



Protecting and Managing SAP HANA in Azure with Commvault

Contents

Protecting and Managing SAP HANA in Azure with Commvault.....	1
Introduction	3
Commvault/Microsoft Partnership.....	3
CommCell Environment Overview	4
Integration with Azure.....	5
Solutions for SAP workloads.....	6
Commvault Public Cloud Architecture Guide for Microsoft Azure.....	6
Protecting SAP HANA in Azure with Commvault.....	7
Test Environment.....	7
Cloud Libraries	7
Storage Policies.....	7
SAP HANA setup.....	8
Testing Results	9
Backup and Restore with Deduplication and Compression	9
Backup and Restore with Compression only.....	9
Backup and Restore without any data reduction.....	10
Summary of testing results.....	10
Storage Requirements.....	12
Reference Architecture.....	14
Storage requirements.....	14
Commvault infrastructure.....	15
Summary and Conclusion.....	16

Introduction

SAP customers looking into managing and migrating their SAP environment to Azure need a solution which can meet or exceed requirements for SAP data management and protection, cost reduction, improved efficiency, as well as remaining compliant with current SLAs. Commvault and Microsoft are leveraging their 20 year partnership and have teamed up to put together, test and document this solution.

Commvault/Microsoft Partnership

Microsoft and Commvault forged a relationship that now spans nearly two decades, delivering industry-leading data management solutions (Commvault) on a powerful, secure infrastructure (Windows Server & Azure). Commvault has built a single platform that unifies and automates data management and protection across all the operating systems and applications your organization relies on every day — simplifying the way you manage, protect, access and share data. Commvault is a leader in data storage and management and is also the first to enable agentless backup and recovery of Azure instances.

Not only have Commvault and Microsoft teamed up to create this SAP document, we're also leveraging that two-decade partnership which includes each company heavily using the others' technology to help run our day-to-day businesses. By leveraging our combined vast experience and scale, as well as our exposure to early-access technology, leads to proven best-practice guidance for our customers.

Commvault has supported Microsoft Azure since 2008, delivering a deep depth and breadth of support for Azure compute and storage that uniquely positions our joint customers to benefit from on the best practices of thousands of clients. Commvault is used extensively by Microsoft within the Office Products Group (since 1999), Xbox and Xbox Live (since 2002), as well as the Azure Data Protection Group (since 2013). Commvault is also a large Azure customer (since 2008), which included having Azure as the exclusive host for Commvault's SaaS offering. With support for hot, cool, page and cold (archive) Blob storage, as well as Express Route, Import/Export, Data Box, Geo-Replication and Azure Stack, we can help deliver an extremely optimized and flexible Azure solution that is agile enough to move with the high demands of digital transformation. Commvault data management software contains a full range of data management and data protection capabilities to/from the Azure cloud and Azure Stack, including best-of-breed data management capabilities for Microsoft products such as Exchange, SharePoint, SQL, Active Directory and Office 365.

In addition to the support of these Microsoft products, Commvault can also be used to manage and protect SAP HANA or any other mainstream database within Microsoft Azure, on-premises or in a hybrid cloud scenario.

Running SAP Business Applications based on SAP HANA in Microsoft Azure is becoming more and more popular and requires an enterprise-grade data management solution. Commvault is "Azure Certified" and "SAP-certified for the integration with SAP HANA" making the Commvault data management solution the right choice to manage, move and protect SAP HANA workloads in Microsoft Azure.

Commvault provides a SAP HANA data protection solution that is tightly SAP-integrated and is designed for universal protection. Not only standalone SAP HANA deployments, but also SAP Business Suite Powered by HANA, SAP BW/4HANA, SAP C/4HANA and SAP S/4HANA environments can be protected effectively. HANA can run in scale-up and scale-out setups, high availability configurations and HANA System Replication is supported as well.

Along with safe-guarding SAP data, Commvault provides a SAP archiving solution and also integrates with SAP Landscape Management (LaMa) tool for 1-click automated refresh of SAP test and development systems from production data. This allows for effective management of SAP landscapes running on SAP HANA and other databases supported by SAP.

As a market leader, Commvault has been named a leader in [Gartner's Magic Quadrant for Data Center Backup](#) and Recovery Solutions for 7 years in a row and has similar positioning from Forrester's Data Resiliency reporting.

SAP on Azure offering

Microsoft Azure is the best public cloud for running SAP workloads. With the largest portfolio of [SAP HANA certified IaaS Cloud offerings](#) customers can scale-up production certified virtual machines from 192GB to 4TB and run HANA Large Instances up to 24 TB of memory within one server. In addition, customers can scale-out OLAP workloads such as BW on HANA and BW/4HANA on virtual machines with 16 nodes (2 TB each) and run HANA Large Instances with up to 60TB scale-out configurations.

The partnership of Microsoft and SAP began more than 20 years ago when SAP certified Windows and Microsoft SQL Server for running SAP R/3. The deep integration of SAP with Microsoft Office (Office 365) and the possibility to connect the core SAP system with Microsoft Azure PaaS offerings such as Azure Active Directory for security and identify management and Analytics and Internet-of-Things (IoT Hub, Power BI) services emphasizes the strong integration capabilities of SAP and Microsoft technologies. Moreover, SAP SE themselves are now running several of their [internal SAP systems on Azure](#).

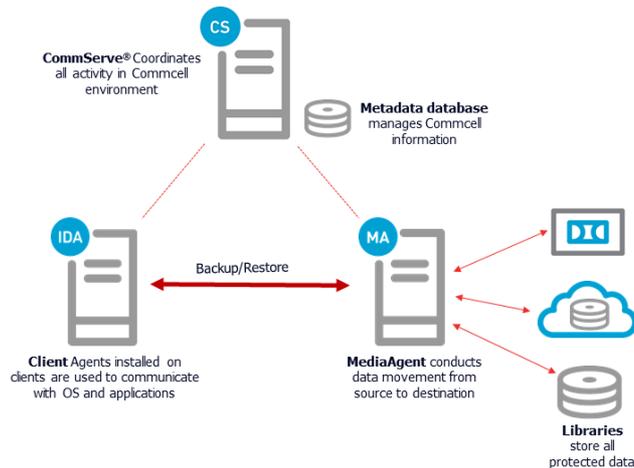
SAP and Microsoft provide a joint support for all SAP workloads using Microsoft technologies; this provides customer confidence that they can run their SAP workloads in the most trusted and secured cloud environment. Microsoft is not only the best cloud for running SAP workloads, Microsoft also knows how to run SAP; SAP ECC is the core ERP for running Microsoft's line of businesses and one of the largest SAP systems worldwide. Microsoft's transformation project to S/4HANA is also now underway. SAP also provides SaaS services to Microsoft such as Ariba and Successfactors and in many cases these services are provided out of Azure datacenters to SAP customers worldwide.

Commvault Architecture

CommCell Environment Overview

Before explaining how Commvault can be used to protect SAP HANA workloads in Microsoft Azure, here is a short overview of the main Commvault components.

A **CommCell environment** is the logical grouping of all software components that protect, move, store, and manage data and information. A CommCell environment contains one CommServe host, one or more MediaAgents, and one or more clients.



The **CommServe** host is the central management component of the CommCell environment. It coordinates and executes all CommCell operations, maintaining Microsoft SQL Server databases that contain all configuration, security, and operational history for the CommCell environment. There can be only one CommServe host in a CommCell environment.

The **MediaAgent** is the data transmission manager in the CommCell environment. It provides high performance data movement and manages the data storage libraries. The CommServe server coordinates MediaAgent tasks. For scalability, there can be more multiple MediaAgents in a CommCell environment.

A **client** is a logical grouping of the software agents that facilitate the protection, management, and movement of data associated with the client.

An **agent** is a software module that is installed on a client computer to protect a specific type of data. Different agent software is available to manage different data types on a client, for example, Windows file system data and SAP HANA databases.

The Virtual Server Agent or VSA is a specialised agent that protects hypervisor and cloud resources.

Integration with Azure

The primary integration point with Azure is through Commvault's Virtual Server Agent (VSA). You can use VSA to perform the following tasks for Azure VMs:

- Backup and recover Azure virtual machines using either the Azure Classic or the Azure Resource Manager deployment model. You can restore full virtual machines or guest files and folders.
- Seamlessly convert backups of Amazon, Hyper-V and VMware virtual machines to Azure virtual machines (Azure Classic or Azure Resource Manager).
- When performing a restore from a backup of an Azure VM, you can choose to restore a VM disk and attach it to a different VM that already exists.

Commvault data management and protection operations are managed through the CommCell Console. Use the CommCell Console to configure a virtualization client and other entities that are used to support operations.

A virtualization client instance is the access point for an Azure subscription and is used to backup full Azure virtual machines. You must define a VSA agent instance for each Azure subscription. If a database or application runs inside the VM on the Linux OS, the VSA approach is not appropriate, as it guarantees crash consistent backups only. For SAP HANA workloads we need

application consistent backups. Therefore we are using the “agent in guest” approach here with the the Commvault SAP HANA agent installed inside the VM. VMs just running SAP application instances (DI, CI, ASCS, etc.) can be protected using the VSA.

When you create a virtualization client instance, the Commvault software automatically creates an Azure instance, a so-called backup set, and a default subclient that can be used to protect all virtual machines. You can create additional subclients to perform separate protection operations for different groups of virtual machines. For example, you can create a different subclient for different guest operating systems and use the default subclient to protect any remaining virtual machines that are not covered by user-defined subclients.

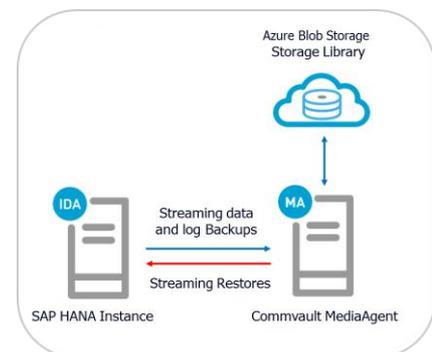
With Commvault VSA agent you can perform full, incremental, or synthetic full backups of virtual machines and perform full virtual machine, disk or guest files and folders at granular level.

Solutions for SAP workloads

SAP HANA uses an In-Memory Database technology that processes massive amounts of real-time data in a short span of time. Using the In-Memory computing engine, HANA processes data stored in RAM as opposed to reading it from a disk. This leads to instantaneous results from customer transactions and data analyses.

Commvault software provides a SAP-certified end-to-end backup and recovery solution for SAP HANA scale-up and scale-out environments using the SAP Backint for HANA interface.

For SAP HANA in Microsoft Azure, Commvault currently supports streaming backups from the SAP HANA node where the agent is installed and sends the backup data to the centralized MediaAgent; it will then store the backup data in the Cloud Storage provider backed by Azure blob storage. Alternatively, it is also possible to install a MediaAgent directly into the SAP HANA VM. You can do this on all nodes in a scale-out configuration. This way, all nodes can backup their local data in parallel to shared blob storage, enabling high-speed backups.



This provides the ability to easily increase storage capacity as and when it is required. Cloud Storage provides centralized data access, better failover capabilities and reduces the day-to-day storage administration tasks. As the data gets transferred over the network, protecting the integrity of data is an important aspect of any cloud storage implementation.

Azure cloud storage protects the integrity of the data using the following features:

- By default, data is transferred through secured channels using HTTPS protocol.
- Data encryption further encrypts the data providing data protection during network transfer as well as storage.

Commvault's Deduplication feature identifies and eliminates redundant data in the backup, thereby reducing not only the volume of data stored in cloud, but also the bandwidth required for data transfer.

Commvault Public Cloud Architecture Guide for Microsoft Azure

Commvault and Microsoft are providing this document as an architecture guide for solutions architects and Commvault customers who are building data protection and management solutions utilizing the Microsoft Azure public cloud and the Commvault platform.

It includes cloud concepts, typical use cases, architectural considerations, and sizing recommendations to support Commvault's platform in Azure. The approach defined in this guide extends existing functionality into easily saleable, re-usable architecture patterns to cover migration to the cloud, disaster recovery to the cloud as well as protecting running workloads in the cloud.

For additional reading, we recommend Commvault's Cloud Architecture Guide for Azure. For sizing, we would like to refer you to the "Architecture sizing" section of the document. The full guide can be downloaded here: <http://bit.ly/2v1BN4>

Protecting SAP HANA in Azure with Commvault

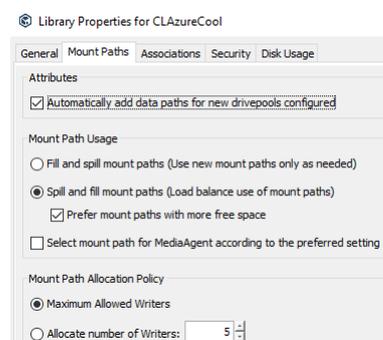
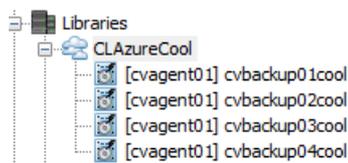
Test Environment

For testing and benchmarking, a test environment was built. This environment consisted of the following Azure resources:

- 1 Standard_D16_v3 VM, for running the Commserve on Windows 2016
- 300 GB of Premium Storage
- 2 Standard_D32_v3 VMs, for running Media Agents on Windows 2016
- 2x 400 GB of Premium Storage for holding the IndexCache and Deduplication databases
- 1 M128ms VM, SAP HANA 2.0 SPS03 Scale-up node running SuSE SLES 12 SP3
- 4 Blob Cool and 4 BLOB Hot Storage Accounts to serve as backup targets

Cloud Libraries

All storage accounts were configured as Commvault cloud libraries consisting of four mount paths each. All test backups will end up in these libraries eventually. In order to achieve maximum throughput and evenly distributed load balancing, the "Spill and Fill" option was activated.



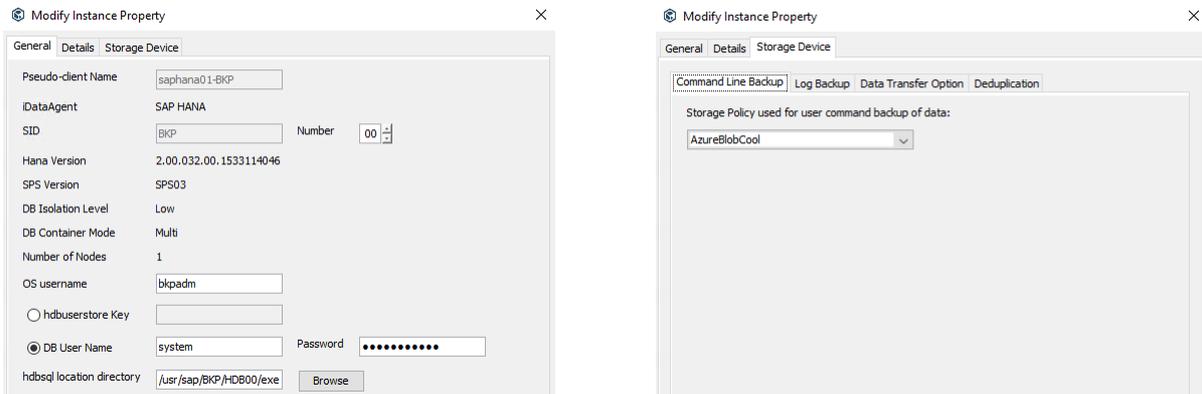
Storage Policies

Using Storage Policies we can define the data management rules which will be applied to the SAP HANA backup data. Storage Policies determine how the data is backed up (deduplicated, compressed or without data reduction), how many copies will be created, where these copies are stored (on the cloud libraries in our case) and how long each copy will be retained. We created Storage Policies for serving 3 test cases:

- HANA Backup and Restore with deduplication and compression
- HANA Backup and Restore with compression only
- HANA Backup and Restore without any data reduction

SAP HANA setup

To perform set up in Commvault, we start with the HANA instance whose SID is named BKP. It is required to specify the instance details alongside with either an hdbuserstore key or a database user so we can connect to HANA. Further, all HANA nodes (in this case just one) need to be added to the "Details" tab and Storage Policies for HANA logs and data needs to be specified. For HANA logs, we configured a separate Storage Policy for HANA logs with just compression enabled, as it doesn't make sense to deduplicate these. Deduplication and compression settings can be fine-tuned using this GUI dialog as well.



While the instance is generated in the Commvault GUI, the HANA agent will also auto-detect all tenants which in turn inherit all settings of the embracing HANA instance. All SAP HANA side configuration changes were already taken care of automatically during agent installation and can be verified through SAP HANA Studio.



To achieve real-world testing results, we loaded 1.5 TB of test data into the SAP HANA BKP tenant database. An important aspect for the testing effort, as well as for customer environments, is the ability to tune backup and restore performance by taking advantage of SAP HANA's multi-streaming capabilities. SAP HANA databases consist of up to four internal services (name server, index server, XS engine, etc.). Each of these services stores its own data and generates its own log files. However, the vast majority of data always resides in the index server. By default, all SAP HANA services are backed up in parallel, each using a single stream. All other HANA services typically finish their data transfer within minutes, while index server backups can run for many hours.

SAP also supports multi-streaming capabilities on the level of index server. This can be configured by specifying the desired number of streams using the global.ini parameter "parallel_data_backup_backint_channels" in combination with the parameter "data_backup_buffer_size", which needs to be set to number of streams multiplied by 512. The screenshot below shows an example of 16 streams.

SYSTEMDB@BKP (SYSTEM) saphana01 00

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

Filter: data

Name	Default	System	Database - BKP
esserver.ini			
database			
global.ini			
backup			
data_backup_buffer_size	512	8192	8192
data_backup_max_chunk_size	0		
data_backup_parameter_file		/usr/sap/BKP/SYS/glo...	/usr/sap/BKP/SYS/glo...
data_backup_savepoint_lock_timeout	7200		
es_data_backup_buffer_size	8		
parallel_data_backup_backint_channels	1	16	16

Please refer to Commvault's Best Practice Guide for SAP HANA for further information and agent installation, configuration and operation.

Testing Results

All testing was done based on SAP HANA full backups.

Backup and Restore with Deduplication and Compression

For the first round of tests we enabled deduplication and compression at the client side. Deduplication provides an efficient method to store data by identifying and eliminating duplicate blocks of data during backups. We enabled deduplication on the client side to make sure the least amount of data will be sent over the network, and to ensure we consume only the smallest amount of object storage, since only unique chunks will be stored.

During the test we have configured multiple streams for the SAP HANA backup and sent the result to Hot and Cool Azure Blob storage. In this setup we used a central deduplication database (DDB) on one of the media agents alongside a small local DDB cache for efficient signature testing on the SAP HANA node. The central DDB with IndexCache was installed on premium storage in order to warrant best performance, and data compression was executed on the client as well. After a couple of test runs, data reduction went up to 99%. This is due to the fact that our tests did not include data changes. Achievable data reduction rates depend on the amount of changed data and are typically lower.

The tests immediately delivered backup throughput beyond 2.5 TB/h resulting in backup job runtimes around 30 minutes for our 1.5 TB test system. 16 streams proved to be a good setting here. Client-side CPU load was a bit elevated, which should be fine given the quick job runtime. If this is still an issue for a critical production system, a backup window needs to be defined at a time with minimal SAP end user and batch system load.

We found that restores from deduplicated backups didn't run at the speed of the corresponding backups. This is because of their non-linear read pattern from the backup media. Customers need to be aware of this in the context of RTO. We also want to make clear that we measured full RTO time, which not only includes the actual data restore, but also the time needed to roll-forward the database based on the logs and for shutdown and startup.

Another interesting observation is that there is no notable performance difference between storing the backup data on BLOB Hot vs. BLOB cool storage. We therefore focussed on BLOB Cool only for the remaining test cases and recommend this option for customer deployments.

Pros:

- Fastest backup option
- Best data reduction

Cons:

- Restores don't achieve backup performance

Backup and Restore with Compression only

For this test, all deduplication was disabled (client-side and media agent) by using a storage policy without deduplication. Client-side compression remained in place. CPU load on the SAP HANA side was lower compared to the the deduplication test.

However, more CPU resources were consumed on the media agent. Using 16 streams, the backup performance was still beyond 2 TB/h.

On the restore front we saw improved performance numbers. The gap between backup and restore speed is becoming very small in this setting, as we are able to read the data chunks sequentially from the cloud library. RTO of less than an hour (including shutdown and startup of the database) for our test setup was achieved quickly. However, for satisfying highest RTO requirements the best solution is to use local Premium Storage on the HANA as backup library and run all restores from there.

Pros:

- High restore performance
- Cost savings as less premium disks are needed
- Best option for production systems

Cons:

- Data reduction to ~30%
- Higher cost for secondary BLOB storage

Backup and Restore without any data reduction

Now we disabled compression, as well as deduplication completely. All data was transferred in original size without data reduction. Compared to the tests with compression only, backup performance was lower with a 16 streams setup. In a further test with 24 streams it was possible to achieve more than 2 TB/h. As no backup data needs to be compressed or deduplicated anymore, CPU consumption on the client-side was pretty low. CPU utilization on the media agent was again elevated.

Restore performance was again a bit lower than backup, but leaving a small gap. However, compared to the results with compression enabled, restore speed was slower. When compared to the results with deduplication and compression, performance was certainly better. CPU loads were basically the same as for backup, which also makes sense.

Pros:

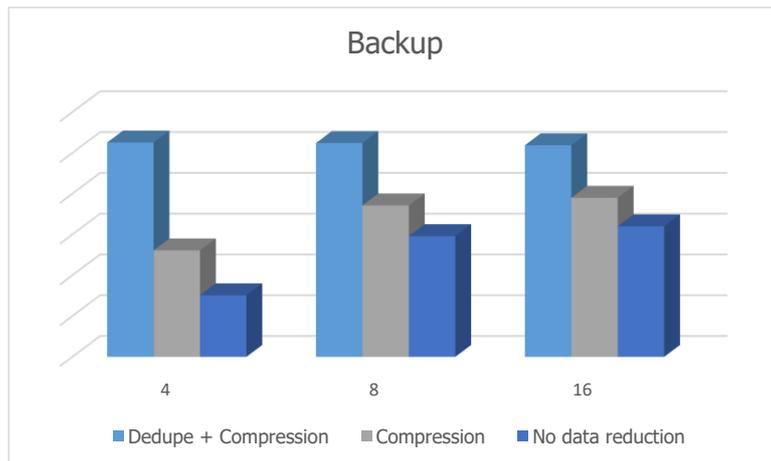
- Cost savings as less premium storage is needed
- Backup speeds beyond 2 TB/h possible
- Good option for short RTO

Cons:

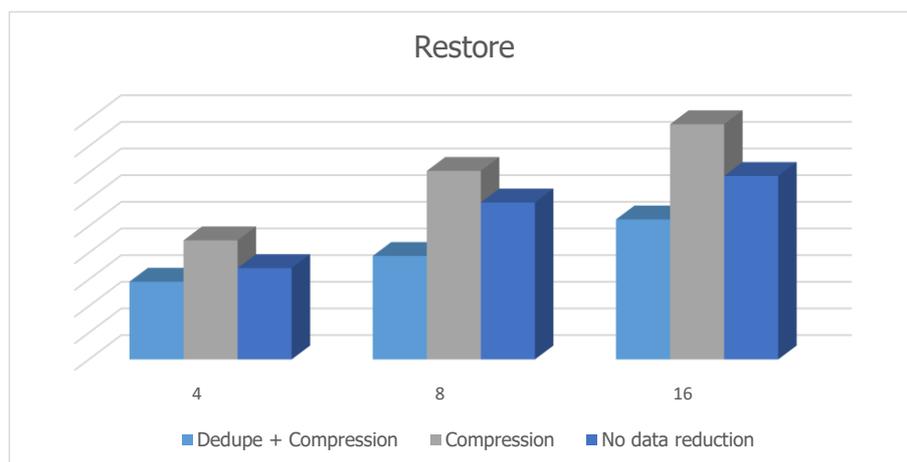
- No data reduction at all
- Highest secondary BLOB storage requirements

Summary of testing results

As the tests have shown, using deduplication in combination with compression primarily helps to reduce BLOB secondary storage consumption (around 70% in practice, depending on data change rate) and is certainly the fastest backup option at well over 2.5 TB/h. However, restores were much slower, making this option less than ideal for critical production systems. CPU utilization at backup time was also elevated.



In cases where we disable deduplication and just use compression, we can observe backup performance going down a bit, but still saving around 30% of BLOB storage space. CPU load is actually moving away from the HANA VM to the media agent. This test case has good restore performance, meaning that RTO of less than an hour is achievable for our test database. Even just one hour or less of SAP downtime can have a huge impact on costs and business continuity when trying to restore a critical production system.



The final test with deduplication and compression being disabled showed no further improvements. Backup and restore performance were actually lagging behind the first two test cases, while storage consumption went up.

To summarize, the recommendation is to look at RTO before making a decision on which of the three backup methods should be used. For a SAP production system, it might be best to use client-side compression only and without deduplication. Bringing back a 5 TB system in 3 hours is possible. Backups can run for longer times as CPU resource impact is still moderate. For test, development and sandbox systems, RTO requirements are usually lower, so deduplication can be used for minimizing the BLOB storage footprint.

Using deduplication along with the previously introduced VSA agent is also the recommended approach for protection all SAP application instances.

If BLOB secondary storage consumption is considered too high, you can combine SAP HANA fulls with differential and incremental backups. In this case you run just 1-2 fulls per week and incrementals and differentials in between, which brings down storage requirements. The following section compares the two options with regards to BLOB storage requirements.

Storage Requirements

Storage requirements for the backup of SAP HANA workloads depends on the backup schema and retention.

The SAP HANA agent backs up the following:

- All database files on each node
- Log files on each node
- SAP HANA backup catalog

For SAP HANA data the following backup types are available:

- Full
- Incremental
- Differential

In order to protect the SAP application instance footprint, the Commvault file server agent which is part of the installation is recommended to be used; another option is using the VSA agent.

- A typical customer with a retention requirement of 2 weeks, running all backups at a full level, would have 14 full backups stored in their Azure Blob storage based cloud library. If the customer descides to run only one full backup per week and incremental or differential backups on the other days, less storage would be consumed. Incremental backups would store all changes since the last backup and the differential stores all changes since the last full backup.

So if performing one full backup per week, a differential backup half way through the week and incremental on all other days, we would see less storage consumed, but an additional cycle will be retained because the incremental and differential jobs rely on their parent backups. This is automatically orchestrated by Commvault as per table below where in the full-only backup cycle the backup of day 6 will be expired on day 20.

Backup cycle	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Only full	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F
Levelled backups	F	I	I	D	I	I	I	F	I	I	D	I	I	F	I	I	D	I	I	I

When using the full-only backup cycle, there will be fewer jobs retained, but since the full backups are bigger than the incremental and differential backups, more Azure blob storage would be consumed. This is illustrated by the example below.

In this example we consider a larger daily change rate (churn) of 10% on the 1.5 TB test database and an average compression ratio of 30% (deduplication has been excluded from this example).

The following amount of Azure Blob storage would be used:

	Only Full	Levelled backups
Full	14.7 (14*1.05 TB)	3.15 (3*1.05 TB)
Incremental	-	1.47 (14*0.105 TB)
Differential	-	0.945 (3*0.315 TB)
Total TB Azure Blob storage consumed	14.7	5.57

Note:

Log file and SAP application instance files are not considered in this calculation:

- Log files backed up with compression only in both scenarios and will have the same retention since they are retained against the full, incremental and differential backup.
- SAP application instance files can be backed up using a storage policy for regular file system or VSA backup with compression and deduplication enabled

Reference Architecture

Now it's time to build a reference architecture for a typical customer environment. Let's consider this example SAP environment consisting of four application landscapes. Depending on the workload use case (PROD, DEV, TEST) we will have different Restore Time Objectives (RTO) and retention requirements to complete the example.

	SAP S/4HANA	SAP BW	SAP CRM	SAP PLM	RTO
Production	1500 GB	3000 GB	800 GB	500 GB	1-3h
Development	1500 GB	3000 GB	800 GB	500 GB	5h
Test	1500 GB	3000 GB	800 GB	500 GB	5h

Total source data volume is approximately 17.4 TB. Data retention time for database and log backups is 14 days. S4/HANA and CRM production systems have a RTO SLA of 1 hour. SAP BW is less critical, but needs to be recovered in 3 hours. All test and development systems have a RTO of 5 hours. As production is highly critical, we need to select a data protection approach which not only solves today's challenges but will continue to meet the SLAs in a year or two as databases grow over time. Leveraging compressed backups and restores will meet this requirements best. For protecting test and development systems we recommend enabling deduplication and compression, which will not only meet the SLA, but also help in reducing the storage footprint.

Storage requirements

So how much Blob Cool storage do we need for implementing this backup strategy? For the production systems we run full backups only and assume a 30% data reduction via compression. The following table summarizes the storage amount needed for the database backups.

Production Systems	Only Full
Full	56.8 TB
Incremental	-
Differential	-
Total TB Azure Blob storage consumed	56.8 TB

For the test and development systems, we are assuming a daily change rate of 5%. With deduplication and compression enabled we can assume a data reduction rate of around 70%. As with production we can run fulls only. Another option would be to run just one full per week, combined with daily incrementals and differential which capture the changes. This leads to slightly longer restore times, but helps to limit blob storage use. To put things in numbers; running fulls only requires 48.7 TB vs. 19.7 TB with the levelled backup approach (based on 5% daily change rate). This represents a factor of 2.5.

Test and Development Systems	Only Full	Levelled backups
Full	48.7 TB	10.4 TB
Incremental	-	5.7 TB
Differential	-	3.6 TB
Total TB Azure Blob storage consumed	48.7 TB	19.7 TB

Incremental and differential backups need to be managed by a non-deduplicated storage policy, as those data is considered unique. The same applies to SAP HANA log files whereas deduplication needs to be enabled for protecting SAP application instance data via the file system or VSA agent. In terms of storage capacity, we recommend adding another 1TB of Blob storage for protecting SAP HANA logs and SAP application instance files.

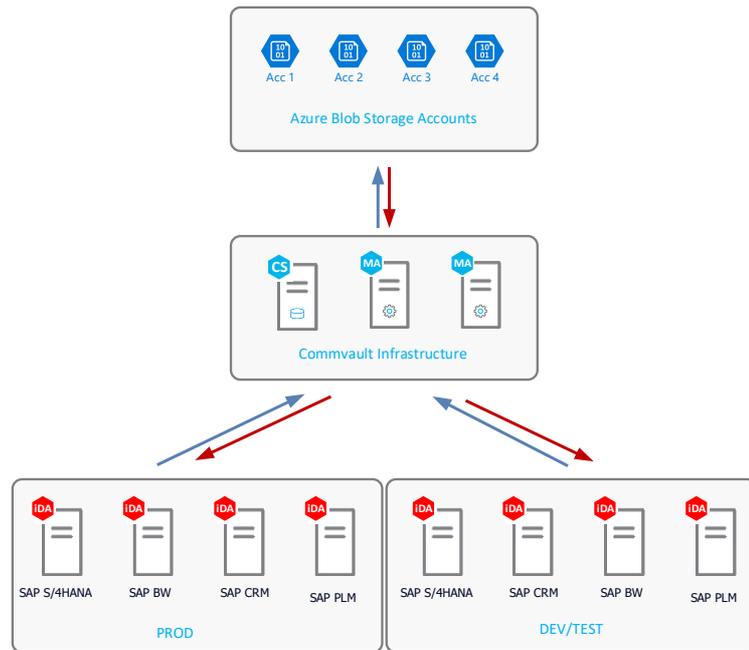
Finally, all storage polices will use the same Commvault cloud library. This uses four Cool Azure Blob storage accounts as mountpoints where backup data is randomly written to each storage account to load balance for read/write performance, and to overcome any storage account throughput limitations (<https://docs.microsoft.com/en-us/azure/storage/common/storage-scalability-targets>).

Commvault infrastructure

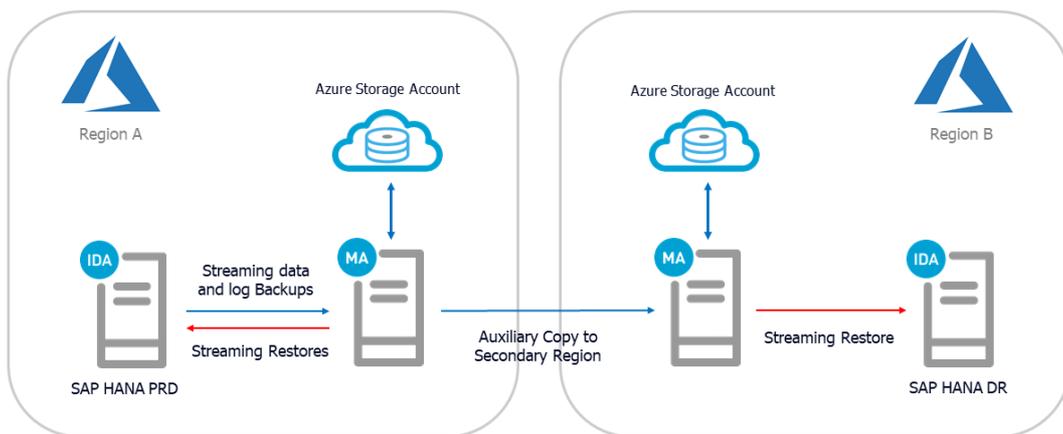
At a minimum, you need a Commvault Commserve and one MediaAgent to control the Cloud Libraries formed by the Blob Storage Accounts. We suggest setting up two MediaAgents and configuring them to handle redundancy/failover and load balancing. This can be achieved via Commvault's GridStor technology. One MediaAgent would be primarily handling SAP HANA logs and the second one all database and file system backups. Commvault's power management can be leveraged to power down one of the MediaAgents, if it's not needed (e.g. outside of the backup window), to help save costs. Restores are still possible using the other MediaAgent which remains in operation for handling the SAP HANA logs. We recommend allocating the following VM types for building the Commvault infrastructure

- Commvserve: 8 CPUs and 32 GB of RAM (e.g. Standard_D8_v3), Windows Server 2016, 150 GB volume for Commvault software and Commserve database
- MediaAgents: 16 CPUs and 64 GB of RAM (e.g. Standard_D16_v3), Windows Server 2016 (or Linux), 2x 200 GB of Premium Storage

Please note that you need another VM for running a so-called VSA proxy, in case you want to protect SAP application instances via the VSA method. As for storage accounts, we recommend to use Azure Blob Cool. The following picture summarizes the discussed reference architecture.



In case a disaster recovery solution is required, the Commvault setup can span multiple Azure regions and copies of all backup data can be automatically made available in all regions, allowing for quick recovery from any possible region failure.



Summary and Conclusion

Starting from an Azure test environment, which was built around the powerful M128s SAP HANA VM and Azure Blob Cool storage, this paper has discussed various approaches on how a SAP environment based on SAP HANA in Microsoft Azure can be protected and SLAs can be met. No matter which SAP workload type you are using, the right solution can be built. We have validated multiple ways to bring back a mission-critical SAP HANA production database of 1.5 TB in less than an hour and how test and development systems can be protected efficiently.

We've also put together a reference architecture for a typical customer environment (including a Disaster Recovery option), discussing various backup strategy options and their impact on Azure Blob storage consumption. Azure Premium Storage showed excellent performance both for running Commvault's internal Deduplication database but also as a backup target.

Many SAP customers are currently looking at migrating their SAP environments running on SAP HANA into Microsoft Azure. These customers are experienced with professional, enterprise-level solutions like Commvault to protect their on-premises SAP environment, and expect the same level of protection when they move to Azure. Commvault and Microsoft have teamed up and demonstrated how Azure and Commvault together can be leveraged to provide a comprehensive solution to make this level of protection available today.