

AWS Whitepaper

VMware vSphere Backups to Amazon S3



VMware vSphere Backups to Amazon S3: AWS Whitepaper

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Abstract and introduction	i
Introduction	1
Are you Well-Architected?	1
AWS Backup	3
Partner solutions	5
Druva Phoenix	5
Logical architecture	6
Feature support for Druva Phoenix for VMware 4.8	7
Further information	8
Cohesity DataPlatform	8
Logical architecture	9
Cohesity DataPlatform 6.4 feature support	10
Further information	12
Rubrik Cloud Data Management	12
Logical architecture	13
Feature support	14
Further information	15
Veeam Backup and Recovery	15
Logical architecture	16
Veeam Backup and Replication 9.5 u4a (and later) feature support	17
Further information	18
Appendix - Terms and definitions	19
vSphere storage APIs - Data protection	19
Transport modes	19
Changed block tracking	21
Data protection solution components	21
Distributed data protection solutions	21
Hyperconverged data protection solutions	23
SaaS-based data protection solutions	24
Important features	24
Contributors	26
Document history	27
Notices	28
AWS Glossary	29

VMware vSphere Backups to Amazon S3

Publication date: **June 15, 2023** ([Document history](#))

This paper acts as an introduction to the features, Amazon Web Services (AWS) services, and architectural components relevant when you back up an on-premises virtualization environment based on VMware vSphere to Amazon Simple Storage Service (Amazon S3). It shows examples that use products and features from AWS Backup and AWS Partner Network (APN) partners Veeam, Rubrik, Cohesity, and Druva to show you common approaches taken by AWS or APN Technology Partners. This paper is intended for solution architects familiar with VMware virtualization. You should have a knowledge level equivalent to a VMware Certified Professional in Data Center Virtualization for vSphere version (VCP-DCV).

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 time since VADP's release, many customers have benefitted from the use of this interface; it has become fundamental for VMware specialists when they address the recoverability needs of a design. At the same time, even more customers want to take advantage of the benefits of storing their backup data in the cloud.

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

For a list of terms and definitions used in this paper, refer to [???](#).

Are you Well-Architected?

The [AWS Well-Architected Framework](#) helps you understand the pros and cons of the decisions you make when building systems in the cloud. The six pillars of the Framework allow you to learn architectural best practices for designing and operating reliable, secure, efficient, cost-effective, and sustainable systems. Using the [AWS Well-Architected Tool](#), available at no charge in the [AWS Management Console](#), you can review your workloads against these best practices by answering a set of questions for each pillar.

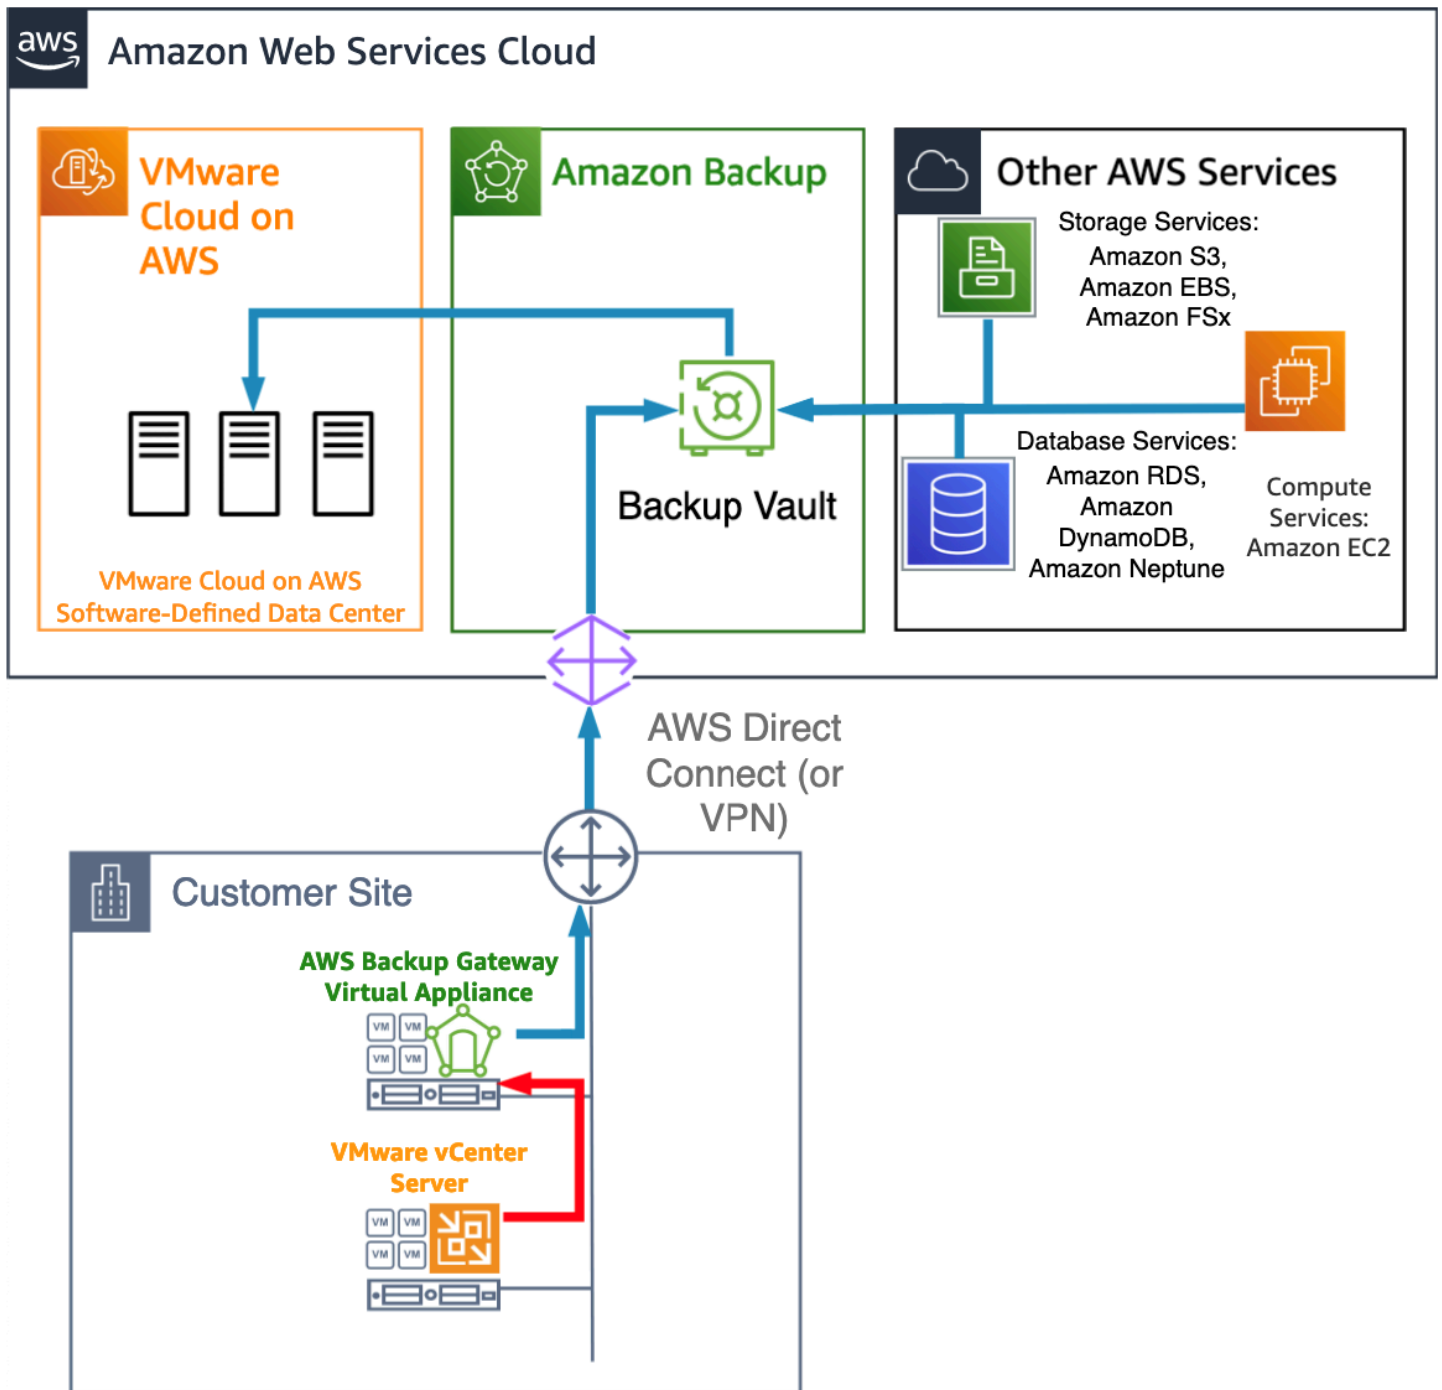
For more expert guidance and best practices for your cloud architecture—reference architecture deployments, diagrams, and whitepapers—refer to the [AWS Architecture Center](#).

AWS Backup

This section describes the architecture of using AWS Backup to protect your on-premises VMware environments. Although this is not being backed up onto a customer owned and managed Amazon S3 bucket, it is backed up onto an AWS Backup data vault, which offers the same levels of durability and service availability as Amazon Simple Storage Service (Amazon S3). Backups created by AWS Backup are also replicated into three Availability Zones within the Region and can be replicated to other AWS Regions.

Consider using AWS Backup to protect on-premises VMware virtual machines if you have already deployed AWS services that are protected with AWS Backup and want to maintain a single utility for protection. This minimizes the administrative overhead of managing multiple backup utilities across your on-premises and cloud environments.

Another benefit of using AWS Backup for on-premises VMware virtual machines is that you can recover them into VMware Cloud on AWS. This gives organizations the ability to validate their backups in an isolated environment, which might be necessary for compliance and can also provide a low-cost option for disaster recovery.



VMware Cloud on AWS

As of November 2022, it is also possible to restore VMware virtual machines to Amazon Elastic Compute Cloud (Amazon EC2). More details can be found on the [AWS Backup now supports restore of VMware workloads to Amazon EC2](#) blog post.

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.

Topics

- [Druva Phoenix](#)
- [Cohesity DataPlatform](#)
- [Rubrik Cloud Data Management](#)
- [Veeam Backup and Recovery](#)

Druva Phoenix

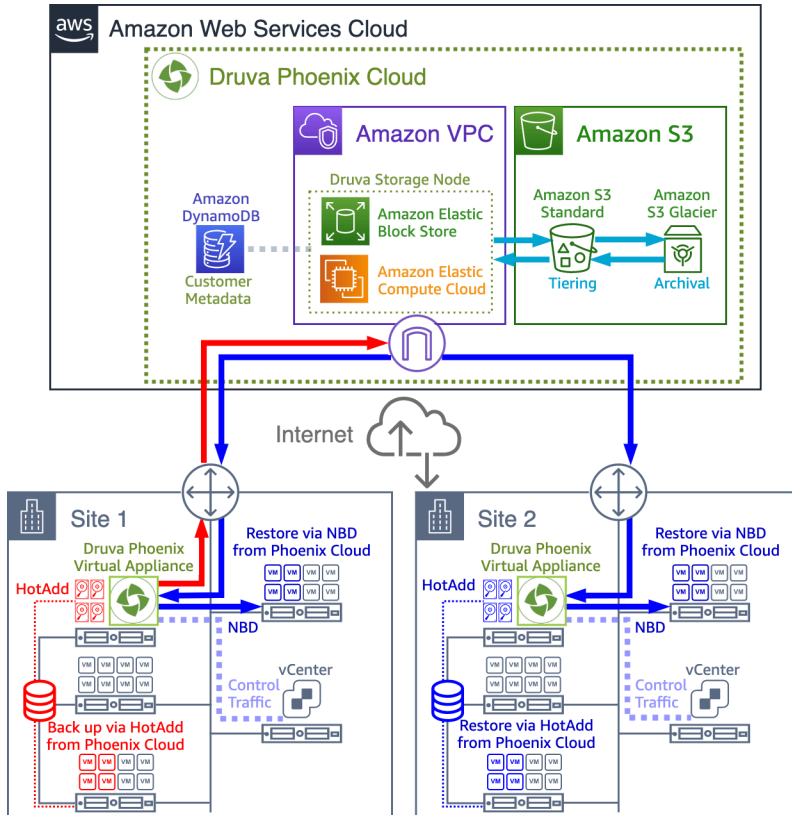
Phoenix for VMware is a hybrid backup solution delivered by a SaaS model. Components of a Phoenix for VMware deployment include:

- **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. CloudCache is installed by 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 can be deployed depending on the number of protected VMs, throughput, and RTO requirements.

Logical architecture



Druva Phoenix for VMware logical architecture

Illustrated workflows

- Backup of VMs to Phoenix Cloud by backup proxy using HotAdd transport
- Tiering down of hot data from storage nodes to warm tier on Amazon S3
- Archival of older backup data from Amazon S3 to 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

Feature support for Druva Phoenix for VMware 4.8

1 Supported

1.1 Backups

1.2 Support

2 Application

2.1 S3

2.2 Consistent

2.3 cy

3 Amazon

3.1 Consistent

3.2 Standard-

3.3 IA

4 Amazon

4.1 Side

4.2 Compressi

4.3 on-

4.4 IA

5 Amazon

5.1 Side

5.2 Deduplica

5.3 tion

5.4 Tierin

5.5 g

6 Global

6.1 Replica

6.2 tion

7 Changed

7.1 Blocker

7.2 Deeping

7.3 Archive

**VMware
Backups
Single
File
Restore
Reg
ion
Replicati
on**

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

Further information

For additional information, refer to:

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

Cohesity DataPlatform

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 Elastic Block Store (Amazon EBS) volumes.

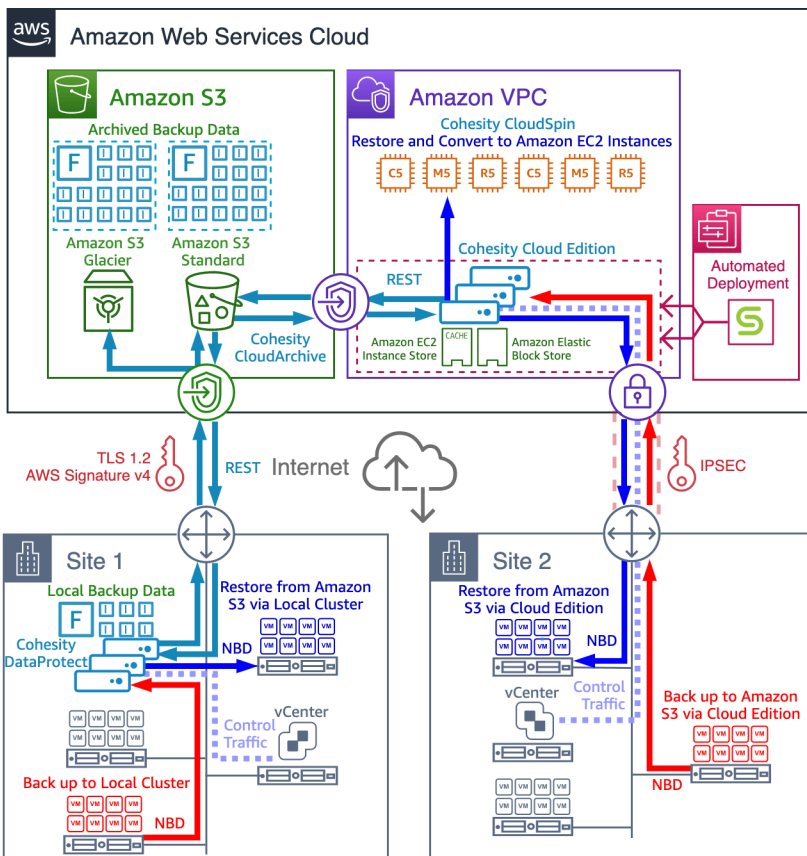
DataPlatform offers a variety of storage-related services. This section is focused on how to use it to conduct VADP-style backups of on-premises vSphere clusters to Amazon S3.

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



Cohesity DataPlatform logical architecture

Illustrated workflows

- Backup of VMs to local DataPlatform cluster in Site 1
- Backup 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.

- Archival of old backups from DataPlatform Cluster in Site 1 to Amazon S3
- Archival of old backups from DataPlatform Cloud Edition cluster to Amazon S3
- Archival by AWS Lifecycle policy from Amazon S3 to 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.

Cohesity DataPlatform 6.4 feature support

VMware vSphere

Backups

Support

Application

S3

Consistent

cy

VMware vSphere

Consistent

Standard-

IA

VMware
Backups
Backups

Amazon
S3
Compressi
Zone-
IA

Amazon
S3
Deduplica
tion
Tierin
g

Global
Deduplica
tion

Changed
Block
Tracking
Archive

Amazon
S3
Restore
Reg
ion
Replicati
on

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.

Further information

For additional information, refer to:

- [Datasheet – Cohesity and VMware](#)
- [Datasheet – Cohesity on AWS](#)
- [Datasheet – Cohesity DataPlatform](#)

Rubrik Cloud Data Management

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 / SSD 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.
- **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



Rubrik Cloud Data Management for VMware logical architecture

Illustrated workflows

- Backup of VMs to Rubrik Cluster in Site 1
- Backup 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.

- Archival of old backups from Rubrik Cluster in Site 1 to Amazon S3
- Archival of 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 Amazon EC2 instances from Rubrik Cloud Cluster

Note

5.0.3 does not support the HotAdd transport mode.

Feature support

Application

Backups

Support

Application

63

Consistent

cy

Version

Consistent

Standard-

IA

Version

Size

Compression

Zone-

IA

Version

Size

Replication

tion

Tiering

g

Global

Replication

tion

VMware
Backup
Agent
for
Linux
and
Windows
Guests
on
vSphere
ESX/ESXi
Hosts
with
Amazon
S3
Storage

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

Further information

For additional information, refer to:

- [Datasheet – Rubrik and VMware](#)
- [Datasheet – Rubrik Cloud Solutions](#)
- [Datasheet – Archival Across Clouds](#)

Veeam Backup and Recovery

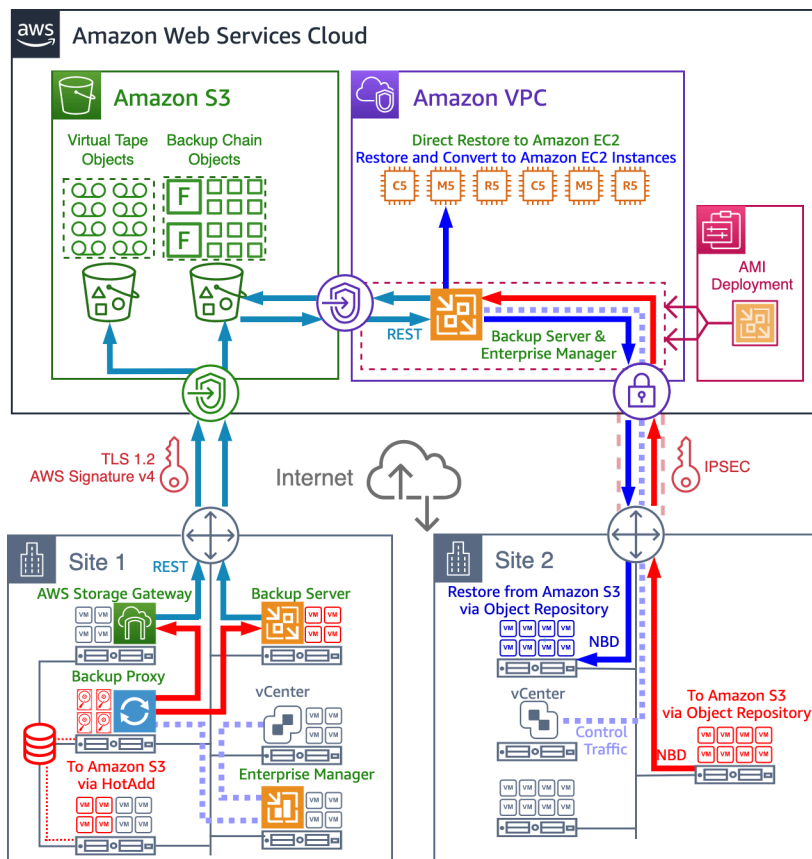
Veeam Backup and Replication is a Windows-based distributed data protection solution. While Veeam has evolved the capability to protect a wide variety of resources, it was originally designed

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 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



Veeam Backup and Recovery logical architecture

Illustrated workflows

- Backup of VMs in Site 1 to Amazon S3 via AWS Storage Gateway in VTL mode using HotAdd transport
- Backup of VMs in Site 1 to Amazon S3 via Object Repository using HotAdd transport
- Backup 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 and Replication 9.5 u4a

Veeam Backup and Replication 9.5 u4a (and later) feature support

Application

Backup

Support

Application

S3

Consistent

cy

Amazon

Consistent

Standard-

IA

Amazon

Site

Compressi

on-

IA

Amazon

Site

Intellige

VMware Tools**Backup****Agent****Deduplica****tion****g****Global****Deduplica****tion****VMware****Backup****Agent****Deeping****Single****File****Restore****Reg****ion****Replicati****on**

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

Further information

For additional information, refer to:

- [Deployment Guide – Veeam using AWS Storage Gateway](#)
- [Configure your infrastructure to use Backup gateway](#)

Appendix - Terms and definitions

Review the terms used in this whitepaper.

Topics

- [vSphere storage APIs - Data protection](#)
- [Data protection solution components](#)

vSphere storage APIs - Data protection

Transport modes

Any data protection solution that uses 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.

- **Storage area network (SAN) transport mode** - A data mover server mounts the running shared storage volume as a logical unit number (LUN). This mode is only relevant in an iSCSI or Fibre Channel based SAN environment. VM files are backed up by the data mover over the SAN fabric directly.
- **Network block device (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.

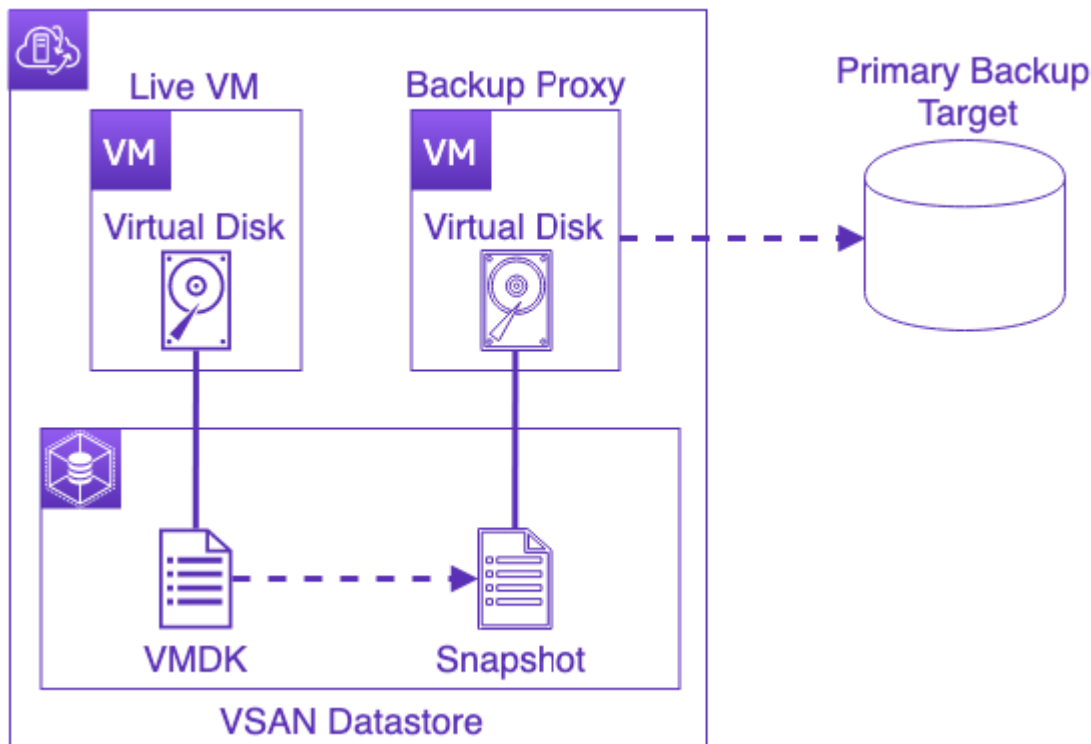
The NBD transport modes use this management interface to copy the VM disk files across the network to the backup target. If more than one VMkernel interface is enabled for management traffic on the host, it will use the one that has a preferred route to the target IP address.

- **NBD secure sockets layer (NBDSSL) transport mode** - NBDSSL is simply an SSL-encrypted version of NBD. It is the default used by ESXi 6.5 and later.

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 uses 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:
 1. Snapshots of the live VM's virtual disks are taken and mounted to the backup proxy as if they were normal virtual disks.
 2. Depending on the vendor implementation, VSS or other freeze/thaw mechanisms might be invoked by using VMware tools to quiesce at the application or volume level when the snapshot is created.
 3. 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.
 4. When the backup is complete, the snapshot is deleted.



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

Changed Block Tracking (CBT) is an optional feature offered by VADP that helps 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.

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 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 include:

- An 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 (such as a disk or cloud repository).
- Deduplication or 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 you use the HotAdd transport. It is a VM that resides within the protected vSphere cluster and 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 eight 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.

Each Amazon S3 storage class has a distinct set of properties meant to optimize cost for a given access pattern 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-Azure 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:

- VMs in a local vSphere environment running Windows or Linux with a large virtual disk backed by a LUN on a SAN array. They can be self-built or virtual appliances.
- 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

The following features are available when using a third-party backup vendor with Amazon S3.

- **Client-side data efficiency** - Data efficiency mechanisms such as deduplication or compression 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 or compressed (MPEG)
- Solution-specific details such as fixed-length or 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, 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 might 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.

Contributors

Contributors to this document include:

- Sean Howard – Sr. Solutions Architect, Amazon Web Services
- Henry Axelrod – Principal Solutions Architect, Amazon Web Services
- Sam Walker – Sr. Partner Solutions Architect, Amazon Web Services

Document history

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Whitepaper updated	Updates to include AWS Backup	June 15, 2023
Initial release	Initial release of the AWS Whitepapers User Guide	December 1, 2019

Note

To subscribe to RSS updates, you must have an RSS plugin enabled for the browser that you are using.

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.

© 2023 Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.