage APIs, specifically the vStorage APIs for Data Protection (VADP) to include full and incremental changed block tracking and application consistent recovery leveraging the VMware integration with VSS providers. A number of leading organizations such as FedEx and NBC Universal are using V-Ray technology in production today. Although this document is focused on NetBackup and VMware environments, the V-Ray technology works with both BackupExec and Microsoft Hyper-V environments as well. And Symantec has plans to extend the technology to other VI platforms.

V-Ray leverages the cornerstone FlashBackup technology previously mentioned to read a given virtual machine volume as it were a raw disk partition. By essentially bypassing the need to interact with the file system and have it perform file open, read, and close commands, the V-Ray technology can effectively read volumes that house VMDK and VHD files and look inside and index those virtual machine files without extensive file operations or disk I/O. This gives NetBackup the ability with a single-pass image-level backup to provide item- or file-level recovery of an object inside the machine and easily search files or file versions, which span backups.

The VMware API itself does not provide insight into what is inside a VM but instead enables backup of VMDKs as a raw blob or an image. Each supplier has come up with its own methodology to provide item-level recovery granularity. Some reply on a copy command to copy the VMDK to an NTFS disk staging area. Symantec was one of the first to market, several years ago, its item recovery from an image backup

capability, and other suppliers followed. However, today it is "how" Symantec achieves this benefit that is unique and efficient.

Image- and File-Level Recovery

The NetBackup Enterprise client is installed on an offhost backup server, and when NetBackup initiates a backup of a VMDK, the client collects all the information required for both an image- and a file-level recovery. No further operations are required on the production VMDK, which is important to offload overhead from production systems. The NetBackup client looks inside the VMDK image backup and identifies all the files inside the backup. The client uses a standard and an embedded indexing technology to perform this inline operation. NetBackup then catalogs all the files within a given image backup for fast search later. Once the backup is complete, NetBackup has full insight into all the contents within the VM. This process works hand in hand with the vStorage APIs for Data Protection (VADP) including Changed Block Tracking (CBT). NetBackup uses VADP to call VMware to create a VM-level snapshot, which is backed up directly from ESX storage. As the snapshot-based backup is in flight, NetBackup performs its inside VM indexing process. Even with incremental CBT backups, the NetBackup process has insight into the files inside the VM and does not need to recompile incremental and full backups in order to find and restore a single file. So NetBackup can extract a single file from either a full or an incremental CBT-based backup. And upon restore, NetBackup is looking inside its own index to locate files, not looking inside the production VMDK or a VMDK copy.

Load Balancing

The NetBackup VMware Intelligent Policy (VIP) feature allows for load balancing of backups against given resources. Administrators can set a policy to standardize on the protection of VMware VMs. A scenario of 50 VMs sharing a single LUN creates significant I/O contention on that LUN when backups for those VMs are started. All VMs are fighting for the same drive head and spindle resources. VIP will autodiscover all VMs and locate the LUNs they reside on (as well as discover newly added VMs). The policy will allow NetBackup to automatically determine how many backups are occurring on each ESX server and load balance backups across the environment. For example, it can throttle the number of simultaneous backups per ESX server, per LUN, and so forth.

Application Consistency

NetBackup gains application consistent recovery by leveraging the VMware tools and specifically the VSS provider, which is installed when installing VMware tools. NetBackup calls the VMware provided VSS provider, which in turn quiesces the virtual machine and the guest application. This process creates an application consistent point in time checkpoint, which can be used for backup and recovery. The challenge with this process is that it provides an "all or nothing" restore of an application or database without any truncating of logs. As a result, Symantec created an augmentative VSS provider to do additional functions such as log truncation and support new operating system versions and newer applications such as Exchange 2010. This augmentative VSS provider also enables more granular restore rather than just an all or nothing restore and can enable application object-level restore

(mailboxes, SharePoint objects, etc.) (the granular restore capability is available in NetBackup 7.5).

Efficient Deduplication

Since the V-Ray technology is embedded into the NetBackup enterprise client, the same client can perform deduplication on data within physical systems and virtual machine images. Common data blocks are deduplicated across physical and virtual resources. For example, files or Windows binaries inside a VMDK, which also resides on a physical server, will be deduplicated in a shared deduplication pool.

Alternative Approaches

As mentioned previously, other suppliers also offer item-level recovery from an image backup. Each supplier does it its own way although almost all leverage the VMware VADP mechanism to create a VSS provider-consistent virtual machine snapshot. However, firms evaluating solutions should be cognizant of the different approaches used to achieve the same gain.

Some vendors enable this item-level recovery by taking a given VMDK and copying it to an NTFS disk, and if a file inside the VMDK is required, the VMDK is mounted like a standard network drive. Administrators can browse for the needed file and copy it over to a desired location. This introduces some limitations. The VMDK backup must be on disk (versus tape) and requires disk capacity. While this may seem trivial for a single VM, for thousands of VMs this can introduce an undesired effect. Last, this does place a restriction on what kind of disk target can be used. This may prove problematic if a firm has invested in a VTL, for example. Firms should also evaluate how easy it is to initiate a search and restore for a file, in particular, whether multiple mounts of multiple VMDK backups are required and what impact a VMDK backup indexing process has on production VMs themselves.

FUTURE OUTLOOK

The Symantec V-Ray technology provides Symantec with an optimal approach to addressing many of the virtual machine backup challenges firms face today. In summary, NetBackup's use of the V-Ray technology provides:

- ☒ Immediate indexing of all files and application objects during the backup process for efficient and fast search during a recovery or a restore process
- ☒ Instant visibility and access to files inside VMDKs once the backup is complete
- ☒ Support for all media targets including VTL, target-side deduplication systems, and tape
- ☒ Simple, user-driven, single-step restore processes for application items and files inside a VMDK backup

While Symantec V-Ray is relevant for backup today, in the future Symantec can leverage the V-Ray technology to determine the optimal information security, data placement, and availability policies based upon the knowledge it has of application objects and files resident inside a virtual machine image.

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights.

Copyright 2011 IDC. Reproduction is forbidden unless authorized. All rights reserved.