vSphere Backups to Amazon S3

Common architectures used by APN technology partners

December 2019

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

Abstract

This paper will act as an introduction to the features, Amazon Web Services (AWS) services, and architectural components relevant when backing up an on-premises virtualization environment based on VMware vSphere to Amazon Simple Storage Service (Amazon S3). Specifically, products and features from AWS Partner Network (APN) partners Veeam, Rubrik, Cohesity, and Druva are described to provide the reader with examples of common approaches taken by APN Technology Partners.

This paper is intended for solution architects familiar with VMware virtualization. Readers are presumed to have a knowledge level equivalent to a VMware Certified Professional in Data Center Virtualization for vSphere version 6 (VCP-DCV6).

Introduction

Since the release of vSphere 4.0, VMware has offered several mechanisms for conducting agentless backups known collectively as the vSphere Storage APIs for Data Protection (VADP).

In the 10 years since VADP's release, many customers have benefitted from the use of this interface that it's become table-stakes for VMware specialists when addressing the recoverability needs of a design. At the same time, even more customers wish to avail themselves of the benefits of storing their backup data in the cloud.

This paper will examine the common architectural approaches taken by APN Technology Partners when using VADP to back up virtual machines (VMs) to Amazon S3.

Terms and Definitions

vSphere Storage APIs – Data Protection

Transport Modes

Any data protection solution which utilizes the vSphere Storage APIs for Data Protection must choose from a collection of data transport mechanisms that specifically define how to copy the VM files from the running shared storage to the backup target.

SAN Transport Mode

In this transport mode, a data mover server mounts the running shared storage volume as a LUN (only relevant in an iSCSI or Fibre Channel based Storage Area Network (SAN) environment). VM files are backed up by the data mover over the SAN fabric directly.

NBD Transport Mode

Each vSphere host has at least one VMkernel IP interface enabled for management traffic. This IP is used by vCenter to control individual vSphere hosts within a cluster, and also provides an interface for an administrator to connect directly to on a host for setup and troubleshooting.

The NBD transport modes utilize this management interface to copy the VM disk files across the network to the backup target. If more than one VMkernel interface is enabled for management traffic on the host, it will use the one that has a preferred route to the target IP address.

NBDSSL Transport Mode

NBDSSL is simply an SSL-encrypted version of NBD. It is the default used by ESXi 6.5 and greater.

While encrypting the backup traffic from the vSphere host to the data mover of choice is generally desirable, consider the following when choosing between NBD and NBDSSL:

- Traffic from the data mover to Amazon S3 is always TLS-encrypted.
- The CPU overhead incurred on each vSphere host results in up to 30% less throughput. For instance, a backup job that takes 15 hours to complete over NBDSSL might take as few as 12 hours using unencrypted NBD.

HotAdd Transport Mode

This transport mode utilizes a backup proxy that is itself a virtual machine running within the same vSphere environment. The backup process is multi-stage and proceeds as follows:

- Snapshot(s) of the live VM's virtual disk(s) are taken and mounted to the backup proxy as if they were normal virtual disks.
- Depending on the vendor implementation, VSS or other freeze/thaw mechanisms may be invoked via VMware tools to quiesce at the application and/or volume level when the snapshot is created.
- The data protection solution then copies the contents of these snapshot-based virtual disks to the backup target across the virtual network interface of the backup proxy VM.
- When the backup is complete, the snapshot is deleted.

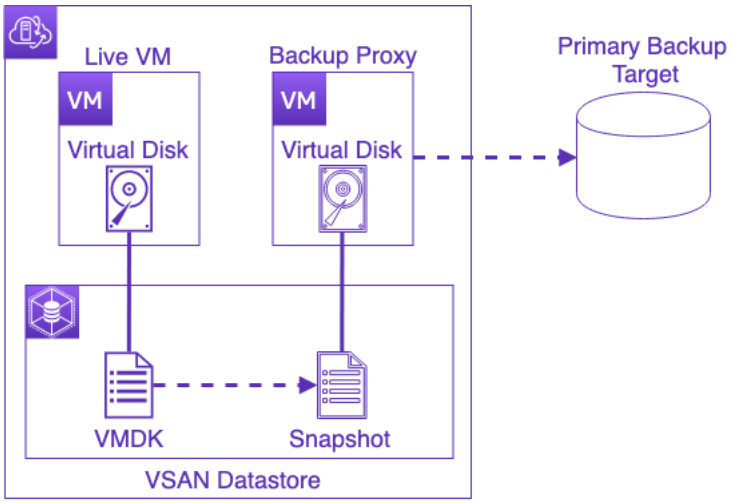

Figure 1 - HotAdd Transport Mode

Normally the HotAdd Backup Proxy mounts multiple snapshots simultaneously. This allows it to back up several VMs in parallel.

Changed Block Tracking (CBT)

Changed Block Tracking is an optional feature offered by VADP that enables vendor solutions to easily conduct incremental backups of virtual disk files. CBT tracks block-level changes to the virtual disk files

between backups. The backup solution is then aware of what specific blocks are different from the last time it backed up that VM.

Data Protection Solution Components

Distributed Data Protection Solutions

Many data protection solutions are deployed as separate components in a distributed architecture. In the context of this paper, each of these components runs as a process within the guest operating system of an Amazon Elastic Compute Cloud (Amazon EC2) instance, or a vSphere-based Virtual Machine (VM).

While these solutions are often deployable on a single server that hosts all components, such configurations are generally limited to protecting small numbers of VMs (~100 or less). The more VMs that need to be backed up, the more distributed these components become.

> **Note:** Fully distributed architectures tend to suit customers with large numbers of VMs and/or complex heterogeneous environments who want a unified solution.

Command and Control Server
This component acts as the management plane, which orchestrates all other components.

Typical functions performed include:

- Administrative interface that handles job scheduling, reporting and alerting.
- Interacting with vCenter to initiate VM snapshots or NBD backup operations.
- Deployment and updating of other components in the system.

Data Mover

Typical functions of data movers include:

- Streaming data from the vSphere environment over the network to the primary backup target (i.e. a Disk or Cloud Repository).
- Deduplication/compression or other data efficiency operations before transmitting to the backup target.

HotAdd Backup Proxy
This is a special type of Data Mover that is required when using the HotAdd transport. It is a VM that resides within the protected vSphere cluster, as it requires direct access to the same shared storage as the VMs it is protecting.

Cloud Repositories – Native Backups to Amazon S3

Amazon S3 provides a highly scalable and cost-effective storage solution that is ideal for backups. It is designed for 99.999999999% durability (eleven nines); all objects stored within an Amazon S3 bucket are automatically copied to multiple devices spanning a minimum of three Availability Zones.

Amazon S3 offers six storage classes, ranging from active classes like Amazon S3 Standard to archive classes like Amazon S3 Glacier Deep Archive. For more information about Amazon S3 storage classes, see the storage class details on the AWS website.[1]

Each Amazon S3 storage class has a distinct set of properties meant to optimize cost for a given access pattern and/or performance requirement. In the case of backup data, considerations regarding Recovery Time Objective (RTO) are additional key drivers – this is particularly relevant when deciding whether or not to place data in Amazon S3 Glacier or Amazon S3 Glacier Deep Archive. Finally, it is important to ensure the data protection solution provided by the APN technology partner of choice supports the desired storage class.

Amazon S3 also supports a wide variety of management features and security controls to help you view, manage and secure the data stored on Amazon S3.

Cloud Repository Proxy – Assisted Backups to Amazon S3

This is a component that front-ends Amazon S3 endpoints for one or all of the following reasons:

- To present a non-Amazon S3-native interface to the backup solution. Virtual Tape Library (VTL) or Network File System (NFS) are the two most common types.
- Mitigation of Bandwidth Delay Product issues when latency is high to the Amazon S3 endpoints. This is usually accomplished through network protocol manipulation.
- Caching backup data locally before transmission to Amazon S3. This assists with RTO adherence in environments with limited throughput available to Amazon S3.

> **Note:** Some solutions use AWS Storage Gateway for this purpose.

Disk Repositories – Backing up to block storage

Disk repositories are servers that directly store backed up VMs on some type of block storage device. Possible examples include:

- VM in a local vSphere environment running Windows or Linux with a large virtual disk back-ended by a LUN on a SAN array. It could be self-built or a virtual appliance.
- Physical appliances (or clusters of appliances) containing direct-attached storage that present a CIFS, NFS, or iSCSI interface to the hosts in a vSphere cluster.
- Amazon EC2 instances (or clusters of them) with a combination of NVME-backed instance storage for caching and Amazon Elastic Block Store (Amazon EBS) volumes for storing backup data.

Hyperconverged Data Protection Solutions

HCI combines the block storage necessary for a disk repository with the data protection solution itself. Consisting of horizontally-scaled clusters, each node added contributes disk storage to the cluster and can perform all of the functions of the components described above.

> **Note:** Hyperconverged solutions tend to fit customers who primarily run workloads on vSphere and are looking to simplify their backup infrastructure.

SaaS-Based Data Protection Solutions

Some partners offer their solution as a fully managed service. Components deployed in the on-premises environment are limited and configuration on the customer's part is minimal. While VM backups are stored on Amazon S3, all AWS services involved are configured, maintained, and billed on the customer's behalf by the SaaS provider.

> **Note:** Customers seeking the simplest solution with the most rapid time-to-value benefit the most from this type of architecture.

Important Features

Client-Side Data Efficiency
Data efficiency mechanisms such as deduplication and/or compression that occur on the source itself before transmitting data to the backup target, for instance, a HotAdd Backup Proxy that eliminates duplicate blocks inside VMDKs before sending.

Backup solution vendors sometimes quote efficiency ratios as high as 50:1. Whether this is achieved in practice varies according to a number of variables including:

- Redundancy and compressibility the data in the protected VMs.
- If data within the VMs is already deduplicated and/or compressed (i.e. MPEG)
- Solution-specific details. e.g. fixed-length vs. variable-length deduplication.
- Resources such as vCPU dedicated to this task on the HotAdd Backup Proxy.

Global Deduplication
The ability of a given solution to deduplicate blocks, objects, or files across all customer data – regardless of the backup repository type. For instance, a global namespace that can deduplicate backups spread across Amazon S3, Amazon EBS, and/or on-premises disk repositories.

Solutions that incorporate this type of feature often greatly reduce the monthly storage expenditure necessary to maintain a given retention and tiering strategy.

Consistency of Volumes and Applications

Raw snapshots of running VMs will capture the point-in-time state of virtual disks, regardless of any incomplete IO operations that might be occurring.

When a snapshot of a VM starts, if the VMware tools are installed, they will communicate with the Volume Shadow Copy Service (VSS) (Windows VMs) or vmsync driver (Linux VMs) to perform a freeze operation on the attached volumes. This commits in-flight IOs via mechanisms such as flushing in-memory write buffers to disk. When this is finished, vSphere is notified, the snapshot occurs, and a thaw operation releases the volume to continue IO.

While this protects the volume itself, application-level operations that might be in progress are unknown to a volume-level quiescence provider. This is known as a *volume-consistent backup*.

If VSS writers or native vendor-provided drivers are registered, supported processes such as Microsoft SQL Server or Oracle RDBMS will be notified. A similar, but application-specific, quiescence procedure then occurs.

A snapshot that quiesces applications in this way is known as an *application-consistent backup*.

During any such process, VMs are stunned for a short period (normally measured in microseconds). Large VMs that are resource-intensive on a consistent basis may experience noticeable stun periods that result in perceptible service interruption (pauses of several seconds are not uncommon). Some vendors provide native application-integrated drivers designed to eliminate this issue.

Partner Solutions

This section describes the architecture of a few data protection solutions that back up on-premises VMware environments to Amazon S3, but is not a comprehensive list of APN Technology Partner offerings. It is merely a sample meant to illustrate common approaches taken by APN Partners in this space.

Druva Phoenix

Overview

Phoenix for VMware is a hybrid backup solution delivered via a SaaS model. Components of a Phoenix for VMware deployment include:

- **Backup Proxy** – Acts as either a HotAdd Backup Proxy or a simple Data Mover depending upon the Transport Mode in use. Delivered as a virtual appliance.

- **Phoenix CloudCache** – Optional component that acts as a Cloud Repository Proxy. It can cache up to 30 days of backup data, which it synchronizes to Phoenix Cloud at configurable intervals – for example, bandwidth consumption could be throttled during working hours, but unlimited in the evening. Installed via MSI into Windows.
- **Phoenix Cloud** – From the customer's perspective, this is a single logical component that handles Command and Control, and presents a single backup repository to which all VM backups are sent. On the back end, AWS services are consumed, but these services are not directly managed by the customer.

Note: Multiple Backup Proxies and CloudCache servers may be deployed depending on the number of protected VMs, throughput and RTO requirements.

Logical Architecture

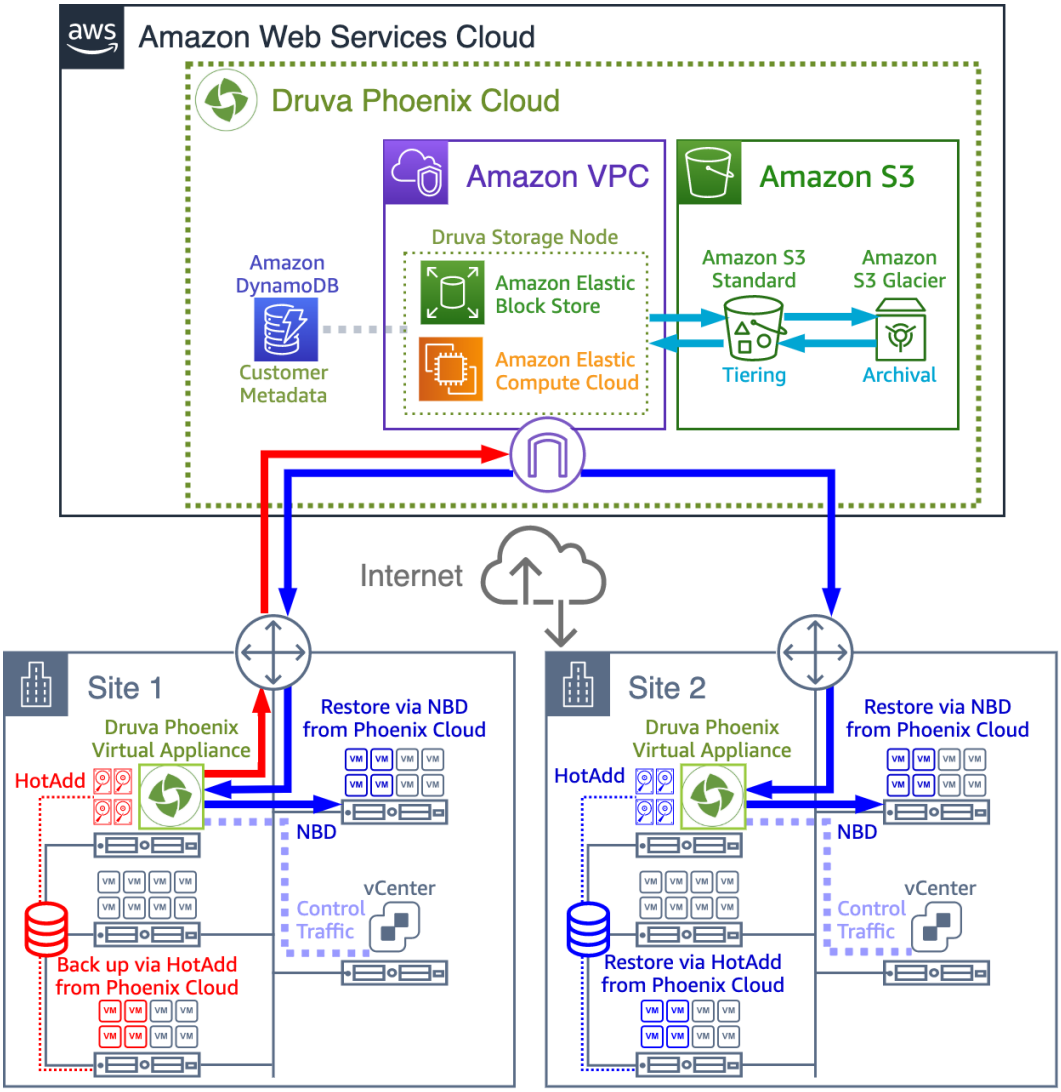

Figure 2 - Druva Phoenix for VMware Logical Architecture

Illustrated Workflows

- Back up of VMs to Phoenix Cloud via backup proxy using HotAdd transport
- Tiering down of hot data from storage nodes to warm tier on Amazon S3
- Archive older backup data from Amazon S3 to Amazon S3 Glacier
- Recovery of VMs from Amazon S3 to the original site via NBD transport mode
- Recovery of VMs from Amazon S3 to Site 2 via HotAdd transport mode
- Recovery of VMs from Amazon S3 to Site2 via NBD transport mode

Feature Support

VADP Backup Features		Amazon S3 Feature Support	Commercial	GovCloud
Application Consistency	Yes [1]	**Amazon S3 Standard**	Yes	Yes
Volume Consistency	Yes [2]	**Amazon S3 Standard-IA**	Yes	Yes
Source Side Compression	Yes [3]	**Amazon S3 One Zone-IA**	Yes	Yes
Source Side Deduplication	Yes [3]	**Amazon S3 Intelligent-Tiering**	No	No
Global Deduplication	Yes	**Amazon S3 Glacier**	Yes [4]	Yes [4]
Changed Block Tracking	Yes	**Amazon S3 Glacier Deep Archive**	No	No
Single File Restore	Yes	**Amazon S3 Cross-Region Replication**	No	No

[1] For registered VSS Writers in Windows Guests running VMware Tools

[2] For Windows Guests running VMware Tools

[3] Source-side in this context refers to the Backup Proxy appliance, not the vSphere hosts

[4] Indirectly supported via tiering after a default aging period of 90 days

Table 1 – Druva Phoenix for VMware 4.8 Feature Support

Further Information

For additional information, see:

- Blog - Backing up VMware Cloud on AWS with Druva Phoenix[2]
- Whitepaper – Cloud-native backup and recovery for VMware[3]

Cohesity DataPlatform

Overview

Cohesity DataPlatform is based upon a hyperconverged architecture. Each appliance in a cluster contributes locally-attached storage to a highly available, dynamically optimized virtual storage pool.

The direct-attached storage in each node is either physical HDD/SDD drives, or a combination of Amazon EC2 instance stores and Amazon EBS volumes.

DataPlatform offers a variety of storage-related services. This section is focused on how to use it to conduct VADP-style backups of on-premises vSphere clusters to Amazon S3.

There is one supported approach relevant to this use case:

- **Amazon S3 as an archive repository** - The most recent backup data (typically 30 days) is kept locally on the cluster. Once it has aged past that, it is then archived to one of the Amazon S3 storage classes for long-term retention.

Note: Archived backup data is fully self-describing. An Amazon S3 bucket containing a complete backup set and a Cohesity cluster that can read it is all a customer needs to conduct restore operations.

Logical Architecture

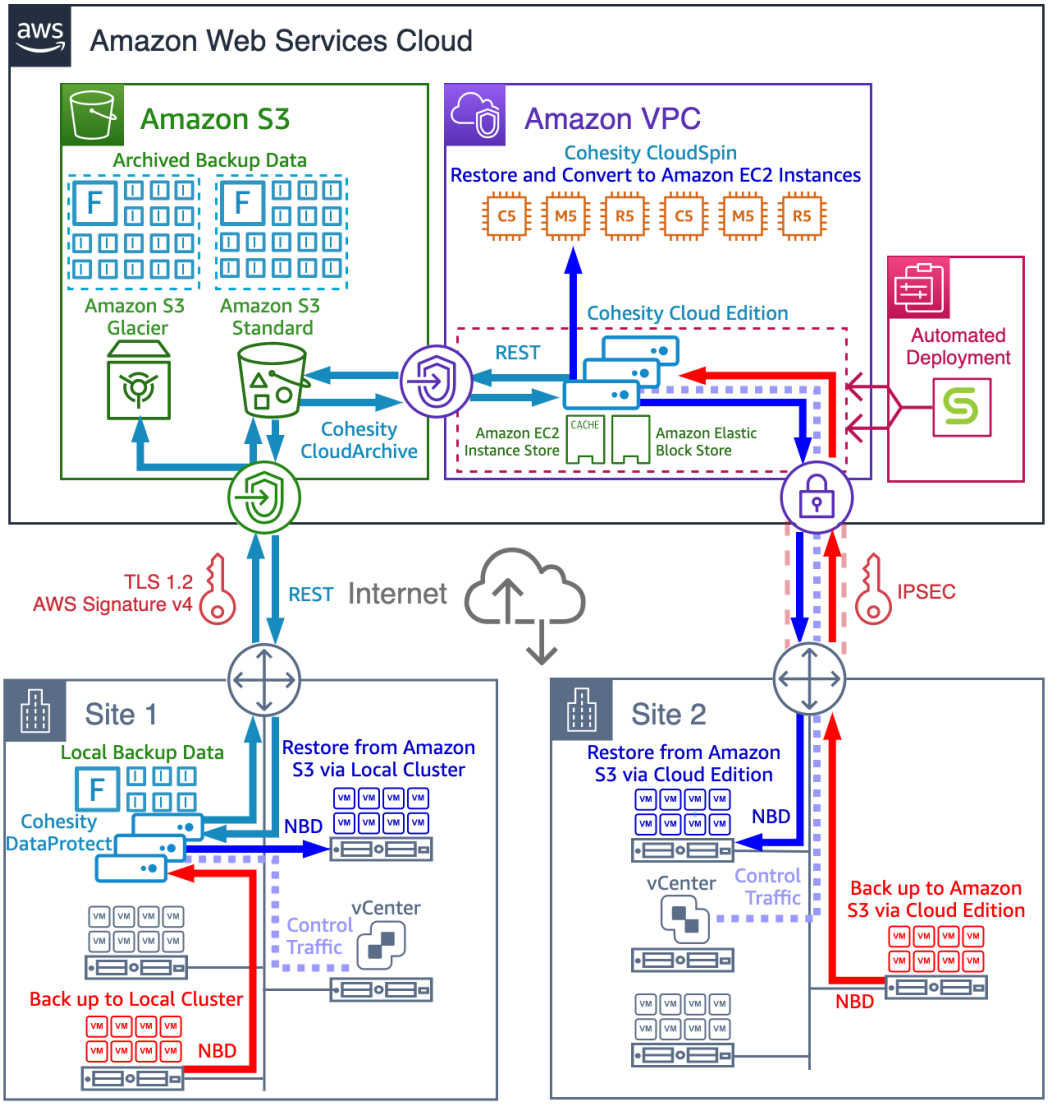

Figure 3 - Cohesity DataPlatform Logical Architecture

Illustrated Workflows

- Back up of VMs to local DataPlatform cluster in Site 1
- Back up of VMs in Site 2 to a DataPlatform Cloud Edition cluster in AWS

> **Note:** Automated deployment of a Cloud Edition cluster is initiated through Helios, a monitoring and management service offered by Cohesity.

- Archive old backups from DataPlatform Cluster in Site 1 to Amazon S3

- Archive old backups from DataPlatform Cloud Edition cluster to Amazon S3
- Archive via AWS Lifecycle policy from Amazon S3 to Amazon S3 Glacier
- Recovery of VMs from Amazon S3 to the original DataPlatform Cluster in Site 1
- Recovery of VMs from Amazon S3 to Site 2 via DataPlatform Cloud Edition
- Conversion of VM backups in Amazon S3 to Amazon EC2 instances leveraging the CloudSpin feature of the DataPlatform Cloud Edition cluster located in AWS

Note: DataPlatform 6.4 does not support the HotAdd Transport Mode.

Feature Support

VADP Backup Features		S3 Feature Support	Commercial	GovCloud
Application Consistency	Yes [1]	**S3 Standard**	Yes	Yes
Volume Consistency	Yes [2]	**S3 Standard-IA**	Yes	No
Source Side Compression	Yes [3]	**S3 One Zone-IA**	Yes	Yes
Source Side Deduplication	Yes [3]	**S3 Intelligent-Tiering**	Yes	Yes
Global Deduplication	Yes	**S3 Glacier**	Yes	No
Changed Block Tracking	Yes	**S3 Glacier Deep Archive**	Yes	No
Single File Restore	Yes	**S3 Cross-Region Replication**	Yes	Yes

[1] Via Microsoft-supplied VSS Writers in Windows Guests running VMware Tools. Via vmsync for Linux.

[2] Via Microsoft-supplied VSS Provider in Windows Guests running VMware Tools.

[3] Source-side in this context refers to the Cohesity appliance, not the vSphere hosts

Table 2 - Cohesity DataPlatform 6.4 Feature Support

Further Information

For additional information, see:

- Datasheet – Cohesity and VMware[4]
- Datasheet – Cohesity on AWS[5]
- Datasheet – Cohesity DataPlatform[6]

Rubrik Cloud Data Management

Overview

Rubrik Cloud Data Management (CDM) for VMware is based upon a hyperconverged architecture. Each appliance (or Brik) in a cluster contributes locally-attached storage to a highly available, dynamically optimized virtual storage pool.

The direct-attached storage in each node is either physical HDD / SDD drives, or a combination of Amazon EC2 instance stores and Amazon EBS volumes.

The CDM platform offers a variety of storage-related services. This section is focused on using CDM to conduct VADP-style backups of on-premises vSphere clusters to Amazon S3.

There are two supported approaches relevant to this use case:

- **Amazon S3 as an archive repository** - The most recent backup data (typically 30 days) is kept locally on the cluster. Once it has aged past that, it is then archived to one of the Amazon S3 storage classes for long-term retention.
- **Direct backups to Amazon S3** – The Instant Archive feature allows a cluster to act as a Cloud Repository Proxy. VM backup data is immediately transmitted to Amazon S3, with the local disk repository acting as a cache.

Note: Archived backup data is fully self-describing. An Amazon S3 bucket containing a full chain and a Rubrik cluster that can read it is all a customer needs to conduct restore operations.

Logical Architecture

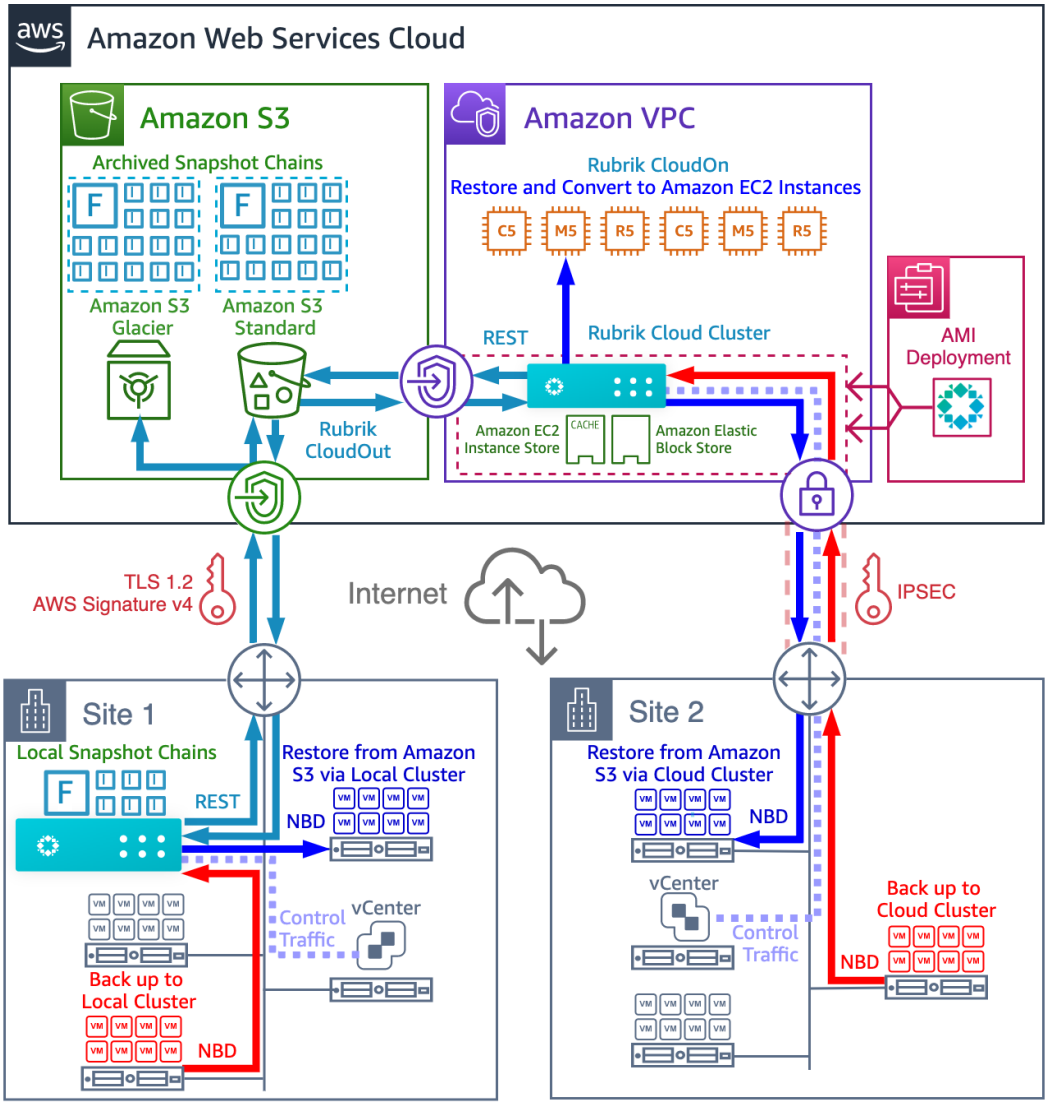

Figure 4 - Rubrik Cloud Data Management for VMware Logical Architecture

Illustrated Workflows

- Back up of VMs to Rubrik Cluster in Site 1
- Back up of VMs in Site 2 to a Rubrik Cloud Cluster in Amazon EC2

> **Note:** Deployment of a Cloud Cluster is done by contacting Rubrik support, who will privately share an AMI with the customer's account.

- Archive old backups from Rubrik Cluster in Site 1 to Amazon S3

- Archive old backups from Rubrik Cloud Cluster in Amazon EC2 to Amazon S3
- Recovery of VMs from Amazon S3 to the same Rubrik Cluster in Site 1
- Recovery of VMs from Amazon S3 to Site 2 via a Rubrik Cloud Cluster
- Recovery and conversion of VMs to EC2 instances from Rubrik Cloud Cluster

Note: CDM 5.0.3 does not support the HotAdd transport mode.

Feature Support

VADP Backup Features		S3 Feature Support	Commercial	GovCloud
Application Consistency	Yes [1]	**S3 Standard**	Yes	Yes
Volume Consistency	Yes [2]	**S3 Standard-IA**	Yes	Yes
Source Side Compression	Yes [3]	**S3 One Zone-IA**	Yes	Yes
Source Side Deduplication	Yes [3]	**S3 Intelligent-Tiering**	Yes	Yes
Global Deduplication	Yes	**S3 Glacier**	Yes	Yes
Changed Block Tracking	Yes	**S3 Glacier Deep Archive**	No	No
Single File Restore	Yes	**S3 Cross-Region Replication**	No	No

[1] MSSQL, Exchange, Sharepoint, AD, and Oracle via Rubrik Adaptive Consistency, or using standard VSS providers

[2] Rubrik Adaptive Consistency available for both Windows and Linux guests, or standard VSS provider in Windows

[3] Source-side in this context refers to the Rubrik appliance, not the vSphere hosts

Table 3 – Rubrik Cloud Data Management 5.0.3 Feature Support

Further Information

For additional information, see:

- Datasheet – Rubrik and VMware[7]

- Datasheet – Rubrik Cloud Solutions[8]
- Datasheet – Archival Across Clouds[9]

Veeam Backup and Recovery

Overview

Veeam Backup and Replication is a Windows-based distributed data protection solution. While Veeam has evolved the capability to protect a wide variety of resources, it was originally designed to back up VMs running on ESX using snapshots. It is one of the most mature solutions in this respect.

There are three supported approaches relevant to this use case:

- **Cloud Repository Proxy as primary backup target** – Using an AWS Storage Gateway appliance (either hardware or virtual) acting as a Virtual Tape Library (VTL). This is the only method that supports Amazon S3 Glacier.
- **Cloud Repository as archive tier** – This method uses a Veeam Object Repository as well, but backup jobs do not directly stream to it. Instead, a disk repository is used as the primary backup target. Data is then migrated over time to Amazon S3 according to criteria specified in Veeam archive policies.

Logical Architecture

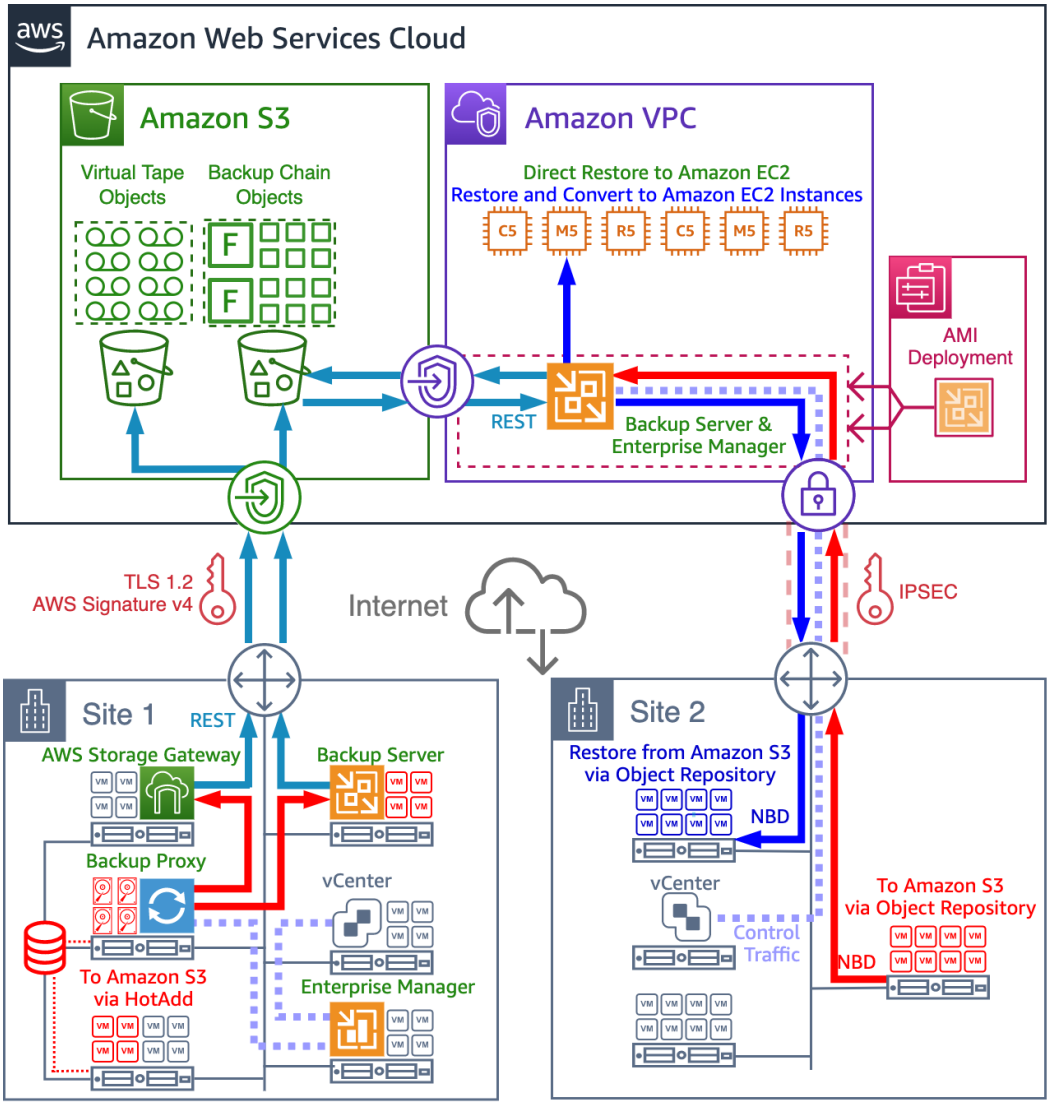

Figure 5 – Veeam Backup and Recovery Logical Architecture

Illustrated Workflows

- Back up of VMs in Site 1 to Amazon S3 via AWS Storage Gateway in VTL mode using HotAdd transport
- Back up of VMs in Site 1 to Amazon S3 via Object Repository using HotAdd transport
- Back up of VMs in Site 2 to Amazon S3 via Object Repository on cloud-based Backup Server using NBD transport

- Recovery of Site 1 VMs that were backed up via Object Repository from Amazon S3 to Site 2 via NBD transport mode
- Recovery of Site 1 VMs that were backed up via Object Repository from Amazon S3 to Amazon EC2 on cloud-based Backup Server using Direct Restore to Amazon EC2 feature in Veeam Backup & Replication 9.5 u4a

Feature Support

VADP Backup Features		S3 Feature Support	Commercial	GovCloud
Application Consistency	Yes [1]	**S3 Standard**	Yes	Yes
Volume Consistency	Yes [2]	**S3 Standard-IA**	Yes	Yes
Source Side Compression	Yes [3]	**S3 One Zone-IA**	No	No
Source Side Deduplication	Yes [3]	**S3 Intelligent-Tiering**	Yes [4]	Yes [4]
Global Deduplication	No	**S3 Glacier**	Yes [4]	Yes [4]
Changed Block Tracking	Yes	**S3 Glacier Deep Archive**	Yes [4]	Yes [4]
Single File Restore	Yes	**S3 Cross-Region Replication**	No	No

[1] For registered VSS Writers in Windows Guests running VMware Tools

[2] For Windows Guests running VMware Tools

[3] Source-side in this context refers to the HotAdd Backup Proxy appliance, not the vSphere hosts

[4] Supported when using AWS Storage Gateway as a Cloud Repository Proxy only

Table 4 – Veeam Backup and Replication 9.5 u4a+ Feature Support

Further Information

For additional information, see:

- Deployment Guide – Veeam using AWS Storage Gateway[10]
- Blog – Ultimate FAQ for Scale-Out Backup Repository[11]

Contributors

Contributors to this document include:

- Sean Howard – Partner Solutions Architect, AWS
- Henry Axelrod – Sr. Partner Solutions Architect, AWS

Document Revisions

Date	Description
December 2019	First publication

Notes

[1] https://aws.amazon.com/s3/storage-classes/

[2] https://aws.amazon.com/blogs/storage/backing-up-vmware-cloud-on-aws-with-druva-phoenix/

[3] https://resources.druva.com/white-papers/cloud-native-data-protection-for-vmware

[4] https://www.cohesity.com/resource-assets/solution-brief/Cohesity-Maximize-Your-VMware-Environment-Solution-Brief.pdf

[5] https://www.cohesity.com/resource-assets/solution-brief/Cohesity-and-AWS-Solution-Brief.pdf

[6] https://www.cohesity.com/resource-assets/datasheets/Cohesity-DataPlatform-Datasheet.pdf

[7] https://www.rubrik.com/wp-content/uploads/2016/12/Joint-Solution-Brief-VMware-and-Rubrik.pdf

[8] https://www.rubrik.com/wp-content/uploads/2017/06/DATA-SHEET-Rubrik-Cloud-Solutions.pdf

[9] https://www.rubrik.com/wp-content/uploads/2017/06/DATA-SHEET-Rubrik-for-Data-Archival.pdf

[10] https://www.veeam.com/using-aws-vtl-gateway-deployment-guide_wpp.pdf

[11] https://www.veeam.com/blog/ultimate-faq-for-scale-out-backup-repository.html