

Public Cloud Architecture guide for Microsoft Azure

Feature release 11.25

September 2021

Table of Contents

- Abstract5**
- Commvault platform design principles for cloud6**
 - Native cloud connectivity 6
 - Scalability 6
 - Design for recovery 7
 - Crash consistency versus application consistency 8
 - Storage-level replication versus discrete independent copies 8
 - Deciding what to protect 8
 - Designed for cloud efficiency 9
 - Cloud power management 9
- Automation 9**
 - Programmatic data management 9
 - Workload auto-detection and auto-protection 10
 - Self-service access and restore 10
- Most common cloud use cases with Commvault software 11**
 - Move data – backup and archive to the cloud 11
 - Move data – Migration of VMs and applications to the cloud 13
 - Manage data – Protection in the cloud 14
 - Use data – Disaster recovery to the cloud 15
- Architecture considerations 16**
 - Networking 16
 - Virtual private cloud/networking 16
 - Bridging on-premises infrastructure – VPN and ExpressRoute 17
 - Firewall ports 17
 - Azure Virtual Network 18
 - Data security 19
 - In-flight 19
 - At-rest 19
 - HTTPS proxies 19
 - Account separation 19
 - Azure Private Link 20



Data seeding.....	20
“Over-the-wire”	21
Drive seeding.....	21
Cost/Consumption.....	22
Network Egress	22
Storage I/O	22
GET/PUT transaction costs	22
Data recall	22
Performance/Storage	23
Multi-streaming with object storage	23
Cloud connector best practices.....	23
Compression vs. deduplication	24
Leveraging multiple mount paths for a cloud library	24
Block blob versus Page blob object storage	24
Partitioned deduplication	24
Micro pruning.....	24
Selecting the right storage class for backup and archive data	25
Azure storage redundancy.....	25
Choosing the correct Azure storage tier.....	26
Azure Immutable Storage (WORM)	30
Storage Accelerator to Azure Blob Storage	32
Performing disaster recovery to the cloud	33
Restoring applications (automated or on-demand)	33
VM Replication (Live Sync).....	33
Replicating other workloads.....	34
Virtual machine recovery from Amazon EC2 to Azure VM.....	35
Azure-specific workloads.....	35
Virtual machine recovery into Azure VM instances	35
Using Commvault Workflows to automate Disaster Recovery	35
Protecting and recovering active workloads in Azure.....	36
Agent-less VM protection (Virtual Server Agent for Azure)	36
Performance tests in Azure using the VSA for Azure.....	40
Agent-in-guest (streaming)	41
Azure snapshots.....	42

Azure blob storage backup	42
Application migration	43
Virtual Machine restore & convert (lift and shift to Azure)	44
Application out-of-place restore (all supported platforms)	44
Deployment	44
Installation basics	44
CommServe® Disaster Recovery solution comparison	44
Pre-packaging Commvault software within a VM template	44
Automating deployment with continuous delivery	45
Cloud library configuration	45
Unsupported cloud storage configurations	46
Architecture sizing.....	46
Azure CommServe® server specifications	46
Azure MediaAgent specification.....	47
Virtual Server Agent specification	51
Example of sizing in cloud	52
Assumptions & costing	52
Additional resources	53
Documentation	53
Online documentation – Cloud Storage	53
Datasheets	53
Solution briefs.....	53
Videos	53
Case studies.....	53



Notices

This document is provided for informational purposes only. It represents Commvault's current product offerings and practices as of the date of issue of this document, of which are subject to change without notice. The responsibilities and liabilities of Commvault to its customers are controlled by Commvault agreements, and this document is not part of, nor does it modify, any agreement between Commvault and its customers.

New in this version

Version	Date	Changes
11.19	April 2020	<ul style="list-style-type: none">• Updated document with SP17-SP19 functionality
11.25	Sept 2021	<ul style="list-style-type: none">• Updated document with 11.20 to 11.25 functionality• Update reference diagram• Azure blob storage metadata backup• Azure immutable storage• Azure Generation 2 VM• VSA Access Node Auto-Scale• Availability Zones support• VM Centric operations• Owner detection• Data verification recommendation

Commvault remains committed to ensuring the **Cloud Architecture Guide** remains relevant to currently available public cloud capabilities.

The latest copy of this document is available at [Virtualization White Papers >](#)

Abstract

This document serves as an architecture guide for solutions architects and Commvault customers who are building data protection and management solutions utilizing public cloud environments and Commvault software.

It includes public cloud concepts, architectural considerations, and sizing recommendations to support Commvault software in public cloud. The approach defined in this guide applies to both running Commvault solely in public cloud environments and extending existing on-premises Commvault functionality into hybrid cloud architectures. The guide covers several common use cases for public cloud including moving data to public cloud, disaster recovery to public cloud, as well as protecting workloads in public cloud.

Currently this guide delivers architecture considerations and sizing recommendations for the Microsoft Azure public cloud platform. Guides for other public cloud environments are available as well.

Commvault platform design principles for cloud

In this section, we provide design principles and architecture principles that have been employed within the Commvault platform to provide an optimal cloud experience for organizations planning to leverage the cloud as part of their data protection and management strategy.

Native cloud connectivity

The Cloud Library feature of a MediaAgent is the native integration within the Commvault platform that directly communicates with object storage such as Azure Blob Storage - Hot, Cool, and Archive storage tiers. Commvault does not require translation devices, gateways, hardware appliances or Virtual Tape Libraries (VTLs). For more details on Azure storage classes see [Azure Blob storage >](#)

This Cloud Library works by communicating directly with the object storage REST API interface over HTTP or HTTPS, allowing for Commvault platform deployments on both virtual and physical compute layers to perform read/write operations directly against cloud storage targets, reducing the TCO of the data management solution. The Cloud Library is part of the native code of the Commvault platform and optimizes the data exchange with cloud object storage platform to maximize the transfer speed while minimizing recall needs and costs.

Since the Cloud Library essentially treats cloud storage akin to a disk target, data management functions such as compression, encryption, deduplication, and data life-cycling can be performed against cloud storage targets to ensure that both costs and risks are managed effectively. This also allows the data to be retained independent to the cloud format thereby enabling optimized recall and movement of data across different cloud platforms for future use cases.

For more information on all the supported vendors, please refer to this comprehensive list located at [Supported Cloud Storage Products >](#)

Scalability

Application environments and the data and VMs that service those environments grow over time, and a data protection and management solution needs to adapt with the change rate to protect the dataset quickly and efficiently, while maintaining an economy of scale that continues to generate business value out of that system.

Commvault addresses scalability in cloud architectures by providing these key constructs:

Deduplication building blocks

Commvault software maintains a scale-up or scale-out “building block” approach for protecting datasets, regardless of the origin or type of data. These blocks are sized based on the front-end data they will ingest, prior to compression and deduplication. This provides clear scale-out and scale-up guideline for the capabilities and requirements for each Commvault MediaAgent that will perform the data movement (both ingestion and storage), compression and deduplication.

Furthermore, these deduplication MediaAgent building blocks may be logically grouped together in a grid formation, providing further global deduplication scale, load balancing, and redundancy across all nodes within the grid.

This software on architecture, with scale-up and scale-out enables cloud adoption to start with a cost-conscious approach however scales to meet SLAs quickly without locking the architecture into a specific unit of operation.

Client-side deduplication

As is the nature of deduplication operations, each data block must be hashed to determine if it is a duplicate block, or unique; and then must be captured. While this is a way to improve the ingest performance of the data mover (MediaAgent), it has the secondary effect of reducing the network traffic stemming from each client communicating through to the data mover.

In public cloud environments where network performance can vary, the use of client-side deduplication can reduce backup windows and drive higher scale, freeing up bandwidth for both production and backup network traffic. By utilizing client-side deduplication, the workload of backup can be distributed across all the VMs, compared to building a larger data protection architecture in cloud. This can also help reduce the recovery points for critical application by enabling more frequency of protection.

VSA Access Node (Proxy) for cloud platforms

Utilizing agents in each cloud operating VM is an approach that distributes the overall workload and cost for data protection across all the VMs. However, in many cases with large scale deployments, management of each VM can become an overhead. The Commvault platform automates the management of such operations from initial deployment to upgrading and disposition. When this approach is deemed insufficient, the Commvault Virtual Server Agent (VSA) software capability can be loaded into a public cloud VM to perform agent-less operations.

Akin to proxy-based protection for on-premises hypervisors, the Commvault VSA* access node, interfaces directly with APIs available from the hypervisor layer of public cloud platforms to perform protection and management operations of VMs within the public cloud platform. The VSA not only manages operations such as snapshot creation and orchestration but can also perform automatic VM identification and selective VM data reads (Change Block Tracking) from cloud platforms that support this capability. The VSA further performs any data format conversions and enables disaster recovery operations for VMs to cloud platforms. Working together with the MediaAgent (data mover), the VSA offers enhanced protection and management of cloud workloads.

*For a complete list of supported VSAs and updated VSA capabilities please review the online

[VSA Feature Comparison Matrix >](#)

Commvault recommends installation of the Virtual Server Agent package on your MediaAgent initially, to provide a single consolidated Access Node for protection and data management. As your cloud protection needs grow, you may look to utilize [Access Node Auto Scaling](#) to automatically create and terminate Access Nodes on-demand.

Design for recovery

Using native cloud provider tools, such as creating a snapshot of a cloud-based VM may be easy to orchestrate but does not always deliver the application-consistency or mobility required by an application or database such as Microsoft SQL Server or Oracle Database residing within the VM. The general approach requires database and application-specific scripting or manual handling to deliver a successful application recovery. Across a large enterprise estate, this bespoke manual management becomes time-consuming and subject to human error.

As part of any data protection and management solution, it is important to ensure that you design for recovery in order to maintain and honor the recovery time objective (RTO) and recovery point objective (RPO) requirements identified for your individual applications groups.

Crash consistency versus application consistency

While crash-consistency within a recovery point may be sufficient for a file-based dataset or cloud VMs such as Azure VMs it is not generally appropriate for an application such as Microsoft SQL Server or Oracle Database where the database instance needs to be quiesced to ensure the database is valid at the time of backup.

Commvault software supports **both** crash and application consistent backups, providing flexibility in your design while assuring VM recoverability coupled with application recovery to a specific point in time. Not only are the most common types of applications covered, but a wide variety of classic applications and cloud applications are supported. For a complete list of updated application support please review the [online documentation](#) for streaming application agent support and [online matrix](#) for application aware VSA support.

Storage-level replication versus discrete independent copies

Azure supports replication at the object storage layer from one region to another with GRS (Geo-Redundant) or RA-GRS (Read Access Geo-Redundant) storage, however, in the circumstance that corrupted blocks are replicated to the secondary region, your recovery points are invalid. Further network and storage costs continue to accumulate regardless of the validity of both “sides” of the data.

While Commvault software can support a replicated cloud library model, in which the secondary storage location for backups is replicated using the Cloud vendors storage-based replication tools, we recommend that you consider Commvault software to create an independent copy of your data, either to another region, or different cloud provider, or back to an on-premises infrastructure to address broader risks. Deduplication is also vital as part of the latter option and this ensures that Commvault software can minimize the cross-region and cross-provider copy time and costs by ensuring only the unique changed blocks are transferred over the network.

This recommendation not only ensures recoverability to multiple points in time, it further manages the cost and risk through the assurance that the data is independent of the platform and ensures that different SLAs for protection and retention can be maintained for different classes of data.

Deciding what to protect

Not all workloads within the cloud need protection – for example, with micro services architectures, or any architecture that involves worker nodes that write out the valued data to an alternate location, presents no value in protecting the worker nodes. Instead, the protection of the gold images and the output of those nodes provides the best value for the business. However, it is important to note that data stored in ephemeral locations may need to be protected prior to termination operations against those VMs to ensure that any valuable data is not lost. Log files, which can be exported to Azure Blob storage, are another example of what may require protecting instead of the entire service. Protect the data where it resides instead of the service. The use of Commvault Cloud Apps for [Azure Blob Storage](#) is an efficient method for this use case.

Designed for cloud efficiency

As already discussed, the ability to provide compression and deduplication for both data on and data in the cloud begins to provide initial cost savings for many of the common use cases for secondary data. However, deduplication savings are closely tied to the type of data being managed and additional methods can result in even more overall cloud efficiency.

A common consideration is to utilize multiple tiers of storage for data as the service life of that data reduces. This has been a common practice on-premises and the Commvault platform extends this capability to cloud platforms. By having not only native integration to primary object storage targets such as Azure Blob Hot storage tier, but also having native access to more cost-effective tiers such as Azure Blob Cool storage tier and Azure Blob Archive storage tier, data lifecycle management can be performed within the cloud. For example, it is not uncommon to see Azure Hot storage tier being used as the primary or secondary copy for short-term retention followed by Cool storage tier or Archive storage tier. Having a data management platform that can utilize SLA policies to orchestrate the data movement and be aware of the location of data for recall and disposition becomes a valuable quality in gaining cloud efficiency.

Cloud power management

Shutdown of VMs in an on-premises data center is a very uncommon practice or design concept, however in cloud environments this type of operation is welcomed by those paying the cloud bills. By having the ability to create policies which monitor resource usage of cloud VMs and can both alert and act by terminating such VMs. However, the risk of data loss is mitigated since it is ensured a copy of ephemeral data is protected before such an operation is performed.

The ability to shutdown VMs is extended to the Commvault platform components running in the public cloud. Referring to the MediaAgent (data movers) referenced above, these VMs can be both shutdown and powered up based around data protection operations. For example, shutting down the Cloud MediaAgent after all protection operations have ceased and restarting prior to the next SLA window can help further reduce operational costs within public cloud environments. When cloud power management is enabled in Azure, the MediaAgent instance is shut down and put in a “stopped (deallocated)” state to keep costs at a minimum.

Automation

The cloud encourages automation, not just because the infrastructure is programmable, but the benefits in having repeatable actions reduces operational overheads, bolsters resilience through known good configurations and allows for greater levels of scale. Commvault software provides this capability through three key tenets:

Programmatic data management

Commvault software provides a robust Application Programming Interface (API) that allows for automated control over deployment, configuration, and backup and restore activities within the solution.

Whether you are designing a continuous delivery model that requires automated deployment of applications, or automating the refresh of a disaster recovery copy, data warehouse or development/testing environment that leverages data from a protection copy, Commvault software provides the controls necessary to reduce administrative overhead and integrate with your toolset of choice.

Beyond API access, the most common use cases for data protection and management are built into the Commvault user interface. Simply enter the cloud credentials and necessary permission and the Commvault platform will query the cloud environment accounts and present wizards with necessary attributes to create VMs and populate with the data required to

support the above uses discussed. Since format conversions are handled by the VSA, the entire operation is orchestrated even if the source of data is an on-premises hypervisor. This reduces the operational overhead and unique skillsets required to on-board cloud usage.

Workload auto-detection and auto-protection

The Commvault Intelligent Data Agents (iDA), whether via the Virtual Server Agent for the various cloud platforms, or the multitude of application and database iDAs, provide auto-detection capabilities to reduce administrative load.

Fresh instances, new volumes recently attached to cloud instances and virtual machines, or databases imported and created into a database instance are some examples of how Commvault software automatically detects new datasets for inclusion in the next data protection SLA window, all without manual intervention. Even agent-in-guest deployments can be auto-detected by Commvault software and included in the next data protection schedule through intelligent Client Computer Groups. This capability is especially valuable in the assurance of data protected in large scale cloud environments where many users can provision workloads in the cloud but may have little or no consideration for the protection of those workloads.

This auto-detection and auto-protection level removes the requirement for a backup or cloud administrator to manually update the solution to protect the newly created datasets. This results in improving your operational excellence, improving resiliency within your cloud infrastructure, and ensuring new data protected Service Level Agreements (SLAs) are maintained.

Self-service access and restore

A common task performed by system administrators is facilitating access to recovery points for end-users and application owners, shifting their attention away from other day-to-day operations and strategic projects.

The Commvault self-service interfaces empower users to access their datasets through a web-based interface, allowing security mapped access to individual files and folders within the protected dataset, freeing up administrators to work on critical tasks. Commvault's robust role-based security function provides assurance that self-servicing users have access to only their data assets, while bespoke auditory reporting capabilities capture how these users are accessing those data assets.

Most common cloud use cases with Commvault software

The most common use cases observed at most customer environments by Commvault related to cloud, fall into three categories depending on the maturity level of initiatives around cloud adoption:

- **Move data to the cloud** – typically involves using public cloud object storage as a target for backups and archive data and moving certain types of VM workload into cloud VMs.
- **Manage data in and across clouds** – protecting and life-cycling data and VMs in cloud, moving data across clouds and back to on-premises in some cases.
- **Use data in the cloud** – utilizing the data stored in public cloud for use cases such as disaster recovery, dev/test, and other production and non-production use cases.

These three primary use cases can be visualized as follows:




Move data

Seamlessly extend the datacenter to the cloud.



Manage data

Unlock IT agility with a comprehensive view of data.



Use data

Enable a more strategic, customer-focused business.

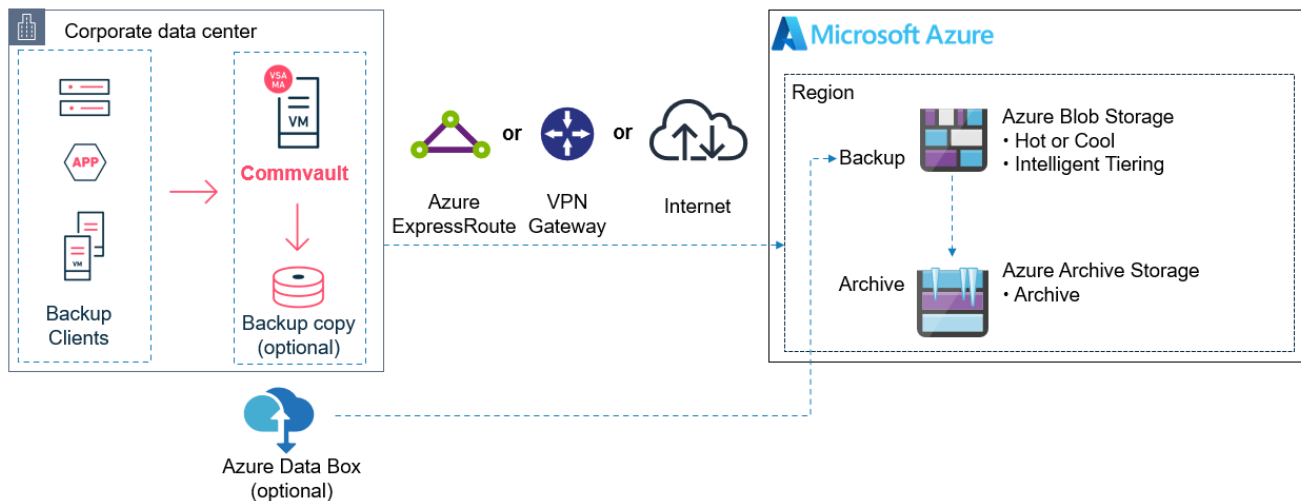
Each use case can have multiple phases and types of data associated. For example, movement could involve simple backup data, but can graduate to workloads being moved back and forth for agility as an extension to on-premises. Management of data can start with basic snapshot management and graduate to complete data lifecycle management with cloud snapshots, operational recovery deduplicated copies, and archive of data coupled with searching and indexing for compliance. The use of data can involve uses such as disaster recovery that eliminate the need to sustain secondary on-premises sites and utilize the agility of the cloud to on-ramp recovery testing and real recoveries.

Move data – backup and archive to the cloud

Business Value: Protecting data at the primary on-premises location by writing directly to an external cloud provider's storage solution or retaining a local copy and replicating the backup and archive data copy (either in full, or only selective portions of that data) into an external cloud provider's storage service suitable for both short and long-term retention configurations.

Backup and Archive

Backup and Archive data is sent online or offline, to Azure blob storage for near-term, long-term, or disaster recovery purposes



Move Data – Backup and Archive to the Cloud

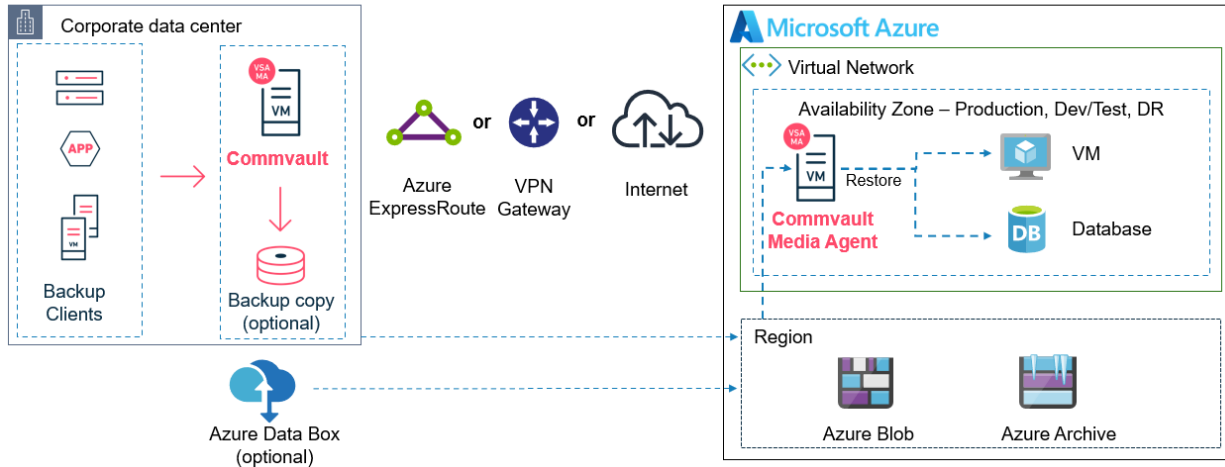
Scenario/Suitability	Requirements
<ul style="list-style-type: none"> • Offsite storage and “tape replacement” scenario – replace long-term on-site retention with cloud storage. • Native, direct connectivity to 40+ object storage endpoints, including Azure Blob Storage – no requirement for translation, gateway, or hardware deduplication devices. • Avoid point solution on a per-application basis. Any data (physical or virtual) that can be backed up by Commvault on-premises can be moved to cloud. • Cloud object storage target can be provided by either public IaaS provider (Microsoft Azure) or via a Managed Service Provider (MSP). • Local backup copy on-premises is not required. • Remote Office Branch Offices (ROBOs) can send data directly to Cloud. 	<ul style="list-style-type: none"> • Minimum 1x MediaAgent on-premises with no VMs in cloud required for backup to the cloud. • Can use direct internet connection, or a dedicated network to cloud provider for optimized data transport performance in a secure manner (e.g. Azure ExpressRoute). • Can use offline transfer method by writing backup, archive data to Azure Data Box device(s), then importing Azure Blob storage. • (optional) In-cloud MediaAgent can be created to support tiering, DR solution in the cloud using the data that is placed in the cloud. This can be done at time of DR or DR Test, as required. • Cloud MediaAgent can be deployed at the time of the Disaster Recovery event or test, as required. • Remote offices do not require onsite Commvault infrastructure, they can upload backup data direct to Azure Blob storage.

Move data – Migration of VMs and applications to the cloud

Business Value: Upon protecting VM and application data at the primary on-premises location, Commvault software orchestrates the migration of application workloads into the cloud, either at the VM container level or the application level. While providing the migration lifecycle and workloads are in a transition phase between on-premises and public cloud, data is still protected on-premises.

Migration of VMs and Applications

On-premises Virtual Machines (VMs) and application data is orchestrated into cloud-native infrastructure or application services



Move Data – Migration of VMs and application to the cloud

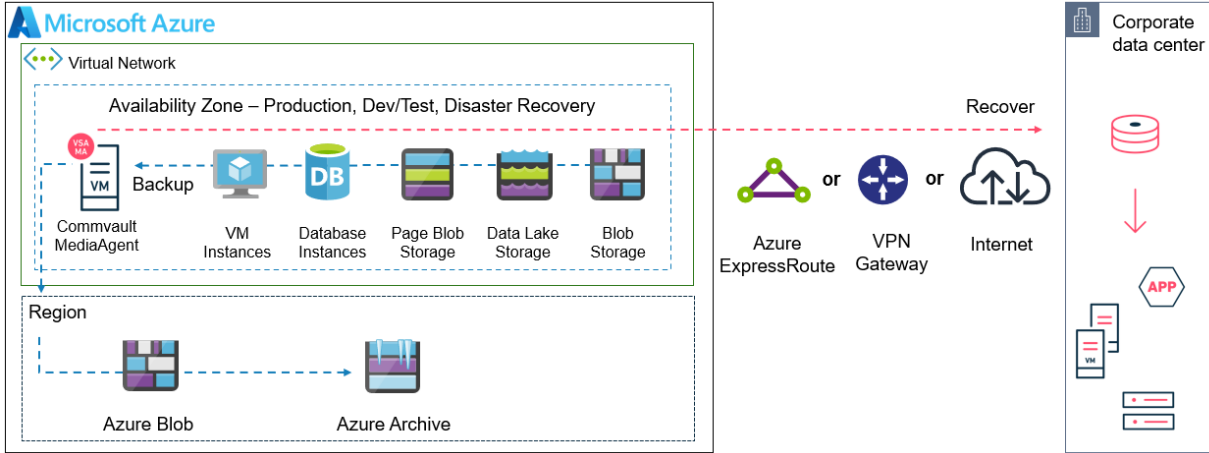
Scenario/Suitability	Requirements
<ul style="list-style-type: none"> Lift & shift of virtual machines – Application-consistent VM backups are used to restore and convert VMware, Hyper-V, Acropolis, Oracle VM, and AWS VMs into Azure as part of a migration with a phased cut-over strategy reducing on-premises downtime. Application restore out-of-place – Leverage Commvault iDataAgents for your supported workload to restore the target application out-of-place to a warm VM residing in cloud. 	<ul style="list-style-type: none"> Minimum 1x MediaAgent on-premises to protect and capture workloads Minimum 1x MediaAgent (& DDB) in cloud to protect workloads post-migration in-cloud, and for optimal migration performance. Highly recommended to use dedicated network to cloud provider for best performance (e.g. Azure ExpressRoute).

Manage data – Protection in the cloud

Business Value: Providing operational recovery for active workloads and data within an external provider’s cloud. Provide the ability to lifecycle data and cloud VMs to meet SLA and cost requirements.

Protection in the Cloud

Protect and recover cloud workloads, cloud-native applications both within cloud and back on-premises



Manage data – Protection in the cloud

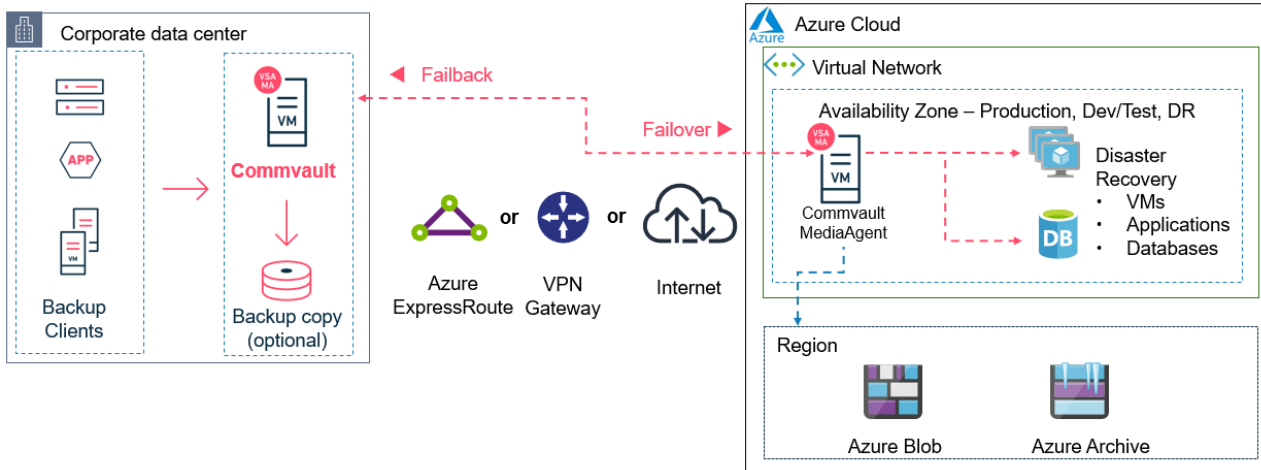
Scenario/Suitability	Requirements
<ul style="list-style-type: none"> Data protection for cloud based workloads – protecting active workloads within an existing IaaS Cloud (Production, Dev/Test, etc.). Azure agentless VM Protection – Protect VMs with an agentless and script-less protection mechanism through the Virtual Server Agent with application awareness. DASH Copy data to another region, cloud, or back to on-premises – complete data mobility by replicating to another geographical region within IaaS provider, a different IaaS provider, or back to on-premises sites. Protect Office 365 – mail, OneDrive for Business and SharePoint online either in cloud or back to on-premises Protect Azure Blob storage – Backup object storage repositories with data created by other 3rd party applications either in cloud, to an alternative provider, or back to on-premises sites. 	<ul style="list-style-type: none"> Azure Virtual Server Agent and MediaAgent deployed on an Access Node within IaaS provider for agentless backup. Applications not supported by VSA application awareness and requiring application-level consistency can be protected via agents-in-guest. Minimum 1x MediaAgent in cloud, and (optional) minimum 1x MediaAgent at secondary site (whether cloud or on-premises) for receiving replicated copy of data. Recommended to use a dedicated network from cloud provider to on-premises for best performance when replicating back to on-premises (Azure ExpressRoute).

Use data – Disaster recovery to the cloud

Business Value: Providing operational recovery of primary site applications to a secondary site from an external cloud provider.

Disaster Recovery to the Cloud

Recover your primary production site (on-premises, cloud) to the cloud for on-demand, near real-time disaster recovery resource



Use Data – Disaster Recovery to the Cloud

Scenario/Suitability	Requirements
<ul style="list-style-type: none"> Off-site storage & cold DR site in the Cloud – Only use the cloud compute infrastructure when a DR event occurs, saving time & money via the elimination of asset allocation with long idle periods between DR operations. Live Sync data replication for Warm Recovery in cloud – Automate the creation of cloud VMs and replication of on-premises VMs to Azure on a periodic cycle basis more frequently than backups. Reduces recovery time to the cloud. VM Restore & Convert – Convert VMware, Hyper-V, Acropolis, Oracle VM, and AWS VMs into Azure VMs on-demand with data intact. This data transformation automation reduces time & complexity costs. Automate Failover and Failback of VMs - From on-premises VMware to Azure. 	<ul style="list-style-type: none"> Database/Files – Restore out-of-place, whether on- demand or scheduled, to refresh DR targets. When combined with job-based reporting, this scheduled operation is of benefit to enterprises that must maintain audit and compliance reporting associated with business continuity reporting. Minimum 1x MediaAgent on-premises, and minimum 1x MediaAgent in cloud MediaAgent in cloud only needs to be powered on for recovery operations Highly recommended to use dedicated network to cloud provider for best performance (Azure ExpressRoute).

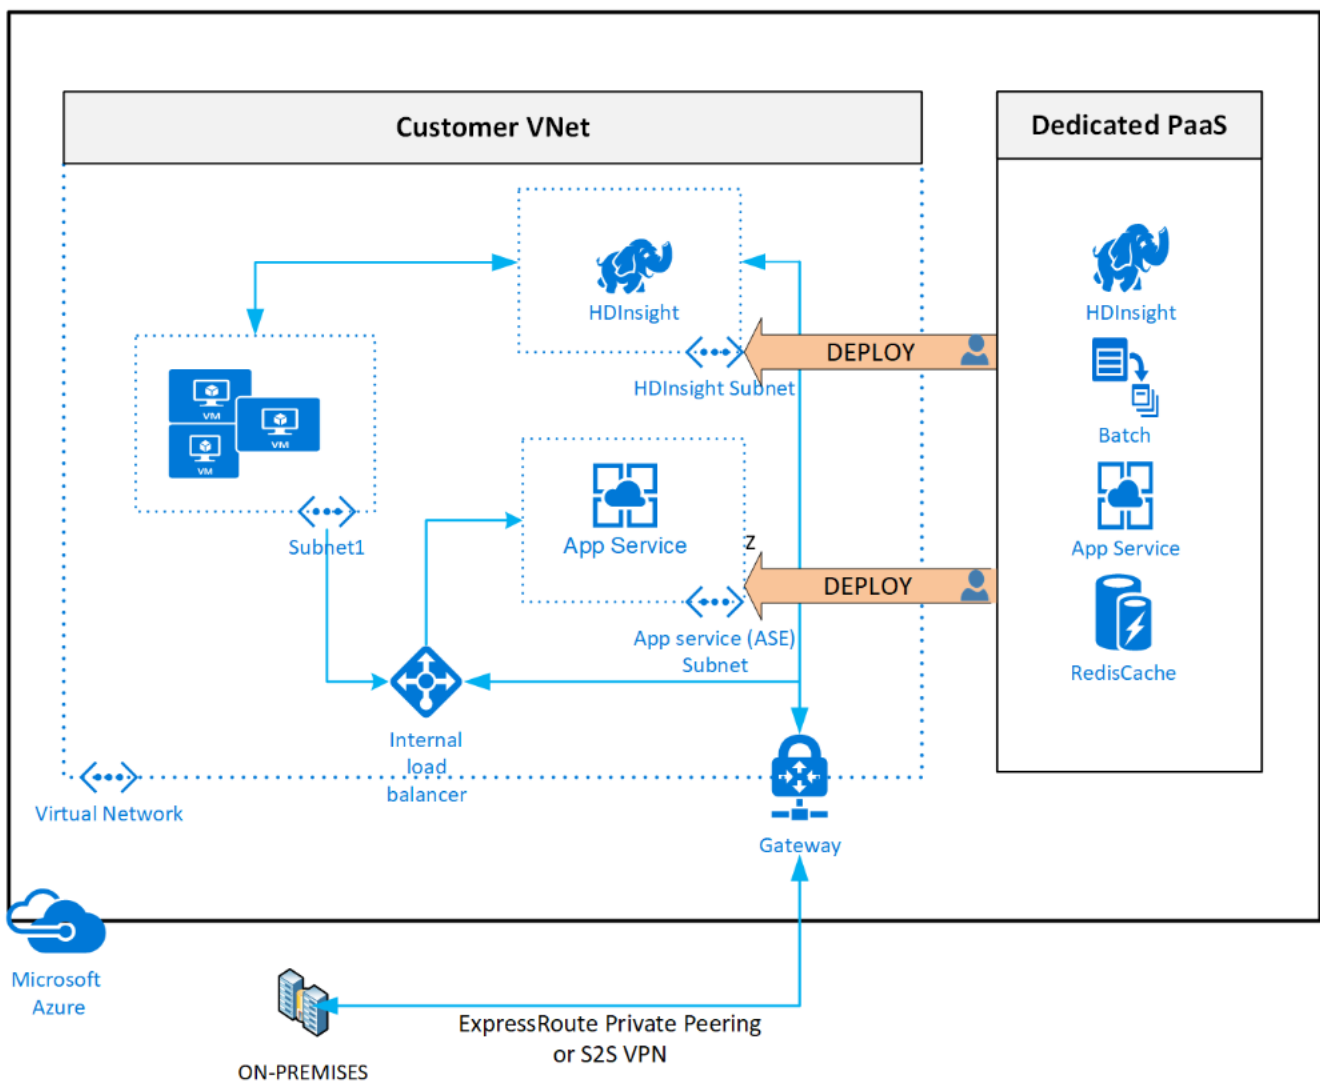
Architecture considerations

Networking

Virtual private cloud/networking

Microsoft Azure (Azure) has the capability to establish an isolated logical network. This is referred to within Azure as an Azure Virtual Network (AVN).

Virtual machines deployed within an AVN, by default, have no access to the Public Internet, and utilize a subnet of the customer's choice. Typically, AVN's are used when creating a backbone between virtual machines, and when establishing a dedicated network route from a customer's existing on-premises network directly into the public cloud provider via Azure ExpressRoute.



Bridging on-premises infrastructure – VPN and ExpressRoute

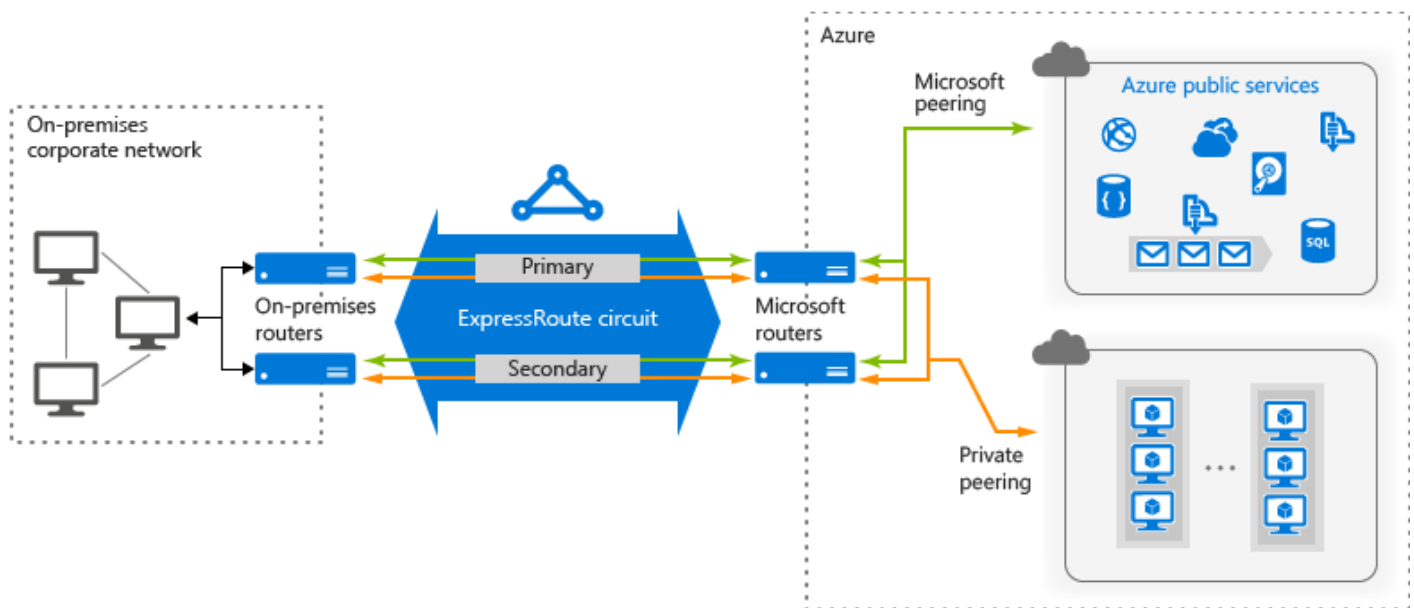
Customers may find a need to bridge their existing on-premises infrastructure to their public cloud provider, or bridge systems and workloads running between different cloud providers to ensure a common network layer between compute nodes and storage endpoints.

This is particularly relevant to solutions where you wish to Backup/Archive directly to the cloud or create deduplicated secondary data copies ([DASH Copy](#)) of existing backup/archive data to object storage within a cloud provider.

To utilize these features there are two primary choices available:

- **VPN Connection** – network traffic is routed between network segments over Public Internet, encapsulated in a secure, encrypted tunnel over the customer's existing Internet Connection. As the connection is shared, bandwidth is limited, and regular data transfer fees apply as per the customer's current contract with their ISP.
- **Azure ExpressRoute** – a dedicated network link is provided at the customer's edge network at an existing on-premises location that provides secure routing into an Azure Virtual Network. Available at various bandwidth speeds for both metered/unmetered charges for outbound data transfers. Details can be found at [Azure ExpressRoute >](#)

Typically, these links are less expensive when compared to a customer's regular internet connection, as pricing is charged on a monthly dual-port fee, with all inbound and outbound data transfers included free of charge (Unlimited Data plan), with bandwidth from 50Mbps to 100Gbps.



Firewall ports

When communicating with cloud services and resources, security is a high priority, and the use of firewalls are recommended. Commvault has registered specific default ports for its services that can be changed if desired. Several Commvault services listen for incoming network traffic on these ports.

Static ports

The CVD service port (8400) is used for communication of the CommServe server, MediaAgents, and clients.

The EvMgrS service port (8401) is used by the CommServe server for receiving events from MediaAgents and/or clients.

The FWD service port (8403) is used for tunneling connections across firewalls.

Commvault service	Port number	Protocol
Commvault Communications Service (GxCVD, found in all client computers)	8400	TCP
Commvault Server Event Manager (GxEvMgrS, available in CommServe server)	8401	TCP
Commvault Firewall (GxFWD, tunnel port for HTTP/HTTPS)	8403	TCP

Dynamic ports

Dynamic ports are opened and closed by the running Commvault software to permit certain types of transient traffic. The CVD service dynamically uses free ports between 49152 - 65535 to communicate during data protection and data recovery jobs. The system dynamically assigns a number of free ports to be used by each job to allow parallel data movement. After the job is finished, if no other job is pending, the dynamic ports are released.

Additional information can be found by referring to [TCP Ports Used for Services >](#)

Azure Virtual Network

Azure Virtual Network (VNETs) enables many types of Azure resources, such as Azure virtual machines, to securely communicate with each other, the internet, and on-premises networks.

Azure Virtual Network provides the following key capabilities:

- Isolation and segmentation
- Communicate with the internet
- Communicate between Azure resources
- Communicate with on-premises resources
- Filter network traffic
- Route network traffic
- Connect virtual networks
 - **Cost reduction** is a top priority for cloud users and there are many ways to reduce these costs. Here are a few key points to keep in mind when designing the networking in Azure:
 - Azure networking costs are limited to egress or outbound bandwidth. All incoming data to Azure data centers is free while the cost of outgoing data is tiered based on usage.
 - Outbound data transfer charges apply when the data transfer is to another region.
 - Region to region networking is performed through the use of Global VNet Peering.

- **Virtual Network peering** allows you to route traffic between virtual networks using private IP addresses and does incur both inbound and outbound bandwidth charges at the rate of \$0.01 per GB.
- VMs deployed in availability zones incur charges for inbound and outbound data transfers between availability zones at a rate of \$0.01 per GB.
- Data transfers within the same availability zone incur no charges.
- Data transfers from a VNet resource in an availability zone and a public address in the same Azure region incur no charges.
- When configuring VNet peering, the address spaces must be non-overlapping. This means you cannot have both VNet1 and VNet2 with an IP subnet of 10.1.0.1. VNet2 would have to have an IP address subnet of 10.2.0.1 as an example.

For more information on Azure Virtual Networks, please refer to this Azure [documentation](#) >

Data security

In-flight

By default, all communication with Cloud Libraries utilizes HTTPS which ensures that all traffic is encrypted while in-flight between the MediaAgent and the Cloud Library endpoint, but traffic between Commvault nodes is not encrypted by default. We recommend that any network communications between Commvault modules routing over public internet space be encrypted to ensure data security. This is employed by using standard Commvault [firewall](#) configurations (two-way and one-way).

At-rest

Data stored in a public cloud is usually on shared infrastructure logically segmented to ensure security. Commvault recommends adding an extra layer of protection by encrypting all data at-rest. To meet this requirement, Azure provides the capability to encrypt IaaS VMs with Azure Disk Encryption. Azure Disk Encryption supports both Windows and Linux VMs and it is integrated with Azure Key Vault to manage and control the disk-encryption keys and secrets. Along with IaaS VM encryption, most cloud providers require that any seeded data is shipped in an encrypted format. Examples of seeding data is with the use of [Azure Data Box](#) >

HTTPS proxies

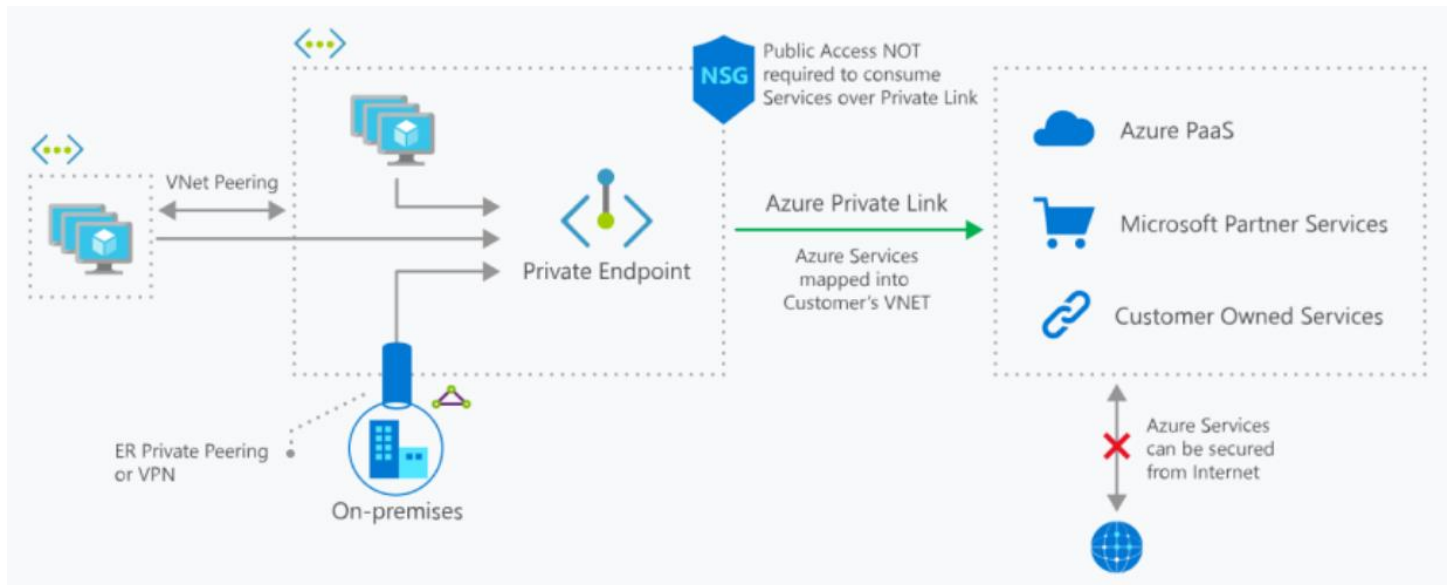
Please take note of any HTTP(S) proxies between MediaAgents and Endpoints, whether via public Internet or private space, as this may have a performance impact upon any backup/restore operations to/from an object storage endpoint. Where possible, Commvault software should be configured to have direct access to an object storage endpoint.

Account separation

Consider utilizing separate Azure accounts for production and non-production data protection activities. Commvault supports cross-account restores, allowing segregation of data access but also authenticated and authorized data mobility where required (e.g. Dev/test re-seeding restores).

Azure Private Link

Private Link is a service from Microsoft that provides private connectivity from a virtual network to Azure platform as a service (PaaS), customer-owned, or Microsoft partner services. It allows you to create a private endpoint in your virtual network and simplifies the network architecture. This endpoint provides you an internal IP address for the resource and all traffic from your virtual network to the resource goes over the Microsoft backbone instead out over the public internet. All traffic to the service can be routed through the private endpoint, so no gateways, NAT devices, ExpressRoute or VPN connections, or public IP addresses are needed. There are several benefits to this service but most importantly data security since only the mapped resource would be accessible.



How does Private Link differ from Service Endpoints? Service Endpoints provide a way to lock down access to resources to a virtual network, but you are still accessing the public endpoint. You are only locking down access to a service and not to a specific resource. With Private Link you are only allowing access to the defined resource, which greatly enhances data security and control of your resources. Private Link creates an endpoint with a private IP address and data flows solely inside your virtual network, which does not require Network Security Group (NSG) rules allowing outbound traffic beyond your virtual network.

There is no cost for the Private Link service, however there is a cost for the Private Endpoint (\$0.01 per hour) as well as all data inbound (\$0.01 per GB) and outbound (\$0.01 per GB).

Data seeding

Data Seeding is the process of moving the initial set of data from its current location to a cloud provider in a method or process that is different from regular or normal operations. There are two primary methods for seeding data to an external cloud provider:

“Over-the-wire”

This is typically performed in a small logical grouping of systems to maximize network utilization in order to more quickly complete the data movement per system. Some organizations will purchase “burst” bandwidth from their network providers for the seeding process to expedite the transfer process.

Major cloud providers offer a direct network connection service option for dedicated network bandwidth from your site to their cloud such as Azure ExpressRoute.

If you require an understanding of the possible throughput between two Azure VMs to assist with transfer times, you can use the Azure [NTTTC](#) tool.

Please see the chart below for estimated payload transfer time for various data sizes and speeds.

Link size	Data set size							
	1 GB	10 GB	100 GB	1 TB	10 TB	100 TB	1 PB	10 PB
10 Mbit	14 min	2.2 hrs	22.2 hrs	9.2 days	92.6 days	-	-	-
100 Mbit	1 min 20 s	13 m 20 s	2.2 hrs	22.2 hrs	9.2 days	92.6 days	-	-
1 Gbit	8 s	1 m 20 s	13 m 20 s	2.2 hrs	22.2 hrs	9.2 days	92.6 days	-
10 Gbit	0.8 s	8 s	1 m 20 s	13 m 20 s	2.2 hrs	22.2 hrs	9.2 days	92.6 days

Drive seeding

If the dataset is too large to copy over the network, or transport over network is too costly, then physical drive seeding is a valid alternative option. Drive seeding is copying the initial data set to external physical media and then shipping it directly to the external cloud provider for local data ingestion. The Azure Data Box family of products is supported and can be used to migrate data to Azure.

Please refer to the Commvault’s online documentation for the [Seeding the Cloud Library](#) procedure for more information.

In addition to this, Azure has their own process for drive seeding:

- [Import Services](#)
- [Azure Databox](#)
- [Migrating Data to Microsoft Azure Using Azure Data Box](#)

Cost/Consumption

Network Egress

Moving data into a cloud provider, in most cases, has no provider cost, however moving data outside the cloud provider, virtual machine instance, or cloud provider region, usually has a cost associated with it. Restoring data from the cloud provider to an external site or replicating data between provider regions are examples of activities that are classified as Network Egress and usually have additional charges. Pay special attention to the tier of storage. Some storage tiers cost more for egress and others are free. This may impact your storage costs enough to decide to choose a higher tier of storage like Hot storage tier instead of Cool or Archive storage tier. **Azure egress fees** are charged per GB per month, so ensure you are aware of the costs and design your data movement in the most efficient way possible.

Storage I/O

The input and output operations to storage attached to the virtual machine. Cloud storage is usually metered with a fixed allowance included per month and per unit “overage” charges beyond the allowance. Frequent restores, active data, and active databases may go beyond a cloud provider’s Storage I/O monthly allowance, which would result in additional charges.

GET/PUT transaction costs

Azure customers usually incur a cost for GET/PUT transactions to cloud object storage. These costs are primarily to enforce good practices for applications when retrieving and placing data in the cloud. As such, the cost when using the Commvault solution is minimal.

When Commvault writes data to a Cloud Library, the Cloud Library splits the data up into a sub-chunk size of 32 MB. Each 32 MB chunk write or read will incur a GET or PUT request. As of April 2020, Azure charges \$0.05 per 10,000 PUT requests and \$0.004 per 10,000 GET requests for its LRS Hot storage tier for example.

A baseline of 2,000 GB with a saving of 40% at 32 MB sub-chunk size would result in an approximately 37,500 PUT requests. At a charge of \$0.05 per 10,000 requests, the cost would be 20 cents.

Note: All cost figures are referenced in USD and based on pricing listed at the [Azure website](#) at the time of this document’s publication.

Data recall

Low-cost cloud storage solutions may have a cost associated with accessing data or deleting data earlier than the minimum days requirements of the tiers. Hot tier has no requirements, Cool tier is 30 days, Archive tier is 180 days. Storing infrequently accessed data on a low-cost cloud storage solution may be attractive upfront, however Commvault recommends modeling realistic data recall scenarios. In some cases, the data recall charges may be more than the potential cost savings vs. an active cloud storage offering.

As a best practice, Commvault recommends developing realistic use case scenarios and modeling cost against the identified scenarios to ensure the cloud solution meets your organization’s SLAs, as well as cost objectives by leveraging the [Azure cost calculator](#) >

Performance/Storage

Multi-streaming with object storage

Object storage performs best with concurrency, and as such with any cloud libraries configured within the Commvault environment, best performance is achieved when configured for multiple readers / streams.

Cloud connector best practices

There are additional Data Path settings and other additional settings used to adjust and fine-tune the performance of the cloud library. The exact settings for best performance may vary between cloud vendors.

The following combined settings are recommended to increase data read performance from cloud libraries utilizing deduplication.

Device Streams > 50

- **Device Streams** are logical channels that connect client data to the media, where data that is secured by backup operations are stored. For a Storage Policy, the number of device streams that is configured must be equal to the number of drives or writers of all libraries that are defined in the storage policy copy. Increasing the number of Device Streams above 50 increases parallel streams during protection activities.

Deduplication block size = 512 KB

- Increase **Deduplication Block Size** to either 256 KB or 512 KB for maximum performance. Cloud object storage is subject to 50% or higher latencies than traditional storage. When requesting blocks from object storage, this delay may reduce overall read performance. To counteract the delay, increase the deduplication block size to allow larger data retrievals for each request. Note that changing existing storage policies will initially cause an increase in storage as new deduplication data re-baselines. Only new data written with the higher block size will benefit from retrieval performance improvements. If the requirement is to keep one copy on-premises and another in the cloud, recommendation is to use 256 KB block size for on-premises and cloud copy. Otherwise, if one or all copies involved will be using cloud storage, recommendation is to use 512 KB block size. The reason for this is, you cannot choose a different deduplication block size for multiple copies within a Storage Policy, allowing this will unnecessarily increase the overhead in creating the secondary copies as data will now have to be rehydrated and re-deduplicated with the new block size. As of April 2018 (Commvault v11 SP11), the block size for a deduplication cloud library will automatically be created to use 512 KB.

Data Readers

- **Data Readers** determine the number of parallel read operations while the data is backed up. Configuring multiple data readers per subclient on disk arrays can improve the backup performance of clients. Increase the Number of Data Readers to increase the backup read performance.
-

Compression vs. deduplication

Deduplication is recommended to be used where possible, with the exception of environments where there are significant bandwidth concerns for re-baselining operations, long-term retention designs in the Archive tier (tape replacements use case), or for Archive-only use cases where the data pattern spread generates no benefit from deduplication operations.

While additional compute resources are required to provide the necessary foundation for optimal deduplication performance, using deduplication in a cloud context can still achieve greater than a 10:1 reduction.

Even with sealing of the deduplication database (DDB), stored data results can achieve a 7:1 reduction in footprint, providing significant network savings and reduced backup/replication windows (DASH Copy).

In comparison, software compression can only achieve 2:1 reduction on average and will constantly consume the same bandwidth when in-flight between endpoints (no DASH Copy).

Leveraging multiple mount paths for a cloud library

Just like regular disk libraries, cloud libraries have the option to leverage multiple mount paths. The benefit of using multiple mount paths depends on the cloud storage vendor. In Azure, using multiple mount paths may help to increase performance but typically it's not required.

Block blob versus Page blob object storage

With Azure, be aware that both block object blobs as well as page blobs, are potential options when creating a storage account. Public IaaS environments allow premium (page blob) based storage to be provisioned and leveraged as disk libraries, by attaching directly to the MediaAgent. The overall cost of those volumes can quickly exceed that of standard (block blob) object storage. While Commvault software supports and can consume page blobs, this option incurs a higher cost compared to standard block blobs, and this cost should be evaluated before making this choice.

With the inclusion of Commvault micro pruning, and its benefit of reducing cost of data stored in object storage, it is highly recommended that standard block object storage be the primary choice for writing data to the cloud, and other forms of storage by exception.

If you are unsure as to which offering to use, you should consume standard [block object storage blobs](#) >

Partitioned deduplication

Like on-premises configurations, making use of partitioned deduplication can provide several benefits. When possible, make use of partitioned deduplication to increase scale, load balancing, and failover. Commvault version 11 allows for the addition of two extra nodes (up to 4) to an existing deduplication store dynamically, allowing for rapid scale-up configurations.

[Configuring Additional Partitions for a Deduplication Database](#) >

Micro pruning

Micro pruning support for object storage is effective for any new data written into the active deduplication database. For customers who have upgraded from Version 10 of Commvault, but have not yet enabled micro pruning support, macro pruning rules will still apply to existing data within the active cloud storage library until the deduplication database has been sealed. Once the active library is sealed, there will no longer be a need for continued periodic sealing against that

library. Micro pruning is supported in the Azure Hot and Cool storage tiers. Micro pruning is **not** supported in the Azure Archive storage tier.

Selecting the right storage class for backup and archive data

Depending on the cloud provider, there may be different tiers of object storage available that offer different levels of cost, performance, and access. This can have a significant impact on both the cost and the user experience for the datasets within the cloud storage.

For example, storing infrequently accessed backups within an intermediate storage tier (Cool blob) can significantly lower the cost of your cloud bill, while storing data in an archive storage tier (Archive blob) may greatly impact accessibility for end-users to the archived data but have the added benefit of reducing storage costs further. To delve into further detail, these storage classes can be broken into three primary categories:

- **Standard storage (Hot)** – this storage class represents the base offering of any object storage platform – inexpensive, instant access to storage on-demand. Offerings in this category include Azure Blob Hot storage tier. Typically, this tier would be used for backup and archive workloads with a very short-term retention configuration or where restores are frequent.
- **Intermediate storage (Cool)** – this is a storage tier that addresses a gap between the standard Hot storage tier offering and Archive storage tier, in that it is offered at a lower price point than Hot storage but is aimed at scenarios where data is infrequently accessed. While the storage is always accessible, similar to the Hot offering, the cost model is structured to enforce an infrequent access use case by charging \$0.01/GB for any retrieval from this storage tier. Offerings in this category include Azure Blob Cool storage tier. This tier would be leveraged for backup workloads in a medium to long-term retention configuration, and for Archive workloads that require instant access to the archived data. This tier functions well as a secondary copy for 30+ days.
- **Archive storage (Cold)** – sometimes referred to as “cold storage”, this tier is intended for data that will probably not be accessed again, but must be retained in the event of compliance, legal action, or another business reason, Azure Archive Blob storage tier is an example of archive storage which Commvault software supports. The cost of this storage class is the lowest compared to all three offerings – \$0.001/GB/month, but as with the intermediate class, the archive class’s cost model is also structured with the expectation that retrievals are infrequent and unusual, and data will be stored for an extended period of time. In addition to the per-GB, per-month charge, any data that is moved to Archive is subject to an early deletion period of 180 days. The charge is prorated based on the amount of time the data has been stored in archive storage. You can think of this class of storage as equivalent to tape and is therefore recommended not to use deduplication.

It is highly recommended that you review the cost options and considerations of each of these storage classes against the use case for your architecture to gain the best value for your cost model. Commvault Professional Services can assist in necessary service class / data class valuations in designing the correct cost value model for your enterprise.

Additional information: [Supported Cloud Storage Products >](#)

Azure storage redundancy

Azure provides redundancy of data stored in a storage account to ensure durability and high availability. The redundancy options available are:

- Locally Redundant Storage (LRS)
- Zone-Redundant Storage (ZRS)
- Geo-Redundant Storage (GRS)
- Geo-Zone-Redundant Storage (GZRS)
- Read-Access Geo-Redundant Storage (RA-GRS).

These **redundancy levels** are transparent to Commvault and can be used if required.

Choosing the correct Azure storage tier

Choosing the correct Azure storage tier for backup data is crucial when considering performance, cost, and accessibility. If, after reading through the information below, you are still unsure where to store backup data, we recommend:

Azure Cool Locally-redundant storage (LRS) General Purpose v2 storage



Performance/Access tier : Standard/Cool
 Replication : Locally-redundant storage (LRS)
 Account kind : StorageV2 (general purpose v2)

It has the best cost/performance of all tiers and should be used for most primary copies. Keep in mind this is just a recommendation and there are always exceptions.

Each tier of storage has specific characteristics which need to be fully understood. However, before choosing the storage tier, you first need to determine what the business requirements demand for accessibility and retention.

For example:

Data storage requirements	Hot	Cool	Archive
Data retention (Days, Months, Years)	Days	Months	Years
Accessibility (Frequent, Occasional, Rare)	Frequent	Occasional	Rare
Time to recover (Minutes, Hours, Days)	Minutes & hours	Minutes & hours	Hours & days

Additional factors:

- Location of data to be protected (On-premises or in the cloud)
- Location of MediaAgent (On-premises or in the cloud)
- Location of backup storage (On-premises or in the cloud)
- Internet bandwidth (Public internet or Azure ExpressRoute)

The answers to these questions may be the same for all data being protected, or they may be different depending on the type of data (databases, healthcare, financial, filesystems, etc). We can then determine the storage tier best suited for the data.

- **Azure Hot** storage is designed for fast and easy access to data, but at the highest cost (\$0.0184 USD) per GB. Data stored in Hot tier is recommended for short term retention (primary backup copy).

- **Azure Cool** storage is technically the same as Hot storage and is designed for fast and easy access to data. However, it's cheaper at ~45% less, (\$0.01 USD) per GB. The reason for the price break is because Azure is expecting the data to remain in Cool storage for a minimum of 30 days. If the data is removed prior to 30 days, the storage cost is prorated. Data stored in Cool tier is recommended for short-term and long-term retention that requires access for restores on frequent or occasional use (primary backup copy and long term retention).
- **Azure Archive** storage is the lowest price storage tier, (\$0.002 USD) per GB. The purpose for this tier is for a very specific use case. Data placed in Archive storage is expected to remain for a minimum of 180 days. Any data removed prior to this is subject to an early deletion cost and is prorated. For example, if data is moved to the Archive tier and then deleted or moved to the Hot/Cool tier after 75 days, the customer is charged an early deletion fee for 105 (180 minus 75) days of storage in the Archive tier.

Azure offers reserved capacity for blob storage, which provides reduced cost of storage through discounts. This is done by committing to either a one-year or three-year term for a fixed amount of storage capacity. The cost savings achieved depend on the duration of your reservation, the total capacity you choose to reserve, and the access tier and type of redundancy that you've chosen for your storage account. Reserved capacity provides a billing discount and doesn't affect the state of your Azure Storage resources. Reserved capacity is purchased in units of 100 TB and 1 PB per month for a one-year or three-year term.

- [Azure Blob Storage pricing](#)
- [Optimize costs for Blob storage with reserved capacity](#)

Commvault Combined Storage tier

In addition to Azure Hot, Cool, and Archive storage, Commvault also provides two additional storage options and are unique to Commvault. We have combined Azure Archive storage with Hot and Cool storage to assist when writing deduplicated data. **Hot/Archive** and **Cool/Archive** are beneficial to those who want to store data (SFILE) in Archive storage, but still have fast access to the chunk metadata (CHUNK_META_DATA) for fast deduplication lookups.

Note: As of Feature Release 11.23 the job index data is stored in the Hot/Cool tier for fast browsing of files in the job for restores. There is additional cloud storage costs for the storage of the index, but the efficiency and smoother user experience far outweighs the small additional costs. This does not improve the restore time of the actual data from Archive tier drastically which could still be upwards of 15 hours. Therefore, Commvault recommends using combined storage tiers for backup data with long term retention using deduplication.

Combined storage tiers are not recommended for primary backup copies. The reason for this is because of the time required to recall the data as well as the cloud costs incurred for the recalls. If you need fast access to the data to recover the business operationally – use the Hot or Cool storage tier and not a combined storage tier.

For customers that are looking to use combined storage tiers and currently have their data in Hot or Cool, you will need to configure and add a combined storage tier to the Commvault Plan. Set the retention policy for the combined storage and define the primary copy as the source. This will aux-copy the data from the primary copy to the combined storage tier. We do not 'move' the data, it's a copy. Once the data has been copied to the combined storage tier, you can change the retention setting of the primary copy, which will prune the data from the primary copy. The result will be reduced cloud storage costs in the primary storage copy. However, you should make sure the primary copy retention is sufficient for typical data recall requirements, so as not to incur Archive recall costs.

Combined Storage Tiers

Commvault metadata stored in Azure Hot\Cool tier



+



+



Backup Job Indexes

Backup Job Indexes store the details of data retained in each individual backup job or job copy (auxiliary copy).

Backup job indexes are mandatory to present a list of valid backups for recovery.

Pruning Indexes

Pruning indexes hold details on which individual blocks in the backup store are used or expired.

Commvault uses pruning indexes to determine when data can be 'deleted' from the backup store.

Deduplication metadata

When **deduplication** is utilized, deduplication metadata is stored in the Hot\Cool tier. This metadata is used for restoration and DDB reconstruction activities.

Commvault stores all indexing data within the 'Hot or Cool' storage class when using **Combined Storage Tiers**, for seamless recovery

Combined Storage Tiers

Restore Types



Automated

- Commvault **alert** triggers Cloud Archive Recall Workflow
- Workflow orchestrates recall of selected data
- Limited to **specific agents**
- Granular restore of files, folder, tables, etc.



Manual / On-Demand

- **Cloud Archive Recall Workflow** is manually executed
- **Job ID** of required job is provided
- **Cloud Archive Recall** recalls the data for selected Job ID.
- Supported by **all commvault agents**


Best Practices



Secondary copies only

Commvault combined storage tiers are designed for infrequently used data.

They are not designed to be used as the **Primary copy** location.



Use for Storage copies

Commvault combined storage tiers represent an additional **storage copy** configured in your Plan/Storage Policy.

Commvault does not move data between cloud libraries



Just add a storage copy...

If combined storage tiers are adopted after a cloud library is already populated...

Individual **jobs** must be copied, then expired from primary location.



Assess data needs day one

To avoid lengthy data migration processes, consider your **near-term, mid-term, and long-term** data accessibility when creating your plans.

	Hot/Archive	Cool/Archive	Archive
Deduplicated backup	Metadata and job indexes will be written in the Hot tier and data will be written in the Archive tier	Metadata and job indexes will be written in the Cool tier and data will be written in the Archive tier	Metadata, job indexes, and data will be written in the Archive tier
Non deduplicated backup	All data will be written in the Hot tier	All data will be written in the Cool tier	All data will be written in the Archive tier
Restore	Workflow will recall the data from the Archive tier to the recall destination. (Recall destination can be set in the Workflow. Can be Hot or Cool. Default value is Cool in the workflow.)	Workflow will recall the data from the Archive tier to the Cool tier. (Recall destination can be set in the Workflow. Can be Hot or Cool. Default value is Cool in the workflow.)	For deduplicated data, Workflow will recall both the metadata and data from the Archive tier, directly to the restore location. For non-deduplicated data, all data will be recalled from the Archive tier, directly to the restore location.
After restore	Data will be moved back to the Archive tier by the Workflow	Data will be moved back to the Archive tier by the Workflow	For deduplicated data, Workflow will move the metadata and data back to the Archive storage, once the recall period provided in the Workflow is met. For non-deduplicated data, all data will be moved back to the archive storage

Note: Micro-pruning is not supported in the Azure Archive storage tier.

Consider the following characteristics when selecting your cloud storage – there is cost in storing, cost in retrieving, and cost associated with the **latency** of your data retrieval. While High Priority retrieval options are possible, these should be considered exceptional events and choosing a warmer storage class is recommended if a likelihood for rapid recovery is expected.

Azure Storage Classes – for Backup & Archive			
Feature	Hot	Warm	Cold
Storage Class	Azure Hot	Azure Cool	Azure Archive
Annual Storage cost (1 TB)	\$220.80	\$120.00	\$24.00
Minimum retention	None	30 days Note: Recalls under 30 days will be pro-rated	180 days Note: Recalls under 180 days will be pro-rated
First byte latency	milliseconds	milliseconds	Up to 15hrs Typically around ~4hrs
Recall cost (1 TB)	\$0.016	\$10.056	<ul style="list-style-type: none"> Standard: \$40.00 High-Priority: \$300

Note: Pricing is for LRS storage from region 'East US 2' and is in \$USD. Minimum retention is the time required for data to remain in storage or the data will be pro-rated and charged early deletion fees. Recall costs are based on a single recall of data from storage and includes the cost of 'Read Operations' and 'Data Retrieval'.

<https://azure.microsoft.com/en-us/pricing/details/storage/blobs/>

Commvault recommends a **tiered approach** to consuming Cloud storage from your Primary application data, Secondary (backup) data, and finally your Tertiary (archive) data vaults.

Warning: Make sure the Azure Life-Cycle policy is disabled for any Storage Account being used for Archive data. There is a conflict where Azure re-archives any data we recall before we can use it for the restore.

More information can be found here:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts?tabs=azure-portal>

Azure Immutable Storage (WORM)

Immutable storage, also known as WORM (Write Once, Read Many) storage, enables users to store business-critical data with the requirement that it cannot be modified or deleted based on a user defined retention period. Designing a solution with offsite copies to protect against ransomware and cyber threats is imperative. Commvault provides the ability to utilize cloud immutable storage with Azure Blob storage for enhanced data security.

Immutable storage for Azure Blob storage supports two types of immutability policies:



Time-based retention policies: a time-based retention policy which users can define for a specified time. Data can be created and read but cannot be modified or deleted until the retention period has expired.

Legal hold policies: a user defined retention policy which explicitly holds data until the policy is cleared. A user can define for a set of data as legal hold. This subset of data can be created and read but cannot be modified or deleted until the legal hold is cleared.

When enabling **Azure immutable storage**, Commvault makes the following changes:

- Configures a default **retention period** > to 2x the data retention period set on the Server Plan (or Storage Policy)
- Enables **default retention mode** > (legal hold) on the container, for newly written objects.
- Configures **periodic DDB sealing** to match the Server Plan (or Storage policy).
- Micro-pruning is disabled on the Cloud Library

Configuration Example

For associated deduplication-enabled storage pools, retention at the container level is the sum of the value at the storage policy copy level and the DDB seal frequency value on the cloud. The seal frequency is set to either a maximum of 180 days or half the retention days of the storage policy copy, whichever is greater. For example, if the retention is set for 60 days at the storage policy copy level, the seal frequency is set to 180 days, and then at the cloud storage level, the retention is set for $60+180=240$ days. If the retention at the storage policy copy level is 2 years, then the seal frequency is set to 1 year, and the cloud storage level retention is set to $2 + 1=3$ years.

- A cloud library server plan with 90 days retention
- Default retention period will be 180 days
- The DDB will be sealed every 90 days

Note: Immutability should ideally be enabled on the backup library before writing any backup data to ensure data is immutable from initial creation.

Effects of DDB sealing

When the DDB is sealed, the sealing process closes the DDB and starts a new DDB. When the new database is started, the next instance of each data block processed creates a new signature tracking entry and the data block is written to the disk again as a new initial baseline.

This will result in a full copy of the backup content being resent to the Cloud Library. This is intentional and provides multiple, segregated data copies to protect from corruption or other unforeseen data access issues.

Additional information can be found here:

- [Configuring the WORM Storage Mode on Cloud Storage](#)
- [Workflow for Configuring WORM Storage Mode on Cloud Storage](#)
- [Sealing the Deduplication Database](#)

Access policy

container1

Save

Immutable blob storage

Policy type ⓘ

Legal hold

i Each legal hold policy needs to be associated with 1 or more tags. Tags are used as a name identifier, such as a case ID, to categorize and view records. Retention policy changes may require some time to take effect. [Learn more about immutable blob storage](#)

Tag

Add tag

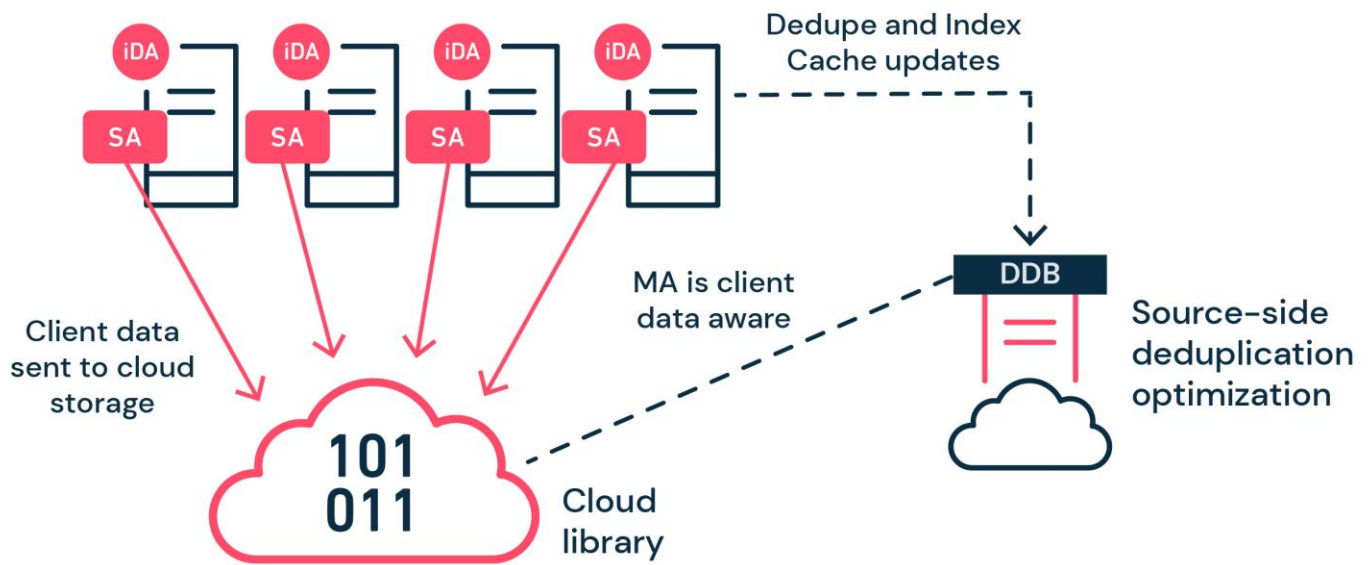
OK

Cancel

Storage Accelerator to Azure Blob Storage

For remote office locations, small cloud environments, roaming devices such as laptops, and any architecture that proves unfeasible or cost prohibitive to implement a traditional or cloud-based MediaAgent, backups can be done directly from the source to a cloud target such as Azure Blob storage, completely bypassing the MediaAgent. Workloads with large datasets may benefit highly by this use case, as an example VMs that are greater than 1TB. This is achieved by installing and enabling the Storage Accelerator feature on the client for direct communication to a storage target and will speed up the backup and reduce costs in these instances. When deduplication is used, the client communicates with the MediaAgent to reference and update the dedupe database, and then sends only the unique blocks directly from the client to the cloud storage blob.

Get additional information on the [Storage Accelerator >](#)



The example above shows the storage accelerator installed on the clients. During a backup job, the clients communicate with the MediaAgent deduplication database (DDB) and only send the unique blocks directly to the cloud storage target.

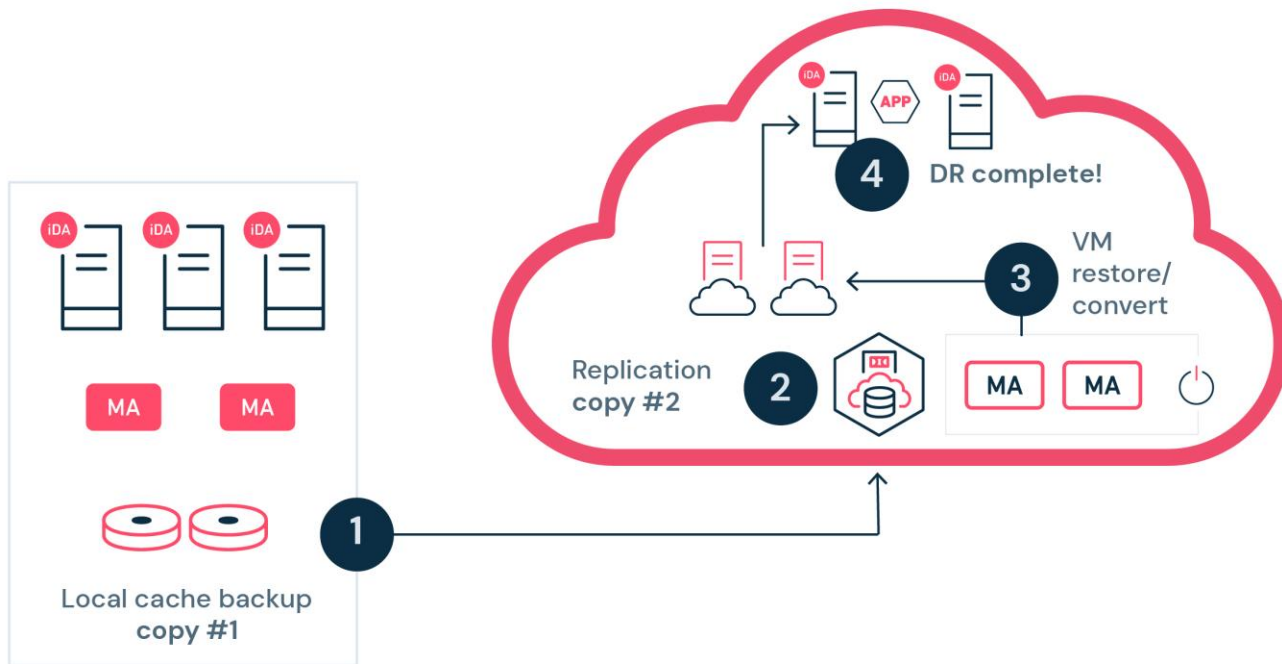
Restores are handled through direct transfer of data from the cloud storage to the client, bypassing the MA, which greatly reduce restore times. This is especially powerful when many restores are happening at the same time to different clients. The throughput of the MediaAgent is not a limiting factor. The throughput will depend on the cloud storage and the client. Since the MediaAgent is only handling dedupe tasks and not the transfer of data for these clients, the load is reduced and can potentially be sized smaller.

During testing the performance was obvious and provided faster backups, 3 to 5 times faster.

Note: The storage accelerator role consumes additional compute and network resources while performing data protection operations and should be managed so as not to interfere with production use.

Performing disaster recovery to the cloud

This section will cover the steps required to perform disaster recovery into the Azure public cloud platform. We examine recovery methods available for both image and agent-based protection. This also addresses different recovery scenarios that may be needed to meet short recovery time objectives.



Restoring applications (automated or on-demand)


An agent-in-guest approach allows for the recovery of a wide variety of operating systems and applications. These can be captured at the primary site and replicated to the cloud based MediaAgent in a deduplicated efficient manner. Once replicated, the data can be held and restored in the event of a disaster recovery scenario or automatically recovered to existing VMs for more critical workloads.

VM Replication (Live Sync)

Live Sync allows you to replicate VMs to the same or a different hypervisor, including public cloud infrastructure. As of December 2016, Azure is a supported cloud infrastructure vendor for Commvault v11. Live Sync combines the VM conversion feature with incremental replication to provide a disaster recovery solution utilizing on-demand cloud infrastructure. As Live Sync to cloud integrates with Dash Copy, highly efficient WAN replication is possible, by reducing the amount of data being replicated. You can replicate to the same Azure subscription or a different subscription, and to the same region or a different region. Live sync for Azure is supported from streaming backups or IntelliSnap backup copies.

Azure provides greater flexibility to important virtual machines and offers superior achievable RPO and RTO targets when used in conjunction with Live Sync. Live Sync can also be used as an alternative to limited bandwidth scenarios where large data transfer isn't achievable. The following hypervisors are supported for Live Sync replication to Azure: VMware, Hyper-V, Azure Stack HCI, and Azure VM.

As such, a good strategy is to identify multiple approaches depending on the business RTO/RPO requirements and implement them accordingly, while also considering the limitations and requirements specific to the cloud vendor. For example, Tier 1 applications may be a better fit for near-continuous replication using Commvault's CDR technology, while Tier 2 applications could make use of Live Sync (VMs, Files, DB), and Tier 3 apps could use on-demand VM conversion from cloud storage when needed. Live Sync supports two configuration options:



Create and maintain a VM in Azure in a stopped state – This provides a low RTO for a customer, but the customer incurs Azure costs for resources (storage, IP, etc.) utilized by the VM.

Create a VM at the time of failover – This saves the customers cost because no resources are used until the VM is created. However, this can elongate the customer's RTO (usually only takes a few minutes) because the VM needs to be created after the failover occurs.

Note: There are limitations when writing to an Azure Managed disk which affects our replication methods. The API provides the ability to create a Managed disk but doesn't allow us to reopen the disk for incremental writes in the same way we can with an Unmanaged disk. In order to overcome this, we stage the VM disk being replicated in blob storage so that we can write the incremental data. Once the VM is required in a Failover state we create the Managed disk from the blob storage.

Get additional information on the [Conversion feature >](#)

Replicating other workloads

Commvault Continuous Data Replicator (CDR) allows near time continuous data replication for critical workloads that must be recovered in adherence to Service Level Agreements that exceed the capabilities associated with Live Sync operations. These VMs require similarly sized Azure VMs to receive any replicated data. In order for CDR to operate, an Azure VM must be running at all times to receive application changes.

Get additional information on [CDR >](#)

Virtual machine recovery from Amazon EC2 to Azure VM

Since the release of Commvault v11 SP7, you can recover Amazon EC2 instances protected with the Virtual Server Agent (VSA) to an Azure VM for disaster recovery or migration purposes. Currently, streaming backups are supported for recovery to both Azure Resource Manager and Azure Classic.

Get additional information on the [Conversion feature >](#)

Azure-specific workloads

Virtual machine recovery into Azure VM instances

The Commvault Virtual Server Agent provides the ability to easily perform direct conversion of protected VMs to Azure VMs from the following hypervisors:

- Amazon
- Azure Classic
- Azure Resource Manager
- Azure Stack Hub
- Hyper-V
- Nutanix AHV
- Oracle VM
- VMware

Backups can be stored either within Azure Blob storage, another Cloud Library, or from an on-premises disk library.

This process is used as part of a disaster recovery strategy using Azure as a Cold disaster recovery site, or as a migration strategy (Lift-and-Shift).

Additional details can be found [Cross-Hypervisor Restores >](#)

Using Commvault Workflows to automate Disaster Recovery

The Commvault Workflow engine provides a framework in which the disaster recovery runbook process, covering the deployment of new VMs, recovery of data and applications, and validation aspects of a DR operation can be automated to deliver a simplified, end-to-end GUI-driven DR process. This can be developed and maintained by your administrators, or with the assistance of the Commvault Personalization Services team.

For more information on Commvault's Personalization Services team, please contact Commvault or the Commvault Partner Account team.

For more information on the Workflow engine, please refer to the [Workflow Overview link](#)

Protecting and recovering active workloads in Azure

This section outlines the basics on protecting active workloads running in Microsoft Azure. This portion of the document outlines the various protection approaches as well as replication and recovery to different geographic regions. This section also reviews cross platform recovery as well as recovery to onsite locations.

Agent-less VM protection (Virtual Server Agent for Azure)

Introduced in Commvault V11 SP4, the Virtual Server Agent (VSA) for Azure delivers an agent-less, block-level capture of Azure VMs and their attached block volumes. Restoration options include Full virtual machine recovery, attaching disks to an existing virtual machine, and granular-level file recovery. Azure VSA optionally includes Changed Block Tracking (CBT), included with v11 SP5, which helps accelerate incremental backup performance. With v11 SP15, Commvault added support for 8 TB VM disks and backup and restoration of VMs that are using Azure Disk Encryption with Azure Key Vault was added with SP15. These features provide customers with additional functionality when protecting and recovering IaaS VMs in Azure. The VSA Proxy is also referred to as the **Access Node**.

There are two types of storage accounts that can be used to provision blob, table, queue, file storage, and virtual machine hard disks. You can create a virtual disk in the Azure cloud by working directly with a storage account or you can let Azure manage the storage account for you with Managed Disks. Azure Managed Disks simplifies disk management for Azure IaaS VMs by managing the storage accounts associated with the VM disks. Specify the type (Premium or Standard) and the size of disk you need, and Azure creates and manages the disk for you.

Commvault software has had agent-less VSA protection for Unmanaged Disks since v11 SP4. As of v11 SP10, protection of Azure Managed Disks is also supported.

Azure Generation 2 VM

Azure added support for Generation 2 VMs in February 2020. Gen 2 VMs use the new UEFI-based boot architecture, as opposed to the BIOS-based architecture used by Generation 1 VMs.

Some of the benefits of Gen 2 VMs:

- Build larger VMs (up to 12 TBs)
- Provision OS disks sizes that exceed 2 TBs
- Supports Windows and Linux Managed VMs

As of FR 11.19 Commvault is able to protect Azure Generation 2 VMs. During protection operations, the "VM Generation" is identified so that a Gen 2 VM is protected with the ability to restore, replicate, and convert.

The following features are supported:

- Streaming backup
 - Snapshot
 - Backup Copy
 - Azure to Hyper-V conversion
 - Hyper-V to Azure conversion
 - Azure to Azure replication
-


- VMware to Azure conversion
- Changed Block Tracking
- Managed VMs only supported by Microsoft

Additional information:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/generation-2>

Azure Availability Zones

An Azure Availability Zone is a high availability offering that protects applications and data from datacenter failures. They are unique physical locations within an Azure region. Each zone is made up of one or more datacenters equipped with independent power, cooling, and networking. To ensure resiliency, there's a minimum of three separate zones in all enabled regions. During VM creation the availability zone can be configured based on the users requirements. Commvault will capture this information during VSA backup. When a restore is initiated, the VM will be restored back to the availability zone originally defined. This feature is enabled by default in FR 11.21 and does not require any configuration. Microsoft currently only supports Managed VMs with availability zones.



Replicate and migrate VMs to a specific Azure Availability Zone

As of Feature Release 11.23 VMs can be replicated and migrated to an Azure Availability Zone from the Commvault Command Center. This reduces the complexity and time required to ensure important workloads are configured with redundancy based on their requirements. When migrating non-Azure VMs to Azure you can specify the Availability Zone (1, 2, 3, None, or Auto). A job will fail if the VM is not able to be migrated or replicated to a specifically defined Availability Zone (1, 2, 3, or None). Leaving it set to 'Auto' will choose the Availability Zone setting of the source VM. If the source VM is not an Azure VM or does not have an Availability Zone defined, the VM will not be placed in an Availability Zone. The 'Auto' setting is a best-effort only and no job failures will be associated with it.

Additional information:

- [Replication with Availability Zones](#)
- [Azure Availability Zones](#)

When to use the VSA for Azure

- Agent-less protection approach for Azure VMs & file-level data – no agents are required in-guest to perform block-level backup to provide application aware VM and File-level recovery. The following matrix shows which applications are supported for application aware protection for Microsoft Azure. Application aware protection for Azure was added with V11 SP16.

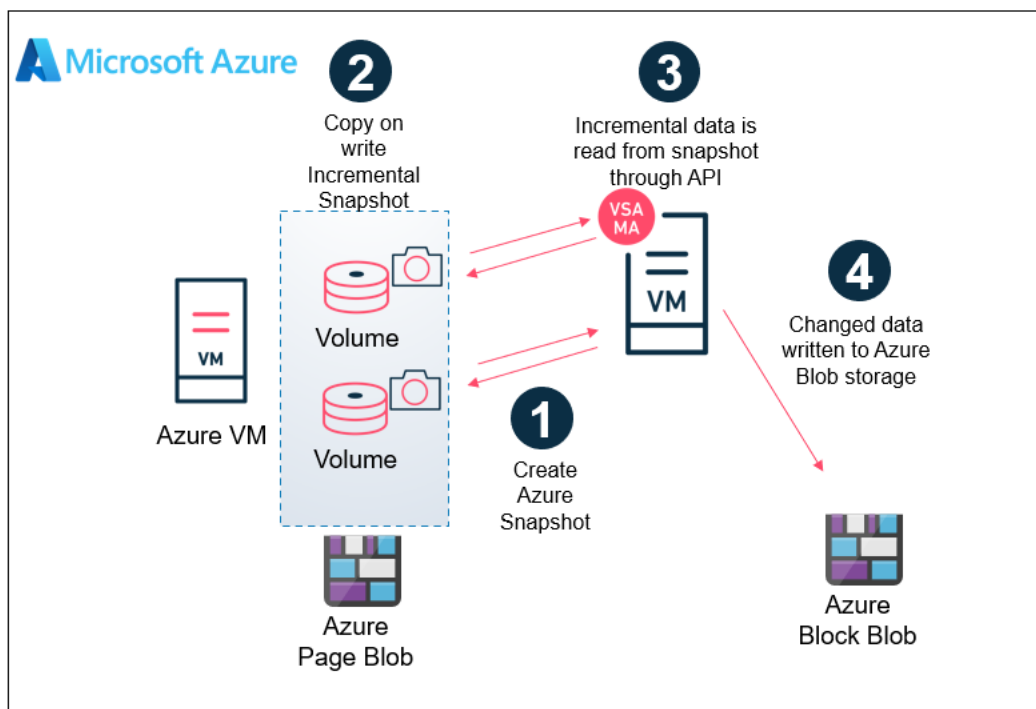
When not to use the VSA for Azure

- The VSA for Azure approach creates a crash-consistent image of the source VM and its block volumes. If you require application consistency, use an agent-in-guest either standalone or in conjunction with the VSA for Azure backup schedule.
- Protecting worker/stateless VMs – Worker nodes may generate valued data that is moved to another centralized repository and the nodes themselves do not require protection. It is recommended to instead target that centralized repository for data protection instead of the individual worker nodes, whether with VSA for Azure or agent-in-guest, depending on the required level of backup (crash vs. application consistent).

How VMs are qualified for protection

- Each VSA can be configured with one or more subclients. Each subclient defines a rule set on which to Auto-Detect and protect Azure VMs, based on a user-defined criteria of VM Name, Resource Group, Region, Storage Account, Tags, or Power state. During the Discovery phase of the backup job, the VSA will use the subclient rule to qualify VMs to add/remove for protection within that job.

Commvault software does not require access to the Azure hypervisor-level, instead using the REST APIs to create snapshots of each block volume, attaching (through APIs) the snapshot to a nominated Access Node (Azure VM-based VSA / MediaAgent) to read and de-duplicate the blocks before writing out to an Azure storage.



Commvault IntelliSnap® functionality introduced in V11 SP7 for the VSA for Azure. Commvault IntelliSnap® allows snapshots to be retained on the VM configured via Storage Policy's Snap Primary retention setting.

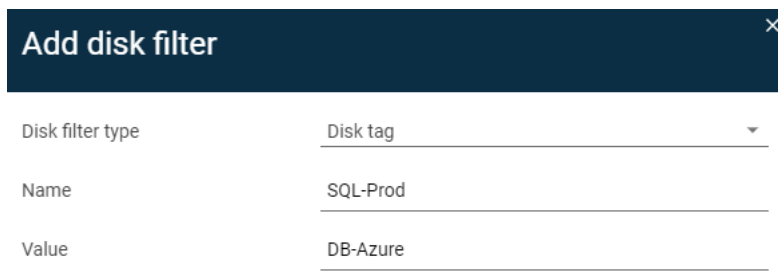
Commvault IntelliSnap® backups enable reducing backup windows considerably, providing fast, snapshot-based restoration capability. Azure also allows for the use of CBT (Changed Block Tracking) which can accelerate incremental backup performance considerably.

Azure changed block tracking (CBT) is available for Azure unmanaged and managed disks and are both supported by Commvault.

Get additional details [IntelliSnap for Microsoft Azure >](#)

Filtering Disks using Azure Tags

- Tagging is a feature used by cloud providers to simplify the tracking of cloud resources. Disk Tags are a user-defined name and value pair which can be assigned to a cloud disk. For example, you can apply a disk tag name of “SQL-Prod” and give it a value of “DB-Azure”. Tagging all disks with this pair will help organize and identify all SQL database disks that are in Azure.
- Commvault supports [Azure Disk Tags](#) and can be used during data protection operations to group/filter VM disks. The Tag Name or Tag Value that are added as filters will be used to search through the Azure subscription associated to the VSA client.



Disk filter type	Disk tag
Name	SQL-Prod
Value	DB-Azure

Get additional details [here >](#)

Architecture requirements for the VSA for Azure

- Minimum 1x VSA/MA per Region, Recommended 1x VSA/MA per Availability Set.
- Each “VSA/MA” node represents a single Azure VM with the Virtual Server Agent VSA and MediaAgent MA modules deployed. The Azure VM specifications should match the MediaAgent specifications within this Architecture Guide.
- When considering the VM instance type for the Access Node, the Commvault platform does not mount the guest VM snap disks on the Access Node for backup. This means that if you are protecting VMs with Premium Disks, you don’t necessarily need to select an Access Node VM type that is capable of mounting a Premium Disk. You can choose a cheaper VM type if it suits your environment better. Although keep in mind that there can be a performance hit if the Access Node VM type is not able to keep up with the required compute and IOPS demanded of it.

Architecture recommendations

- Use of the Commvault IntelliSnap® configuration is highly recommended to improve backup and restore times by enabling native Azure snapshot retention. Use of this method does mean that snapshots remain available for longer than the backup window, however Commvault recommends that snapshots are retained only for as long as required. The Snap Primary can be configured to retain at least one snapshot, keeping snapshot costs at a minimum while providing fast backup and restoration capabilities.

- By default, VSA backups of Azure VMs are crash consistent. To receive application consistent backups, leverage the [AppAware](#) feature.
- While the readers count can be increased to improve concurrency per VSA/MA node, consider scaling out with multiple VSA proxies. Azure recommendations mention that for optimal performance, you will want to limit the number of highly utilized disks attached per worker node to avoid possible throttling.
- Use of Premium Storage SSD for the deduplication database and index cache used by the MediaAgent module is highly recommended for optimal performance.
- (Backup Copy / Streaming) To restore from advanced Linux files systems such as EXT4, XFS and others, you can deploy a file recovery enabler by [converting an existing Linux MediaAgent to a FREL](#). When browsing data on advanced file systems, Commvault software will leverage the FREL to access the file system.
- (Backup Copy / Streaming) Enable [CBT](#) to improve incremental backup performance.
- (Backup Copy / Streaming) Disable Granular Recovery of files if granular recovery is not required, or agents-in-guest are used to collect large file system datasets. This will improve the backup window by removing the need to 'walk' the file system structure within the Azure VM volumes.

Performance tests in Azure using the VSA for Azure

Commvault has performed lab-based testing using a hybrid MediaAgent + Access Node configuration and tested the throughput and cost when using single larger nodes (i.e. medium size) vs. multiple smaller nodes (i.e. 2 x small). The results show that starting smaller (Small 1 Node) will minimize day one cost. As data volumes grow, and Recovery Point Objectives (RPOs) must be met across multiple hosts, a 'Medium 2 Node' offers great performance and resiliency. The cost increase of multi-node is observed in a drastic reduction of backup runtime.

Configuration	Backup Type	Data (TB)	Runtime (Hrs)	Throughput (GB/Hr)
Small (1 Node) 1 x D4s_v3 4 vCPU 16GB Ram	Full	9.92	31.18	325.56
	Incremental (70% new)	11.54	51.45	229.76
	Incremental (5% new)	0.498	2.72	182.96
	Synthetic Full	10.1	4.85	2128.79
	Restore	5.21	29.00	183.96
Small (2 Node) 2 x D4s_v3 4 vCPU 16GB Ram	Full	10.47	15.07	721.95
	Incremental (70% new)	9.45	14.27	689.89
	Incremental (5% new)	0.464	1.87	437.14
	Synthetic Full	10.23	6.25	1717.43
	Restore	5.18	17.65	294.85

Medium (1 Node) 1 x D8s_v3 8 vCPU 32GB Ram	Full	10.23	21.58	492.39
	Incremental (70% new)	9.38	24.00	447.62
	Incremental (5% new)	0.443	1.27	423.75
	Synthetic Full	10.21	4.42	2455.26
	Restore	5.05	24.00	294.85
Medium (2 Node) 2 x D8s_v3 8 vCPU 32GB Ram	Full	10.23	8.88	1215.06
	Incremental (70% new)	9.32	9.38	1046.00
	Incremental (5% new)	0.437	1.22	453.74
	Synthetic Full	10.21	6.93	1836.11
	Restore	4.95	11.62	423.32

Note: Randomized data was used to exclude deduplication benefits

Agent-in-guest (streaming)

An agent-in-guest approach can be used to protect a wide variety of operating systems and applications. These can be captured on the production workload and protected to the MediaAgent residing in Azure, using client-side deduplication to reduce the network consumption within the cloud. These can also be replicated to a secondary MediaAgent residing in a different geographic region. Once replicated the data can be held and restored in the event of a DR scenario or automatically recovered to existing VMs for the more critical workloads.

When to use agent-in-guest approach:

- When you require application-consistent backups – Deployment of agents can either be pushed/installed by Commvault software or baked into an Azure template using de-coupled installation or deployed as part of a continuous deployment method (i.e. Puppet/Chef/Ansible).
- When you require granular-level protection and restoration features for applications – the Commvault iDataAgents can deliver granular-level protection for supported application workloads, such as SQL Server or Oracle Database, in comparison to a Full VM or File-level approach.

Protect Azure SQL Databases

Commvault provides a complete data protection solution for Azure SQL databases by automating backup operations and by providing the following recovery methods:

- Restore from Azure SQL to Azure SQL
- Restore from Azure SQL to an on-premises MS-SQL server
- Restore from an on-premises MS-SQL server to Azure SQL

Best practice: The SQL server on the proxy client should run the latest SQL server release to ensure the proxy server is in sync with the Azure instance. Add the SQL server instance to the CommCell Console or Command Center. When you restore an on-premises database to the Azure cloud, the restored database uses the standard tier model.

Get additional details [here >](#)

Architecture requirements for agent-in-guest:

- Minimum 1x iDataAgent per VM for the intended dataset (i.e. SQL, File). Multiple iDataAgents can be deployed on the same VM.
- Minimum 1x MediaAgent per region. MediaAgents connect to the target object storage, and can either be deployed on the same VM, or on a dedicated host for a fan-in configuration. The Azure VM specifications of the MediaAgent should match the MediaAgent specifications within this Architecture Guide as a reference.
- Check the Systems Requirements section in Commvault documentation to determine if the iDataAgent supports your application ([Backup Agents](#)).

Azure snapshots

Azure snapshots allow for a crash consistent point-in-time copy of an Azure disk and can be automated with the use of Workflows. Snapshots are Copy-On-Write (COW) so it is recommended to keep more than one snapshot if you require the ability to restore from snap in the case of corruption or loss of data on the production VM. Commvault IntelliSnap[®] supports Azure managed disks and unmanaged disks.

With Commvault IntelliSnap[®] support for Microsoft Azure, you can:

- Perform full and incremental snapshot backups of virtual machines.
- Perform backup copy jobs from snapshot backups.
- Perform backup copy jobs with Changed Block Tracking (CBT).
- Perform full VM restores from snapshot backups, and full VM and file level restores (using Live Browse) from backup copy jobs.

Get additional details [here >](#)

Azure blob storage backup

Introduced in Version 11 Service Pack 6 is the ability to perform backups of Azure Blob storage. This capability will allow Commvault software to capture the data contained inside an Azure blob container, allowing for full or granular restore back to a blob container or file system client.

For more information refer to [Azure Blob Storage Overview >](#)

When to use Azure Blob storage backup:

- Backing up object storage – Protect data/objects created and stored in Azure Blob storage.

Architecture Recommendations

- For large datasets, consider using multiple subclients to increase scan performance and lower the amount of time taken to traverse and scan the bucket contents.
- Configure data operations for multi-streaming using multiple readers for best performance.
- To improve backup performance, we recommend you disable logging or to redirect the logging to another bucket in a user-defined subclient.
- Metadata is protected during backup

When not to use this approach:

- Protecting Commvault cloud libraries (Backup Data) – To protect Commvault backup data contained in cloud libraries, use secondary copies (DASH copies or [Aux Copies](#)) in the storage policy instead.

Azure File Share

Azure File Share is a feature derived from Azure Blob storage. Commvault can protect Azure blob storage and can also protect Azure Files. You can protect it from the Blob perspective and backup everything in the file/folder structure, or you can direct the connector to protect the File share folder of your choice, including the metadata at a folder level. You can back up an [Azure Files Share](#) by providing the Azure File Share URL (Host URL=**file.core.windows.net**) when you create an Azure File virtual client from the Command Center or the CommCell® Console. Once configured, you'll have Full and Granular restore capabilities of the data, in-place or out-of-place.

This feature is currently available for streaming backups and snapshots. [Azure File Share Snapshot](#) backup was added in FR 11.23 and offers protection of locked files by default. You can also choose to backup all files using Azure File Share snapshots.

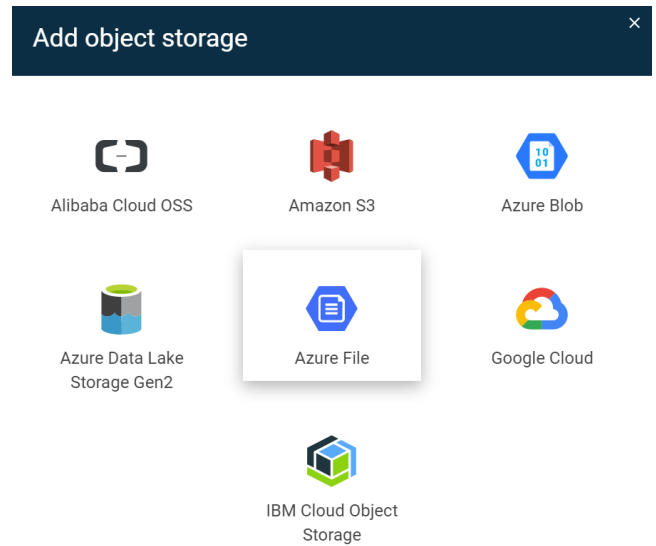
For more information refer to [Azure File Storage](#)

Architecture recommendations:

- For large datasets, consider using multiple subclients to increase scan performance and lower the amount of time taken to traverse the container contents.
- Multi-stream using multiple readers for best performance.

Application migration

Commvault can assist in application migration efforts when shifting from on-premises facilities to public cloud providers such as Microsoft Azure. By leveraging the power of the Data Management platform, workloads can be migrated through several methods.



Virtual Machine restore & convert (lift and shift to Azure)

The Virtual Server Agent can capture virtual machines from VMware, Hyper-V, Nutanix AHV, Oracle VM and AWS based platforms in an application-consistent method to ensure that a consistent image of the guest, and the applications residing within, are captured correctly.

With this image, the Virtual Server Agent can then restore and convert the virtual machine into Azure VMs directly, and the process can handle single or multiple virtual machines.

This process is performed interactively through the CommCell® Console, via Commvault Workflow or API calls.

Application out-of-place restore (all supported platforms)

All application iDataAgents support the capability to restore a given source dataset out-of-place to an alternate location. In this method, the data is captured from the source system (physical, or virtual), and then either directly from the source copy or replicated to cloud (DASH Copy), a restore to the destination is submitted.

The process requires the supported iDataAgent to be deployed on both the source VM, and the destination Azure VM.

This process is performed interactively through the CommCell® Console, via Commvault Workflow or API calls.

Deployment

Installation basics

The following links cover the steps when installing the CommServe® server in the cloud. This is only needed when the primary CommServe® server will be running on the hosted cloud VM or used for DR recovery. Multiple modules can be deployed in a single installation pass to streamline deployment.

- [Installation Overview](#)
- [Installing the CommServe®](#)
- [Installing the MediaAgent](#)
- [Installing the Virtual Server Agent \(Azure\)](#)

CommServe® Disaster Recovery solution comparison

As of v11 SP16, Azure storage becomes a 3rd candidate as an additional export destination. Learn more about CommServe® DR Solution comparisons for building a standby DR CommServe® server in the cloud, or simply restoring on-demand (DR backup restore), [Configuration of Disaster Recovery \(DR\) Backups](#)

Pre-packaging Commvault software within a VM template

For environments where deployment time is reduced by preparing software and configuration within VM templates, the Commvault iDataAgents can also be deployed in Decoupled mode. This means the iDataAgent is deployed within the VM but will only be activated upon registration with the CommServe® server.

For more information, please refer to the Installing the Custom Package instructions within Online Documentation:

- [Installing the Custom Package on Windows](#)
- [Installing the Custom Package on Linux](#)

Automating deployment with continuous delivery

For environments using Continuous Delivery toolsets such as Puppet, Chef or Ansible, Commvault® supports deployment methods that allow administrators to both control agent deployment and configuration to provide an automated deploy-and-protect outcome for applications and servers.

For more information on creating an unattended installation package for inclusion in a recipe, please refer to the Unattended Installation guide within Commvault® Books Online:

- [Unattended Installation](#)

For more information on using Commvault® software's XML / REST API interface to control configuration post-deployment, please refer to the online documentation links below to review options available for each iDataAgent:

- [REST API – Overview](#)
- [Command Line – Overview](#)

Cloud library configuration

This section covers the steps needed to configure cloud storage as a primary, secondary, tertiary, etc. storage target. Please keep in mind that use cases outside of archive will require Commvault infrastructure in the cloud to recover any protected data.

For most on-premises backup use cases (except for very small environments limited to 100 GB in payload size), cloud as a direct storage target for the primary copy is not recommended. For performance and responsiveness, a primary copy should be stored on an on-site disk library and a secondary copy should be hosted on the cloud storage. The secondary copy should be setup as an encrypted network optimized DASH copy to the cloud.

Verification of deduplicated data:

Data verification can be performed on deduplicated data to ensure the unique blocks on disk are consistent with the information contained in the deduplication database. Verifying deduplicated data ensures all jobs written to backend storage are valid for restore or auxiliary copy operations. While this is periodically recommended with on-premises disk libraries there are caveats when running a data verification when using cloud storage libraries. By default, data verification is **disabled** for cloud libraries because of the read operations from the cloud are slower than disk as well as the cost implications. These read operations are charged by the cloud vendors and can increase your cloud costs significantly. When the MediaAgent is on-premises and a data verification process is run on a cloud library, the MediaAgent will read all data in the library which will incur egress fees. **Azure egress fees** are charged per GB per month and as an example the first 5 GB/Month are free but from 5 GB to 10 TB the cost is \$0.0875 GB/Month. So if you have a storage library with 100 TBs of data and you run a data verification job from an on-premises MediaAgent, the cost will be approximately \$8,480 USD.

- [Supported Cloud Storage](#)
 - [Cloud Storage - Overview](#)
 - [Cloud Library Performance Tuning](#)
-

- [Verification of Deduplicated Data](#)

Unsupported cloud storage configurations

If a cloud storage target is not listed in the [Cloud Storage – Support](#) table, but the cloud storage endpoints are publicly accessible, and provide either an S3-compatible or OpenStack-compatible REST API, you can verify the compatibility of the storage offering with Commvault.

Depending upon your cloud device type you may choose to verify the compatibility between:

- [Supported Cloud Storage - Amazon S3 Vendors](#)
- [Supported Cloud Storage - EMC Atmos Vendors](#)
- [Supported Cloud Storage - OpenStack Object Storage Vendors](#)

For devices that are not publicly accessible, please contact your account manager or Commvault support for more information on the Cloud Storage certification process.

Architecture sizing

Remember the nature of cloud is to size the environment for today’s need, this is a change from traditional design practice to size for 3-5 years of growth. When sizing in the cloud, you need to adopt the mindset of ‘pay-per-use’ and grow for tomorrow. Commvault infrastructure components can be elastic to fit your environment’s needs. The tables below are suggestions to handle the upper limits of each size category.

Start with the smallest category in the tables below for Azure instance sizes to fit your current data protection size. Then upgrade them as needed to meet your environment’s requirements. (e.g., start with a Small MediaAgent, then upgrade to Medium MediaAgent as you onboard more client computers).

For the CommServe® server, adjust the instance size upwards when CPU and RAM loads become consistently high; add more space to the second volume as needed to accommodate the size of the CommServe® server database as you add more clients and jobs to the CommCell. CPU and RAM load can be monitored in the Azure Portal, or in the CommCell Command Center using the [Infrastructure Load Report >](#)

Azure CommServe® server specifications

Small (100 VMs or 25 Servers)	Medium (1000 VMs or 500 Servers)	Large (5000 VMs or 2500 Servers)	Extra large (20,000 VMs or 10,000 Servers)
<ul style="list-style-type: none"> • Standard_D8s_v4 (4 vCPU, 24 GB RAM) • 100 GB Premium Storage volume for CommServe database • Windows Server 2016 or Windows Server 2019 	<ul style="list-style-type: none"> • Standard_D8s_v4 (8 vCPU, 32 GB RAM) • 200 GB Premium Storage volume for CommServe database • Windows Server 2016 or Windows Server 2019 	<ul style="list-style-type: none"> • Standard_D16s_v4 (16 vCPU, 32 GB RAM) • 300 GB Premium Storage volume for CommServe database (P20 type) 2300 IOPS • Windows Server 2016 or Windows Server 2019 	<ul style="list-style-type: none"> • Standard_D32s_v4 (32 vCPU, 128 GB RAM) • 500 GB Premium Storage volume for CommServe database (P20 type) 2300 IOPS • Windows Server 2016 or Windows Server 2019

- Azure Marketplace
Windows Server 2019, 2016, or 2012 R2

Note: For more detailed information, please refer to the following link: [Hardware Specifications for the CommServe Server](#)

Azure MediaAgent specification

For MediaAgent servers, monitor CPU, RAM, and network I/O on the VM with the Azure Portal. Upgrade to a larger VM instance when you start encountering bottlenecks in any of these resources.

The recommendations below serve as a baseline. As environments continue to grow, inspect DDB performance health check to ensure it remains in a healthy state. If DDB health becomes non-optimal provide additional capacity to the DDB disk. This method increases allowable IOPs for the disk when additional performance is necessary. Get more detailed information [here](#), but remember the increased dedupe block size for cloud targets changes the backend supported size.

Commvault has identified the Azure [Ds_v4-series](#) VM as the recommended instance type due to the high performance CPU and Memory combination, but also its higher network bandwidth for faster data transfer.

Azure MediaAgent instance sizes			
Extra Small	Small	Medium	Large
BET: Up to 50 TB	BET: Up to 100 TB	BET: Up to 300 TB	BET: Up to 600 TB
<ul style="list-style-type: none"> • D2s_v4 (2 vCPU, 8 GB RAM) • 200 GB Premium Storage volume for DDB (P15 type) 1100 IOPS • 400 GB Storage volume for Index Cache (non- premium) 	<ul style="list-style-type: none"> • D4s_v4 (4 vCPU, 16 GB RAM) • 400 GB Premium Storage volume for DDB (P20 type) 2300 IOPS • 400 GB Storage volume for Index Cache (non- premium) 	<ul style="list-style-type: none"> • D8s_v4 (8 vCPU, 32 GB RAM) • 600 GB Premium Storage volume for DDB (P30 type) 5000 IOPS • 1 TB Storage volume for Index Cache (non- premium) 	<ul style="list-style-type: none"> • D16s_v4 (16 vCPU, 64 GB RAM) • 1.2 TB Premium Storage volume for DDB (P40 type) 7500 IOPS • 1 TB Storage Premium volume for Index Cache
<ul style="list-style-type: none"> • Azure Marketplace Windows Server 2019, 2016, or 2012 R2 Linux CentOS 7.4 / RHEL 7.5 (recommended) 			

Important: Some Azure VMs provide a single, fast SSD free of charge, however this storage is temporary, and all data is lost if the VM is moved to another host or is rebooted. For performance reasons, you can move the CommServe® TempDB SQL Database to this volume, and additional automation is required to ensure that on-reboot any required directory structures are re-created prior to SQL Server startup; otherwise, the SQL Instance (and the CommServe® services) will not successfully start. Commvault Professional Services can provide this bespoke automation as part of its solution delivery process. See [Using SSDs in Azure VMs to store SQL Server TempDB](#) for more information.

BET refers to a Back-End Terabyte, a TB measured within the Commvault backup system, once processed.

The Back-End Terabytes (BET) and Front-End Terabytes (FET) sizing for maximum capacity are based on a **512KB** deduplication block size, which is the default for writing to cloud libraries.

MediaAgent Grids

A single MediaAgent can service up to 1 PB of stored backup data before requiring an additional MediaAgent be added. MediaAgents may also be deployed in resilient grids which spread workload across multiple nodes and scale to larger protected data volumes (see below).

Grid	Extra large	Extra large	Large	Large	Medium	Medium	Small	Extra Small
	1 DDB Disk	2 DDB Disks	1 DDB Disk	2 DDB Disks	1 DDB Disk	2 DDB Disks	1 DDB Disk	1 DDB Disk
1 Node	Up to 500 TB	Up to 1000 TB	Up to 300 TB	Up to 600 TB	Up to 150 TB	Up to 300 TB	Up to 100 TB	Up to 50 TB
2 Node	Up to 1000 TB	Up to 2000 TB	Up to 600 TB	Up to 1200 TB	Up to 300 TB	Up to 600 TB	Up to 200 TB	Up to 100 TB
3 Node	Up to 1500 TB	Up to 3000 TB	Up to 900 TB	Up to 1800 TB	Up to 450 TB	Up to 900 TB		
4 Node	Up to 2000 TB	Up to 4000 TB	Up to 1200 TB	Up to 2400 TB	Up to 600 TB	Up to 1200 TB		

Source: [Hardware Specifications for Deduplication Mode](#)

For reference, Commvault has sized the MediaAgent grids based on the following standard backup retention:

- 80% VM/File data
- 20% Database data
- Daily change rate 2%
- Compression rate 50%
- Daily backups are retained for 30 days

This configuration results in roughly a **1.0 - 1.2x** increase from front-end to back-end storage, after deduplication and compression is applied.

Grid based configurations may be scaled both vertically (increasing CPU, RAM, and Disk) and horizontally (adding additional nodes of same specification). Each node configuration has a maximum amount of backend data that can be managed. If you plan on scaling your environment to 2-node, 3-node, or 4-node MediaAgent configurations, you can plan to ensure minimal cost scaling in future. Employ these architectural best practices to scale efficiently:

- Plan your MediaAgent deployment per Availability Zone
- Deploy a deduplication partition for the total number of nodes you want to support at maximum scale
- Each DDB partition may be deployed in a directory or dedicated volume on a single MediaAgent (day 1)
- When the MediaAgent exceeds its maximum FET or BET measurement, move the DDB partition (or volume) to a newly provisioned MediaAgent

If you do not know your **backup data** size, multiply your client data size by 1.6x for an estimate of back-end data for a 90-day retention period

Architectural best practices for MediaAgent sizing

- Perform sizing on a per availability zone basis
- Deploy MediaAgents across multiple Availability Zones (within a region) for resilience from AZ outage.
- (Note: if a MediaAgent is down, blocks previously deduplicated by that MediaAgent will be re-written during backup by one of the available MediaAgents)
- Start with the smallest MediaAgent that will meet your immediate needs (3-6 months)
- Pre-populate your MediaAgent with DDB partitions (or volumes) for the total number of nodes you plan to support within the grid.
- Distribute DDB partition data (or volumes) across deployed nodes to balance deduplication workload across the grid.
- Scale horizontally by adding additional MediaAgents when configuration exceeds [Commvault MediaAgent sizing specifications](#)
- [Optionally] Scale your MediaAgent(s) vertically if you simply need additional client data protected, but protection operations are completing within your business accepted backup window.
- Do not mix MediaAgent Instance Types within a Grid (i.e. a Medium and an Extra-Small)
- Do not distribute MediaAgents across regions. You will require a grid per region, per availability zone.
- Do not enable data verification on cloud storage libraries.

Access Node Sizing

For standalone Access Nodes (without the MediaAgent package), monitor the total Front-End Terabytes (FET) protected by a single Access Node and either scale-up the specification of an existing Access Node, or add an additional Access Node for more backup throughput to your MediaAgents. Remember that each VM instance provides additional volume bandwidth and network bandwidth, so scaling horizontally with smaller Access Nodes is recommended over vertical scaling.

Horizontal scaling

Most small virtual machines will be protected with a small Access Node. More backup parallelism may be achieved with more concurrent streams and volume mounts, by adding additional Access Nodes.

For larger VMs (> 15TB), with multi-TB disks or multiple volumes, scale-up access node to gain more IOPS and associated network bandwidth to complete streaming activity within the backup window.

Commvault recommends using small Access Nodes as a low-cost auto-scaling Access Node.

Automatic Access Node Scaling

As of Commvault v11 Feature Release 11.21 Commvault has the ability to perform [Automatic Scaling for Azure Access Nodes](#) during backup activities. Automatic scaling allows further cost optimization of your Azure VM protection infrastructure by only having Access Nodes running during protection activities.

Automatic Access Node scaling performs as follows:

- Commvault assesses the total size of the protection activities (number of hosts, TB to protect, estimated backup throughput)
- Commvault automatically spawns the required number of Access Nodes to complete the backup
- Access Nodes are Linux-based and run same release as the CommServe
- Access Nodes are automatically registered with Commvault for the period of the backup
- After protection activities completes, the nodes are shutdown, and then deleted
- For additional information refer to:
- [Automatic Scaling for Azure Access Nodes](#)

Access Node Specifications

These specifications are based on Hardware Specifications for Virtual Server Agent and refer to stand-alone Access Nodes that do not contain the MediaAgent software package.

Commvault recommends utilizing separate Access Nodes, which will send unique data hashes to a centralized MediaAgent and store data in Azure object storage for short/long-term retention. Commvault recommends compute optimized instances with premium SSD, high IOPS volumes. Testing has shown that many smaller instances vs. fewer larger instances benefit concurrent protection activities.

Azure Access Node Specifications		
Extra small	Small	Medium
5 – 10 FETB	10 – 25 FETB	25 – 50 FETB
Standard_F2s_v2 (2 vCPU, 4 GB RAM)	Standard_F4s_v2 (4 vCPU, 8 GB RAM)	Standard_F8s_v2 (8 vCPU, 16 GB RAM)
1 x 120 GB Premium SSD	1 x 120 GB Premium SSD	1 x 120 GB Premium SSD
Windows Server 2019, 2016 or Linux / CentOS 7.4 / RHEL 7.5, RHEL 7.6 (recommended)		

Source: [Hardware Specifications for Virtual Server Agent](#)

Access Nodes are only required during protection activities (backup, recovery) and therefore are not considered long-running infrastructure. By co-locating your Access Node with your MediaAgent, and using the Cloud MediaAgent Power Management, the running time of the Access Node can be reduced to predominantly only your nightly backup windows.

Alternatively, if the number of parallel backup activities exceeds the number of MediaAgents available, utilize Azure Access Node auto-scaling to auto-create and remove Access Nodes on demand. It is more cost effective to deploy multiple smaller Access Nodes, vs. scaling a single Access Node.

Virtual Server Agent specification

The Virtual Server Agent (VSA) can be installed on the same VM as the MediaAgent. However, depending on the scale and design of the environment, utilizing separate VSA access nodes/proxies may provide increased performance or flexibility. In this situation, each VSA Access Node will send their data to a centralized MediaAgent.

These specifications are based on [Hardware Specifications for Virtual Server Agent](#).

Commvault recommends Azure [Ds_v4-series](#) instances with **Premium Disk** volumes. Testing has shown that more smaller instances vs. fewer larger instances benefit concurrent protection activities ([requirements](#)).

Azure Virtual Server Agent (VSA) instance sizes			
Extra small	Small	Medium	Large
FET: 1 – 10 TB	FET: 10 – 30 TB	FET: 30 – 60 TB	FET: 60 – 125 TB
<ul style="list-style-type: none"> • Standard_D2s_v4 (2 vCPU, 8 GB RAM) 	<ul style="list-style-type: none"> • Standard_D4s_v4 (4 vCPU, 16 GB RAM) 	<ul style="list-style-type: none"> • Standard_D8s_v4 (8 vCPU, 32 GB RAM) 	<ul style="list-style-type: none"> • Standard_D16s_v4 (16 vCPU, 64 GB RAM)
<ul style="list-style-type: none"> • Azure Marketplace Windows Server 2019, 2016, or 2012 R2 Linux CentOS 7.4 / RHEL 7.5 (recommended) 			

Best practices:

- If the Azure subscription includes multiple regions, deploy at least one VSA Access Node per region
- Deploy the VSA Access Node and MediaAgent on virtual machines in the Azure cloud
- Deploy the VSA Access Node on an Azure VM that is optimized for I/O intensive workloads to support faster backups
- Enable Azure accelerated networking on the VSA Access Node /MediaAgent machines in Azure. This step must be completed at the time of deploying the virtual machine.
- Enable service endpoints for Microsoft Storage on the Azure virtual network subnet where the Access Node and MediaAgent are connected. This will ensure that all network traffic from the Access Node machine to the Azure storage account is securely flowing through the Microsoft Azure backbone network.
- Enable Changed Block Tracking (CBT) for Azure. CBT for Azure provides better backup performance than traditional cyclic redundancy check (CRC) backups. You can use CBT with unmanaged and managed disks.

Example of sizing in cloud

Assumptions & costing

For this example configuration, we want to protect approximately 100 TB of Front-End (FET) capacity. The average size of each virtual machine instance is assumed to be about 100 GB and each instance has 2 volumes - one for operating system and the other for applications and data. This equates to approximately 1000 VM (100 TB at 100 GB each, using base10 rather than base 2 for simplicity in approximation).

It is also assumed that the daily change rate is ~2% of net new data that is created per day, or ~2 TB worth of new data before it is compressed and deduplicated. The change rate in an environment varies greatly by environment and coupled with retention and deduplication ratio of the data. Both of which are also highly dependent on the specific environment, these three factors affect the back-end storage capacity that is ultimately required. For this example, we shall assume 90 days of retention and ~5:1 deduplication ratio. Both are typically observed within most virtual environments running a mix of Windows and Linux operating systems. With retention it is important to note that data that is 90 days old or the first full backup are not deleted until the most recent one is fully committed to the backup target. This accounts for retention+1 storage requirement. This results in approximately 117 TB for the backend.

It will also be assumed that absolutely no infrastructure to manage the cloud protection environment is present outside the cloud and that Commvault cloud MediaAgent power management feature is enabled, enabling shutdown of resources when backups and recoveries are not occurring. While most backup windows are usually 8 hrs., the assumption is that with restore account for another 4 hrs. per day allowing for power management to operate for only half of a given day.

Using publicly available pricing for Azure resources the cost of performing protection in Azure by utilizing any combination of iDataAgents, agentless VSA based backups, and coupled with Commvault IntelliSnap® for Azure snapshots, the following becomes a rough estimate of the cost of the Azure infrastructure required for a period of 1 year:

Commvault component	Qty.	Azure type	Azure cost/hr or /GB	Azure cost/year	Azure cost/year w/poweroff
CommServe® Medium VM	1	D8s_v3	\$0.597	\$5,229.72	\$5,229.72
CommServe® OS Disk	1	S15_256GB	\$0.016	\$137.53	\$137.53
Standard Dedupe MA Medium VM	1	D8s_v3	\$0.597	\$5,229.72	\$2,614.86
MA + VSA OS Disk	1	P20_512GB	\$0.102	\$890.89	\$890.89
MA DDB Disk	1	P30_600GB	\$0.102	\$897.02	\$897.02
MA Index Disk	1	P30_1TB	\$0.171	\$1,497.96	\$1,497.96
Disk Library (TB) (Full 90 Days)	230	Cool Blob	\$0.01	\$14,568.65	\$14,568.65
Totals:				\$28,959.58*	\$26,090.68*

*It must be noted that this is a sample configuration utilizing estimated sizing data and that actual costs will vary depending on data type, retention, and other factors. This assumes scaled up to 100 TB FET, starting with a much smaller footprint and growing as the source grows is perfectly acceptable and recommended.

Additional resources

Documentation

Online documentation – Cloud Storage

The Cloud Storage section from Commvault's Online Documentation covers technical procedures and information on Supported Cloud Targets, Advanced procedures, Troubleshooting and FAQ sections for Commvault customers.

[Learn more >](#)

Datasheets

- [Commvault for Microsoft® Azure Cloud](#)
- [Go ahead, move your most important applications and databases to Azure](#)
- [Accelerating data migration to Azure](#)
- [We manage and protect mission-critical apps running on Azure Ultra SSD](#)
- [The right choice for enterprise class recovery of SAP workloads on Azure](#)
- [Commvault and Azure NetApp files](#)
- [Thinking of managing your open source databases on Microsoft Azure PaaS?](#)
- [Microsoft Datasheet: AS/400 Backup to Azure](#)
- [A checklist for cloud data management](#)
- [Microsoft trusts Commvault. Shouldn't you?](#)

Solution briefs

- [Rapid data migration for Azure](#)
- [Azure Stack Backup and Recovery](#)

Videos

- [VMware to Azure Migrations With Commvault](#)
- [Virtual Machine Backup Demo Using the Commvault Command Center](#)
- [Keys to Fast, Secure Azure Migration With Data Protection](#)
- [Shikun and Binui: Cloud Data Protection Across Azure, Office 365 and SaaS environments](#)
- [Commvault Customer Champions: UConn Health](#)

Case studies

- [University of Central Florida](#)
- [HarperCollins](#)
- [Yura Corporation reduces costs and boosts business continuity](#)
- [Chart Industries](#)
- [UConn Health](#)
- [Laing O'Rourke](#)