VMware VDR and Cloud Storage: A Winning Backup/DR Combination

7/29/2010

CloudArray™, from TwinStrata, and VMware Data Recovery combine to provide simple, fast and secure backup: On-site and Off-site

The advantages of Deduplication, Block-Level updates, and VSS integration, with a solution providing both local and Cloud based backup to disk.

www.TwinStrata.Com
24 Prime Parkway
Natick, MA 01760
508-651-0199

Executive Summary

CloudArray™, from TwinStrata, and VMware Data Recovery (VDR) provide simple, fast and secure off-site backups for VMware environments to your choice of cloud storage providers - without the hassles of Tape.

- Simple: CloudArray interfaces natively to VDR
- Efficient: De-duplication results in only changes sent over the network
- Secure: Data encrypted at-rest and in-flight
- Fast: CloudArray cache or local copy for fast restores
- Available: CloudArray snapshots create additional copies in the cloud
- Economical: Establishes a pay-as-you-go usage model with unlimited elastic capacity

Introduction

A common headache for VMware administrators is backup. With many Virtual Machines (VMs) to back up on a vSphere server, you can easily consume valuable system resources pushing multiple backups to a traditional backup server. VMware consolidated backup (VCB) is one way to consolidate and simplify the backup process, but it requires a substantial investment in backup application software that is compatible, and a tape or local backup to disk (B2) solution to manage. Also, VCB is being phased out in favor of the new vStorage API initiative with vSphere.

With vSphere, you have the alternative no-charge option of using VMware Data Recovery (VDR) to backup VM images to disk in a consolidated manner. VDR uses changed track updates with a native deduplication capability to optimize backup performance while minimizing capacity requirements.

TwinStrata CloudArray enhances VDR further, enabling VDR to backup to a CloudArray volume that is stored at a remote Cloud Service Provider or to your own Private Cloud infrastructure. This hybrid solution is termed Backup to Disk to Cloud (B2D2C) and seamlessly enables off-site protection for VMware backups.

Most disk based Disaster Recovery (DR) solutions require large investments to fit a recovery site with hardware and personnel. A Cloud Storage solution, enabled by CloudArray, eliminates that upfront investment. Your data is safe, secure and available for recovery any time and from virtually anywhere.

CloudArray B2D2C enables significant cost savings for off-site backup. When compared to the costs associated with traditional tape backup (Tape Library costs, software, media, transportation and storage, etc), the use of VDR with CloudArray is a clear winner. You save about 46% over the traditional tape model for a 5 TB target configuration. It's even better compared to commercial (not consumer) online backup vendors or remotely hosted backup to disk services.

-

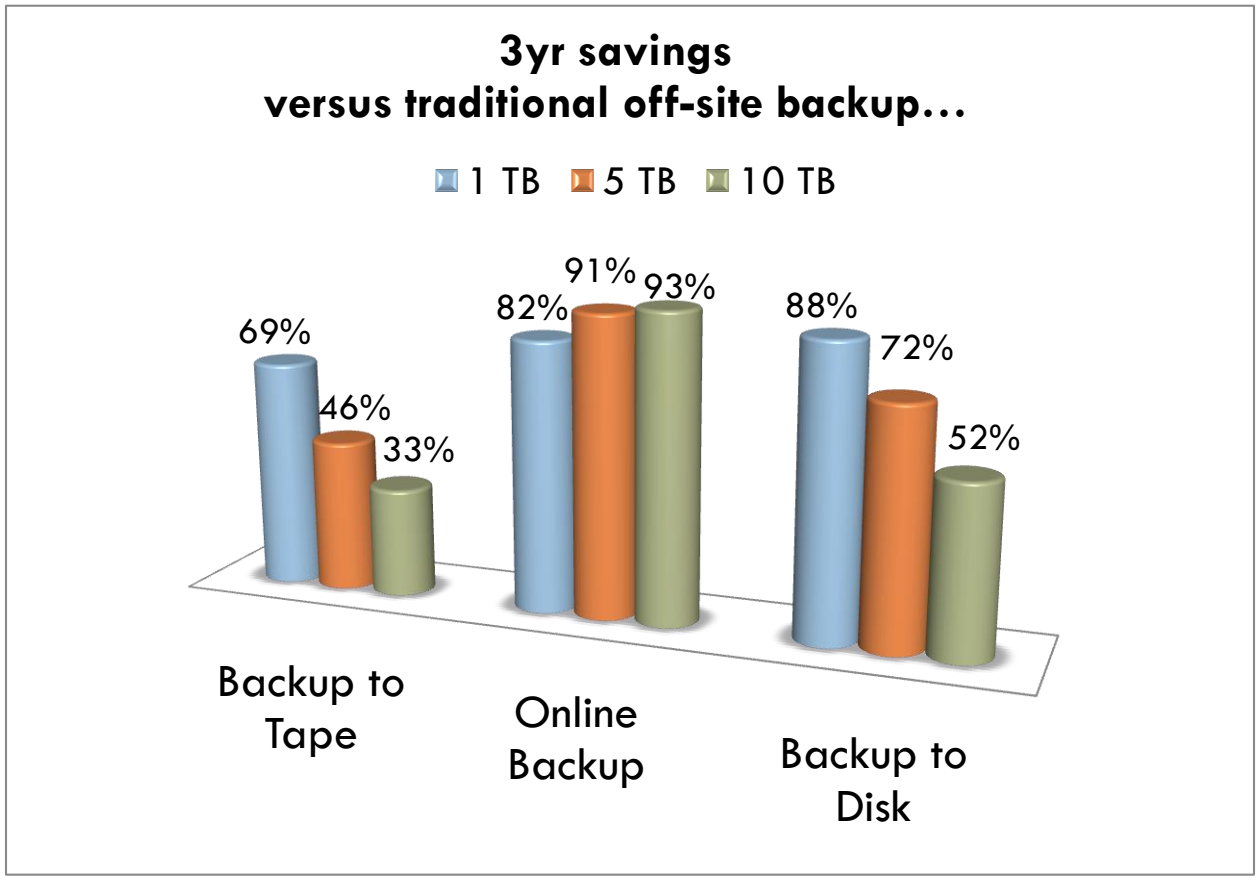

VMware VDR Overview

VMware Data Recovery (VDR) isn't new. It's been around since the 2009 release of vSphere as a no-charge option, but it is gaining some new attention.

VDR performs de-duplicated B2D. And now those backups can be stored remotely on Cloud Storage simply, securely and cost effectively. VDR with CloudArray: B2D2C in one easy step.

VDR, like the CloudArray, runs as a VM under vSphere. They are easy to deploy and configure. Both can be installed, configured, and running in less than 30 minutes. That's a fully configured solution; not a test configuration.

VDR Changes the Backup Paradigm

VDR not only performs de-duplication, it also uses changed track updates. What this means is that the notion of full and incremental backups is gone. All backups are incremental after the first instance, and all updates are deduplicated against all VM's backed up to that datastore.

VDR deduplication is automatic: you can't turn it off. During backup, VDR creates a snapshot of the VM in a data consistent state (quiesced). Using these methods, the applications and file systems are quiesced and the backup data sets are "crash consistent" and recoverable. There is no need to shut down the VM's prior to running the backups. Of note, VDR also supports VSS (Microsoft's Volume copy Shadow Service) for Exchange Server, SQL Server, and with vSphere 4.1, Windows Server 2008. Additional information on VSS integration can be found in the VDR Administration Guide (VMware EN-000193-00).

After you create the first image backup (and remember, it is doing block level de-dupe so even that is a much smaller image than your source), every backup after that only sends the changed data and also de-duped. In addition, VDR maintains restore points based on every backup run, so you still have multiple restore points to choose from, just like with traditional backup solutions. Of course, it's all disk based so there is no need to manage tapes.

Backup to Disk is Good, but Backup to Cloud at the same time is even better

Combine those capabilities with CloudArray, and you can remotely replicate backups to a Cloud, Public or Private, with a choice of providers. You get the peace of mind of a

disaster-tolerant infrastructure without the added expense of a second location. Your backups are securely replicated offsite without any manual intervention.

Keep BOTH a Local and a Remote Copy at the Same Time

With CloudArray, you have the option to maintain either a full local copy, or a partial copy in CloudArray cache. The advantage of a full local copy is that you will never have to go to cloud storage to retrieve your backed up VM's unless you have a major site disruption. In either case, you have the option to take snapshots in the cloud, thus preserving older restore points and allowing you to maintain the consistency of the backup set in the event something happens to your on-line copy.

If needed, these remote backups can be restored to any location you choose, even in the cloud, from instances of CloudArray running on Amazon EC2 (Elastic Cloud Compute service). This greatly expands your options if a disaster recovery site ever needs to be stood up.

Of course, being a B2D2C solution, the need for tape is eliminated, as is the need to manage those tapes or transport them to offsite storage facilities. And since the backups are on-line, you can restore from either local cache or the Cloud without having to schedule a delivery to go and try to find them.

VDR Features

VMware VDR supports an impressive list of features, including:

- File level restore client for Windows and Linux virtual machines
- Ability to backup up to 100 VM's per VDR appliance
- Ability to run up to 10 VDR appliances per vCenter Server instance
- Ability to fast switch between the deployed VDR VM's via the vSphere Client plug-in
- Deduplication and track-level updates
- Multiple restore points
- Enhanced VSS integration for Exchange, SQL Server, and Windows Server

A quick note on VDR limitations

VDR has limitations: it only supports 100 VM's and 2TB of de-duped/compressed target data per VDR instance (version 1.2 supports up to 10 VDR instances per vCenter server); and it doesn't support linked vCenter servers or linked clones. But even with those, it can be a viable solution to real headaches you are fighting today. And, given that it's a

first generation product, you can be sure VMware will improve and expand its capabilities.

To expand on this briefly, the current (Version 1.2) limitations for VDR are that it can support up to 1000 VM's per vCenter server, with a total of 20TB of compressed and Deduped target storage (representing up to 400 TB of primary storage).

CloudArray Overview

CloudArray, TwinStrata's Cloud SAN software gives cloud storage the look, feel and performance of local data storage. CloudArray provides comprehensive enterprise-class storage features that enable delivery of cost-effective and simply managed off-site storage optimized for performance, availability and elastic storage capacity.

CloudArray is a virtual appliance supporting all the market-leading hypervisors: VMware ESX/ESXi; Citrix XenServer; and Microsoft Hyper-V. It is also available in a physical appliance for customers that want a dedicated physical server to run CloudArray.

CloudArray will natively run as a VM in vSphere along side VDR, and will provide native iSCSI storage to both VDR and to vSphere in the form of RDM (Raw Data Mapping) or VMFS (Virtual Machine File System) volumes. CloudArray can also run natively on Amazons EC2 Cloud Computing service.

CloudArray supports internal or private storage clouds for companies that require data on-premise or within the walls of their own datacenter. Companies can take advantage of the on-demand storage distribution model of storage for each line-of-business.

CloudArray also supports on-demand storage-as-a-service from cloud providers. CloudArray is installed locally in the data center and connects to the cloud provider via IP (Internet Protocol). The data is safely stored offsite, and retrieved from any site of your choosing if you have the proper credentials and keys.

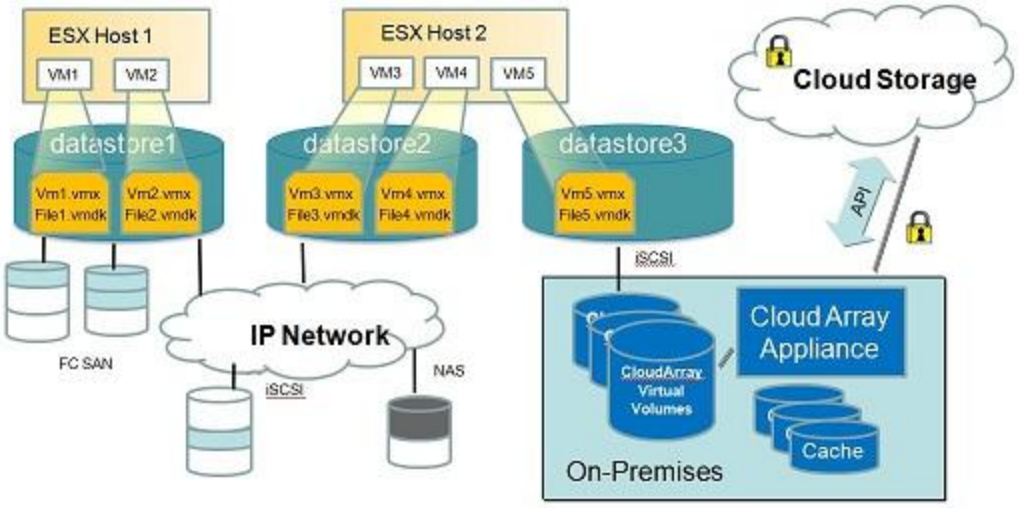

Seamless Integration with Applications, Physical and Virtual Servers

CloudArray seamlessly and securely integrates existing business applications with cloud storage without any changes to applications or programming of cloud provider APIs (Application Programming Interface). CloudArray supports applications running on both physical and virtual servers.

In the VMware environment, CloudArray volumes can be presented as native iSCSI devices to VMware and used as either RDM or VMFS devices by any VM. The CloudArray volumes can be shared across VSPHERE servers, so they can be configured to support HA, DRS, vMotion and even Storage vMotion.

Additionally, the CloudArray VM itself can be configured using either HA or VMware Hot standby for added availability across vSphere servers.

Cloud Storage Volumes

CloudArray volume policies allow data to be cached locally for quick access with local performance levels and replicated to Cloud storage for high availability. CloudArray provides data snapshots for point-in-time copies for rollback, test and development, data analytics and for other applications. CloudArray volumes can be configured to be either partially or fully cached, depending upon the application. The latter option allows for consistent copies of volumes in two separate locations; local to your data center as well as in the Cloud (private or public).

Data Availability and Security

CloudArray optionally compresses and encrypts data prior to transporting it to public cloud storage to ensure data security and privacy. CloudArray leverages the Advanced Encryption Standard (AES 256 with HMAC).

Configuring and Using the VDR and CloudArray together

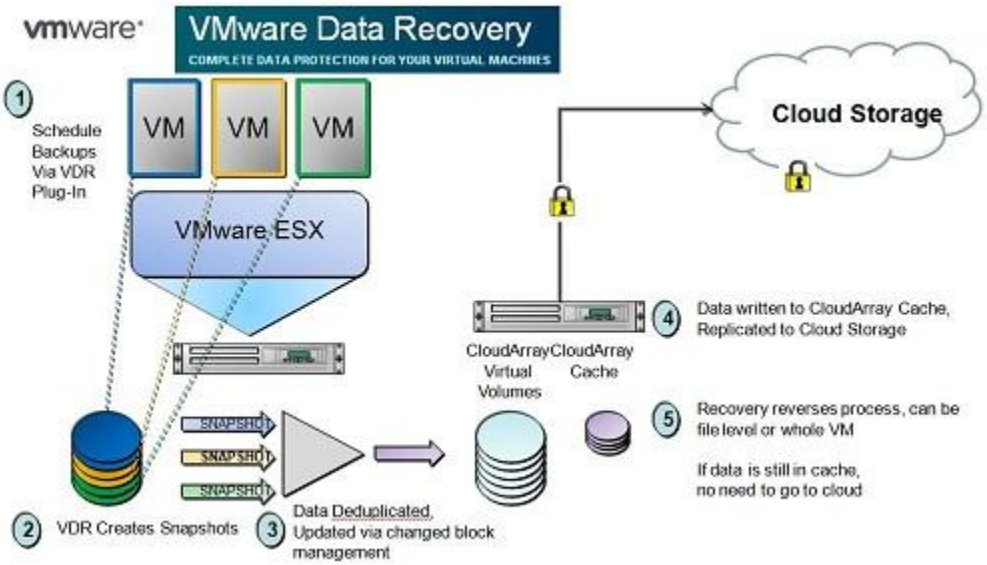

It's easy to perform a basic configuration install and configuration of VDR and CloudArray. Both are deployed as VMs using an OVF (Open Virtualization Format) distribution. VDR is managed using a VMware plugin. CloudArray is managed using CloudControl, a small client application that runs on Windows and Linux operating systems. You can find a video showing how to install and configure CloudArray at http://www.twinstrata.com/Cloudarray_evaluation.

Once CloudArray is installed, the next step is to assign it local cache. Since the largest target volume that VDR will support is currently 2TB per VDR instance, the local cache could be any size up to 2TB. In practice, larger local cache volumes will deliver higher local performance, and a fully cached volume (a 2TB local cache volume) in this case will also mean that VDR will never have to restore from Cloud Storage unless there has been a catastrophic event of some kind and the local copy is destroyed.

Next, a CloudArray "policy" is created, which basically assigns a Cloud Storage provider and selected options (such as encryption) to a local cache pool.

Finally, you create virtual cloud volumes that are then presented to VSPHERE using iSCSI mount points. The resulting datastores can be assigned to VDR as either RDM devices or VMFS devices. You would then format and mount the device from within VDR.

At this point, VDR is ready to roll. Any backups it writes to the new datastore will be replicated to your Cloud Storage provider.

Other Use Cases

Archival

Another interesting use case is archival of older VM's, Templates, or infrequently used VM's such as test environments. In these cases, you can use Storage vMotion to migrate the datastore for these VM's from your primary SAN to Cloud Storage. They will still appear as local to your ESX server, but the data will be kept resident in offsite Cloud Storage.

If you need to use them they can either be used with their datastore on Cloud Storage or they could be migrated back. You can also Clone VM's, or Clone VM's to a Template with the target being cloud storage volumes.

Migration/Replication

By now you are seeing that Cloud Storage can be an excellent way to store VM's you don't need immediate access to, either by using a backup product such as VDR or by simply archiving your VM's. But there is another valuable use as well: as a migration "swing" set. For example, if you have some VM's that you want to transfer between datacenters, or even between VSPHERE hosts, you can use Cloud Storage as a place to first write the VM files from one host or site, and then read them to another host or site.

Conclusion

VMware and Cloud Storage are a natural fit. Virtualization has proven itself as an excellent means to get efficient use out of data center resources; memory, CPU and network. Extending this model by virtualizing your storage to include Cloud Storage for backups has

compelling economics, cost-effectively extending your data center and backup copies off-site.

Cloud Storage provides elastic, unlimited capacity on a pay as you go, self-provisioning model. CloudArray optimizes your use of Cloud Storage even further in a VMware environment. It is VM based, interfaces with ESX/ESXi with the use of native iSCSI connections, and provides important value added features from encryption to in cloud snapshots.

Visit www.TwinStrata .com for more information and take advantage of our Free 30 day trial.