

SAP HANA Backup and Recovery with FlashBlade//S

Rapid backup and recovery of SAP HANA
instances with Commvault Storage
Accelerator and FlashBlade//S

Contents

- Executive summary** 3
- Solution overview** 3
 - Commvault Storage Accelerator for additional horizontal scaling 4
 - Solution benefits 4
- Technology overview** 5
 - FlashBlade 5
 - Multiple channels for SAP HANA backup and restore 6
 - Commvault object storage usage 7
 - SAP Backint 7
 - Data compression and encryption 7
- Performance validation** 9
 - Test environment 9
 - Throughput determination for backup and restore 9
 - Single instance (scale-up) and distributed instance (scale-out) 9
 - Test results 10
- Best practice settings** 13
 - Commvault settings 13
 - SAP HANA settings 14
- Conclusion** 15

Executive summary

As enterprises increasingly rely on SAP HANA to power mission-critical, real-time analytics and transactional workloads, ensuring robust backup and recovery capabilities has become a strategic necessity. One of the primary challenges with SAP HANA backup and recovery is the need to meet aggressive recovery time objectives (RTOs) and recovery point objectives (RPOs) while handling rapidly growing data sets. Meeting SLAs is becoming increasingly difficult due to ever-growing database sizes and shrinking backup windows. Another critical concern is cyber resilience as ransomware and insider threats increasingly target backup data as a means of preventing recovery.

Everpure™ FlashBlade®, with a scale-out architecture purpose-built for high-throughput workloads, provides the backup and recovery throughput necessary for SAP HANA to meet aggressive RTOs. With FlashBlade, a full restore of eight HANA instances in parallel (2.4TB total data) can be completed in 10 minutes.

Additionally, features such as object immutability (via S3 Object Lock) and SafeMode™ retention policies deliver robust protection against ransomware by preventing unauthorized deletion or modification of backup data.

This paper explores a solution using Everpure FlashBlade//S™ as the backup target for SAP HANA instances with Commvault® Complete™ Backup & Recovery. It provides performance metrics, parameter settings, and best practices. Actual performance will vary based on database characteristics, network configuration, and infrastructure. Whether performing backup and restore for scale-up or scale-out instances, FlashBlade delivers the performance required to meet enterprise backup and recovery SLAs.

Solution overview

This paper validates an approach using Everpure FlashBlade Object Store (S3) as the backup target for SAP HANA with the Commvault Storage Accelerator feature. The HANA data and log volumes reside on SAP-certified Everpure FlashArray//XL™ and FlashArray//X™ R5 model arrays. Figure 1 provides the high-level architecture.




FIGURE 1 SAP HANA/Commvault high-level architecture

FlashBlade, with its distributed, blade-based storage design, is ideally suited to take advantage of the multichannel capabilities of SAP BackInt and Commvault data streams. The data streams generate multiple TCP connections that are established to the network interfaces on multiple blades, resulting in high throughput for backup and recovery operations.

Commvault Storage Accelerator for additional horizontal scaling

Storage Accelerator can be used on clients to directly send the backup data to object storage, bypassing the MediaAgent. Without the Storage Accelerator, backup jobs send the client data to the MediaAgent first, and then the MediaAgent stores the data in the storage device.

The same data path is used for restore operations, with data sent directly from object storage to the client and written to primary storage.

Storage Accelerator reduces the load on the MediaAgent by eliminating the need for processing backup data. Pruning, data verification, and defrag operations continue to be processed by the destination MediaAgent. The MediaAgent controlling jobs (referred to as a data server in this configuration) can manage more streams with fewer resources.

Figure 2 illustrates the architectural differences with and without Storage Accelerator.




FIGURE 2 Traditional and Storage Accelerator architectures

Storage Accelerator improves overall backup performance across an environment, but it also increases the CPU load on an individual client. Using Storage Accelerator may even degrade backup and restore speed for clients with CPUs that are heavily loaded during backups. In these cases, you should consider using the traditional architecture and deploying Storage Accelerator on other less-loaded clients. Using the same plans and policies for clients with and without Storage Accelerator is fully supported and does not require any extra configuration. Simply deploy Storage Accelerator on the desired clients and Commvault will automatically use it during backup and recovery.

Commvault CommServe® automatically deploys the Storage Accelerator package (beginning with release 11.26). Clients backing up to FlashBlade will receive the package on their first backup job. Automatic installation of the Storage Accelerator package can be disabled.

Solution benefits

FlashBlade enables SAP HANA backup and recovery environments to:

- **Meet backup windows:** The FlashBlade scale-out storage system supports massive parallel read/write operations. Multiple Backint channels and Commvault streams keep the network pipeline full, providing very high throughput for backup ingestion.
- **Meet aggressive RTOs:** FlashBlade is positioned specifically for rapid restore. The object tier is designed as a fast object S3 platform for backup/restore, not a slow archival appliance. For HANA, this throughput and low-latency S3 access translates into faster database recovery and log replay windows.
- **Scale object storage for large HANA landscapes:** FlashBlade provides a scale-out architecture where storage capacity and performance grow linearly by adding blades. Blades and DirectFlash® Modules are hot-swappable, so capacity and performance can be expanded in place without downtime. This aligns with Evergreen® upgrade models and simplifies lifecycle management.
- **Support concurrent operations:** Maintain consistent performance with the backup and/or restore of multiple HANA instances simultaneously.
- **Provide immutable backups and ransomware protection:** The Object Lock and SafeMode Retention Lock features provide immediate protection against accidental or malicious deletion of object data. Replication of S3 buckets to a secondary site or cloud extends protection for disaster recovery.

Technology overview

FlashBlade

Everpure FlashBlade is a scale-out file and object storage platform designed for high-throughput workloads. Unlike traditional network-attached storage systems that scale capacity and performance together, FlashBlade allows independent scaling of both dimensions.

The FlashBlade employs a massively parallel, distributed-everything design. Metadata, data, and control functions are distributed across all blades, allowing you to scale performance and capacity with each blade you add. This avoids traditional controller bottlenecks and supports very high throughput and low latency for AI, analytics, backup/recovery, and similar workloads.

FlashBlade//S represents the evolution in enterprise scale-out storage, offering a blend of high density, capacity, performance, and scalability to meet the demands of modern applications.

FlashBlade//E™ is a cost-efficient storage platform, aiming to provide effective performance, uncompromising reliability, and all-flash capabilities to a broader audience, especially those with budget considerations.

This paper focuses on the following FlashBlade models for large-scale SAP HANA data protection:

- FlashBlade//S200: high-performance; designed for demanding workloads and larger data sets
- FlashBlade//S100: single-chassis, entry-level scale-out; standard performance and moderate data sets

Figure 3 shows the FlashBlade product family and key features of each model.




FIGURE 3 Technical specifications of the FlashBlade product family

FlashBlade Object Lock

Object Lock on FlashBlade//S seamlessly adds object-level immutability, managed by Commvault. The SafeMode Retention Lock and Freeze Objects features add extra protection against direct attacks on the object buckets and their contents. With these features enabled, locked objects cannot be destroyed or overwritten, and the bucket itself cannot be deleted, even if credentials are compromised.

Object SafeMode (system-wide) for data protection

FlashBlade offers a set of features known as SafeMode. For FlashBlade Object Store, when SafeMode is enabled, manual eradication of buckets can be disabled and object replacement/deletion operations can be blocked. This makes those actions possible only with Everpure Support involvement. Object SafeMode can be configured array-wide (all buckets are protected from manual eradication) or per bucket (action required to protect individual buckets).

Multiple channels for SAP HANA backup and restore

There are three parameters that can have a major impact on backup and restore performance:

- `parallel_data_backup_backint_channels`
- Number of Data Streams
- Data Transfer Pipeline Buffers

SAP HANA splits a data backup into multiple streams. Each stream is sent through a Backint channel to Commvault. Each channel typically produces one backup object/save set. The SAP HANA Backint parameter `parallel_data_backup_backint_channels` controls how many parallel streams (channels) are used when performing a data backup or restore via Backint. The parameter range is 1–32.

The Commvault Number of Data Streams parameter defines how many parallel backup data streams a client or backup job can use simultaneously to send data to storage.

The Commvault Data Transfer Pipeline Buffers parameter (internal setting `nNumPipelineBuffers`) is used in the data movement pipeline during backup, restore, or copy operations. Read buffers ingest data from HANA. Chunking buffers then segment that data into deduplicated blocks (if enabled). Finally, send buffers queue the data for transmission to storage. The HANA client uses these buffers to write to the FlashBlade during backup operations and read from the FlashBlade during restore operations.

Due to the parallelism accomplished with multiple Backint channels and Commvault streams, multiple TCP connections are established between the HANA client and the FlashBlade during backup and restore operations. These connections are established with the network interfaces on multiple blades, providing for the rapid ingest of data during backups and reading of data during restores.

Figure 4 shows the parallelism with multiple backup/recovery channels using multiple TCP connections to Everpure FlashBlade.




FIGURE 4 Backup channels and TCP connections to FlashBlade for backup and restore operations

The Backint channels control HANA-side parallelism, reading data from or writing data to primary storage. The Commvault streams map closely to those channels and manage data movement. The pipeline buffers (client-side only) handle flow control and throughput. Data is written to or read from the FlashBlade via multiple TCP connections.

Commvault object storage usage

Commvault uses a very different process to read and write to object storage than it uses with file storage. With file storage, each backup or restore stream is broken into chunks. These chunks are written sequentially into large data files, with a 1:1 relationship between active files and streams.

With object storage, the data files are broken into smaller binary large objects (BLOBs) before they are written. The system creates a thread pool that is shared across all streams. As each thread is activated, it opens a TCP connection to storage, and the threads write BLOBs in a highly parallel manner. Commvault automatically expands the thread pool as needed (up to a tunable maximum) to improve throughput and can reuse the same thread for multiple objects. Because the threads each have their own TCP connections, Commvault's model results in excellent load distribution across blades. In most cases, Commvault does not benefit from increasing the maximum thread count, but in certain resource-limited environments, it can increase backup throughput.

SAP Backint

Backint for SAP HANA is an API that enables third-party tool vendors to directly connect their backup agents to the SAP HANA database. Backups are transferred via pipe from the SAP HANA database to the third-party backup agent running on the SAP HANA database server. The agent then sends the backups to the third-party backup server. SAP Backint acts as a bridge between SAP databases and enterprise backup systems, enabling high-performance, scalable, and centrally managed backups.

Key features and benefits

- Backup jobs can be controlled and monitored through enterprise backup systems, enabling centralized policies, scheduling, and reporting. Enterprise backup tools allow policy-based retention and compliance controls for SAP data.
- Backups can be written directly to storage systems to avoid unnecessary intermediate storage on SAP servers. Since backups go directly to external storage, local disk space usage is minimized.
- Backint enables parallel data streams, improving backup performance for large SAP databases.
- Backint supports encrypted data transfer and secure communication between SAP and backup systems.
- Backint provides continuous transaction log backups, which are essential for point-in-time recovery in SAP environments.
- Organizations achieve greater cost efficiency by reusing their existing backup infrastructure instead of deploying separate SAP-specific solutions.

Data compression and encryption

There are three infrastructure components where encryption and data reduction can be performed: on the FlashBlade, Commvault MediaAgents, or HANA client servers. The best location for encryption and data reduction (compression and deduplication) is determined by the location of infrastructure resource constraints.

Compression

The always-on, inline compression of FlashBlade is applied to file and object data as it is ingested. The same compression engine is applied uniformly to every FlashBlade filesystem and object bucket, providing consistent data reduction behavior across all tenants and protocols on the system. FlashBlade attempts to compress whatever is written to it, including already-encrypted data. If client data is not compressible, it is stored raw to avoid unnecessary CPU overhead. However, encrypted data is effectively noncompressible, so data reduction in practice is approximately 1:1 on those data sets and most of the compression benefit is lost.

With client-side compression, data is compressed before leaving the HANA server. Client-side compression reduces network bandwidth usage because less data is sent to the MediaAgents and storage (especially beneficial for large HANA backups). However, it does consume CPU on the HANA server, so it is not ideal if the HANA nodes are already CPU-bound. It can also interfere with backup throughput if the CPU becomes a bottleneck. This option is best when the network is the bottleneck and the HANA server has spare CPU resources.

With MediaAgent-side compression, the MediaAgent performs compression before writing data to FlashBlade. MediaAgent-side compression offloads CPU from the HANA servers and centralizes compression on fewer but more scalable servers. However, MediaAgent-side compression requires higher network usage than client-side compression. MediaAgents can become a bottleneck when serving a large number of client systems. Therefore, the MediaAgents must be sized properly. MediaAgent-side compression is suitable for environments where the HANA servers are CPU-sensitive, a high-bandwidth network is available, and there is a strong MediaAgent infrastructure.

When using Commvault deduplication, compression can still be enabled, but the exact behavior depends on the deduplicated storage policy.

Encryption

FlashBlade//S encrypts all data at rest on DirectFlash Modules with AES-256 (XTS/CTR) hardware-based encryption. This is included as part of the base system and does not require an extra license. Encryption keys are managed automatically and refreshed regularly. Because encryption occurs after compression, you get strong security without sacrificing compression. FlashBlade data-at-rest encryption protects against disk theft and physical compromise. Encryption is offloaded from the MediaAgents and HANA clients and integrated into the data path. There is no impact on backup performance, and the process is fully transparent to Commvault. However, this approach only protects data at rest. Data remains unencrypted over the network and while inside MediaAgent memory/processing.

Encryption at the HANA client database layer is the native SAP HANA encryption (backup encryption or persistence encryption) that occurs before Commvault sees the data. It provides end-to-end security from the source (data is encrypted before leaving HANA), is independent of backup software, and covers all HANA data paths (not just backups). While it provides the strongest security boundary, it is also the worst option for performance and storage efficiency. If it is enabled, neither Commvault nor the FlashBlade can compress the data. In addition, it is generally not recommended with Commvault deduplication architectures. Per Commvault's SAP HANA best practices, encryption should be enabled client-side or on the MediaAgent.

Encryption at the Commvault client/MediaAgent (software encryption) comes in two encryption modes:

- Agent-side (network and media)—encrypt before sending and store encrypted
- MediaAgent-side (media only)—send unencrypted and encrypt before storage

The agent-side encryption (HANA host/client) encrypts data in transit and at rest, maintains Commvault control over the encryption keys, and works with deduplication (unlike HANA encryption). However, there is CPU overhead on the HANA host with this mode and, depending on the server workload, there may be a slight impact on backup performance.

MediaAgent-side encryption offloads CPU cycles from the HANA server to the MediaAgent. It provides a centralized encryption point for multiple clients. However, it has a weaker security posture because data travels unencrypted over the network (unless you add TLS separately).

Depending on your environment, the options with the best balance are likely agent-side or MediaAgent-side encryption, with FlashBlade encryption to protect data at rest and data in transit (with TLS if using MediaAgent-side encryption). The decision is whether the encryption overhead is on the HANA client or the MediaAgent. If compute resources on the HANA clients and MediaAgents are constrained but the network has plenty of available bandwidth, then utilizing FlashBlade compression and data-at-rest encryption may be the best solution, assuming it meets compliance requirements.

Regardless of where compression, deduplication, and encryption are used, S3 Object Lock and SafeMode should be enabled on the HANA backup buckets so objects are immutable even if credentials are compromised.

Performance validation

This section presents measured performance from testing SAP HANA backup and restore operations on FlashBlade.

Test environment

Table 1 details the HANA environment configuration.

Component	Details
Compute	Servers: 8 VMware guests vCPUs: 32 Memory: 512GB RAM
Operating system	SUSE Linux Enterprise Server 15 SP6 Kernel: 6.4.0-150600.23.78-default
Network	2x 100GbE NICs per ESXi host (8 ESXi hosts) Bonding: IEEE 802.3ad LACP, layer 3+4 hash Aggregate bandwidth: 200Gbps
Primary storage	1x Everpure FlashArray//XL170 R1 1x Everpure FlashArray//X90 R5 (4 HANA instances per array)
Backup storage	FlashBlade//S100, FlashBlade//S200 Protocol: Amazon S3; one S3 bucket configured per array FlashBlade//S100 configuration: 10 blades, 4 DirectFlash Modules per blade FlashBlade//S200 configuration: 7 blades, 4 DirectFlash Modules per blade
SAP HANA database	SAP HANA 2.0 SPS 08 Database size: 384TB on each HANA node
Commvault	Commvault Complete Backup & Recovery, version 11.40.34

TABLE 1 HANA test environment components

Throughput determination for backup and restore

The throughput data reported for the test results is the average throughput reported in the Commvault job details. SAP Backint also records throughput results, but the two measurements are not the same.

The throughput reported by Backint measures the data handed over by SAP to Backint. This typically reflects the logical database size with SAP-side compression.

The throughput reported by the Commvault MediaAgent measures the data actually written/transferred to the storage device. The transmitted data includes the effects of deduplication, compression, network transfer behavior, and disk/tape write performance. The throughput number is end-to-end throughput, not just network speed or disk write speed. It reflects the entire pipeline efficiency.

Single instance (scale-up) and distributed instance (scale-out)

Our testing environment consisted of eight scale-up HANA instances. This allowed the flexibility to efficiently test backup and restore on an increasing number of instances with two different FlashBlade models. The process to back up and restore a scale-out instance is no different from a scale-up instance. A backup or restore is initiated for the systemDB or a tenant, and Backint manages the process among all of the worker nodes.

Backup and restore performance on scale-out instances should not vary greatly from the performance on the same number of scale-up instances executed in parallel. The parallel backup or recovery operations were launched simultaneously using schedules or by executing `hdbsql` commands to start the procedure on each HANA node at the same time. Backup and restore performance could potentially be faster for a scale-out instance due to the parallel operation of the procedure on a single database on the worker nodes.

There are additional considerations and complexities with the restore of a scale-out HANA instance. If the restore is to a new environment, there must be the same number of worker nodes as the original configuration. In addition, if restoring a database, the entire database must be restored on all worker nodes. It is not possible to restore the database files on a single worker node.

However, using the CommServe GUI to launch the backup or restore of a scale-out or scale-up instance is the same. Backint manages the procedure among the HANA nodes via the Backint integration. Commvault provides detailed instructions to configure a scale-out instance for backup/recovery.

Test results

This section provides a backup and restore throughput comparison obtained from validations performed on HANA instances using S3 object storage as backup targets on FlashBlade//S100 and FlashBlade//S200. The following test scenarios were executed:

- Single-instance testing to verify parameter settings for best throughput
- Parallel backup and recovery on multiple HANA instances

SAP HANA and Commvault parameter optimization

Testing explored the impact of the following three parameters on backup and recovery performance:

- SAP HANA: `parallel_data_backup_backint_channels` (defined in `global.ini`)
- Commvault: Number of Data Streams (set in the HANA database default subclient)
- Commvault: Data Transfer Pipeline Buffers (set on the HANA clients and the MediaAgents)

Commvault sets the default values for `parallel_data_backup_backint_channels` and Number of Data Streams to 4. Tests were run with these parameters set to 4, 8, 16, and 32 to find the setting that allowed the most throughput. The testing was performed with the Storage Accelerator role enabled to send all of the backup data directly to the FlashBlade. The highest backup throughput was achieved with both parameters set to 32.

The Data Transfer Pipeline Buffers parameter (internal setting `nNumPipelineBuffers`) controls how many buffers are used in the data movement pipeline during backup, restore, or copy operations. The pipeline buffers are used for data transfer between components (client, MediaAgent, and storage). These buffers are allocated in shared memory on each server participating in the data movement process. The pipeline allows parallel reading, transferring, and writing of data, improving throughput. The size of each buffer is 64KB. Increasing the number of buffers may increase performance, but it also consumes more RAM on the server because each buffer is allocated in shared memory.

Commvault recommends that the number of pipeline buffers be set to 600 if the storage is a high-throughput device and the clients and MediaAgents are configured with NICs at 10GbE or higher. The Data Transfer Pipeline Buffers parameter was tested at 600 and 900 on the clients and MediaAgents. There was no performance increase with 900 buffers, so all testing was completed with 600 pipeline buffers.

Single-node backup and recovery performance when scaling the number of BackInt channels and client data streams from 4 to 32 is shown in Figure 5.

Backup and Restore Performance / Channel Scaling




FIGURE 5 Backup and restore channel scaling

Increasing the number of data channels from 4 to 8 increased the backup throughput very slightly. Backup throughput for 8 and 16 channels was about the same (8.4TB/hour), but throughput increased significantly when the number of channels was increased to 32 (11.3TB/hour). With data sent over 32 channels, network bandwidth between the HANA clients and the storage devices was more efficiently utilized.

Restores of the SAP HANA database used the same number of channels used for backup. Restore throughput remained consistent across channel counts, indicating the bottleneck was elsewhere in the environment (CPU, network, or source storage) rather than with the HANA and Commvault channel parallelism. However, multi-instance backup and restore demonstrates that aggregated throughput increases, highlighting the ability of the FlashBlade scale-out architecture to accommodate an increasing number of parallel data streams.

HANA database backup

Testing measured the aggregate throughput when running concurrent backups across multiple scale-up HANA instances. The parameter tunings determined during the single-node testing were applied for all HANA instances used in these tests.

With the Storage Accelerator role enabled on the HANA clients, all backup data is written directly to the FlashBlade, bypassing the MediaAgents. Only metadata is written to the MediaAgents. Data reduction is performed on the FlashBlade. There is no deduplication processing overhead on the MediaAgents or the HANA clients.

Tests were run with concurrent backups on two, four, and eight HANA nodes, with the results shown in Figure 6. With eight concurrent instances, FlashBlade//S200 delivered 45.4TB/hour (12.6GB/s) of aggregated backup throughput. The longest parallel backup duration was 5 minutes and 8 seconds.

FlashBlade//S100 delivered 23.7TB/hour (6.6TB/s) of aggregated throughput. The longest parallel backup duration was 8 minutes and 33 seconds.

The performance scaling from two to eight instances was greater on the FlashBlade//S200 due to its higher bandwidth and performance capability.

Backup Throughput




FIGURE 6 Throughput for concurrent backups on two, four, and eight HANA instances

HANA database restore

By default, restore operations use the same number of channels (streams) as backup operations due to how data is striped, tracked, and reassembled.

During backup, SAP HANA divides the database into multiple data streams (channels). Each channel writes a portion of the data in parallel to Commvault. Commvault stores these as separate backup streams, often distributed across media, chunks, or storage blocks. The result is a striped, parallelized backup set.

During restore, HANA expects to read those streams back in parallel in the same logical way they were written. Each channel corresponds to a portion of the backup. SAP HANA recovery is highly parallel and time-sensitive. Commvault mapping between backup streams and restore streams is 1:1 optimized. For our testing, the same number of streams (32) was used for backup and recovery.

Figure 7 shows aggregate throughput for concurrent restores to two, four, and eight SAP HANA instances, with each instance residing on separate virtual machine guests.

Restore Throughput




FIGURE 7 Throughput for concurrent restores on two, four, and eight instances

Executing restores on eight SAP HANA instances concurrently produced aggregated throughput of 19.9TB/hour (5.53GB/s) on FlashBlade//S100 and 20.4TB/hour (5.64GB/s) on FlashBlade//S200.

Full restore of eight HANA instances (2.4TB) in parallel completed in approximately 9 minutes on both the FlashBlade//S100 and FlashBlade//S200.

Best practice settings

Commvault settings

The best write performance on FlashBlade storage is achieved by spreading the load across as many blades as possible. Configuring Commvault for maximum writers and distributed streams yields the best outcome with a simple setup.

Number of Data Streams

The Number of Data Streams parameter defines how many parallel backup data streams a client or backup job can use simultaneously to send data to storage. It is a key parameter for backup performance and throughput. During a backup job, the HANA client splits the data into multiple streams. Each stream runs as an independent process, and the streams transfer data simultaneously to the MediaAgent. The value of the SAP HANA Backint parameter `parallel_data_backup_backint_channels` should normally match or be lower than the value specified for the Number of Data Streams parameter. If the number of Commvault streams is lower, SAP HANA may request more streams than the backup software can handle, which can limit performance.

Relationship: Commvault Number of Data Streams $\geq$ SAP HANA `parallel_data_backup_backint_channels`

Data Transfer Pipeline Buffers

The Data Transfer Pipeline Buffers parameter (internal setting `nNumPipelineBuffers`) sets the number of pipeline buffers used for data transfer between components (client, MediaAgent, and storage) during backup, restore, or copy operations. These buffers are allocated in shared memory on each server participating in the data movement process. The size of each buffer is 64K. The default number of buffers is 90, the minimum number is 30, and the maximum number is 1020.

Increasing this value can improve performance because more buffers allow more data to be queued in the pipeline. Data movement stages (read, network transfer, and write) can overlap more effectively. However, increasing the number of buffers also consumes more RAM on the server because each buffer is allocated in shared memory.

The following are typical scenarios where administrators may increase this value:

- High-speed networks (10GbE or higher)
- Tape backups requiring large streaming buffers
- High-throughput disk libraries or dedupe storage
- Environments where the MediaAgent becomes the bottleneck

Because the HANA clients and the MediaAgents are configured with 10GbE NICs and the Everpure FlashBlade is a scale-out, high-throughput device, the Data Transfer Pipeline Buffers parameter was set to 600 on those servers. Tests were performed with the number of pipeline buffers set to 900, but there was no increase in throughput.

Deduplication block size reduction

Deduplication block size affects performance and storage consumption. Larger block sizes yield better performance but consume more storage. Commvault's recommendation for FlashBlade object storage is to set the block size to 128KB to maximize data reduction. Storage policies tied to cloud storage targets default to 512KB in most cases, so you will need to change the setting manually. Deduplication block size is configured at the storage pool, or at the storage policy if it is not tied to a pool, using the CommCell Console interface.

Note: Beginning with release 11.36, Commvault automatically sets the deduplication block size to 128KB on new storage pools with the Everpure FlashBlade storage type.

SAP HANA settings

During the process of adding a HANA instance in Commvault, the following parameters are automatically added to the HANA instance global.ini files for both the SYSTEMDB and the tenant DB:

```
[backup]
parallel_data_backup_backint_channels = 4
log_backup_using_backint = true
log_backup_parameter_file = /usr/sap/<SID>/SYS/global/hdb/opt/hdbconfig/param
data_backup_parameter_file = /usr/sap/<SID>/SYS/global/hdb/opt/hdbconfig/param
backint_response_timeout = 1800
catalog_backup_using_backint = true
catalog_backup_parameter_file = /usr/sap/<SID>/SYS/global/hdb/opt/hdbconfig/param
parallel_data_backup_backint_size_threshold = 32
```

Per Commvault SAP HANA documentation, the following parameters should also be set:

```
Data_backup_buffer_size = 2048
tcp_backlog = 2048
```

Backint parameter

The SAP HANA Backint parameter `parallel_data_backup_backint_channels` controls how many parallel streams (channels) are used when performing a data backup via Backint. It is used to improve backup performance by splitting the backup into multiple streams. The parameter range is 1–32.

Parallel channels are used **only when the backup size exceeds a threshold** defined by `parallel_data_backup_backint_size_threshold`. If the backup is smaller than the threshold, **only one channel is used**. The default value for `parallel_data_backup_backint_size_threshold` is 128GB.

A backup needs a buffer for each channel, plus one additional buffer, with the size `data_backup_buffer_size`. Therefore, increasing channels increases memory consumption on the HANA client.

During recovery, SAP HANA automatically uses the same number of streams used during backup. The parameter value itself is not checked during recovery.

The value of `parallel_data_backup_backint_channels` should normally match or be lower than the value specified for the Commvault **Number of Data Streams** parameter. If the number of Commvault streams is lower, SAP HANA may request more streams than the backup software can handle, which can limit performance.

Additional best practice settings can be found in [Configuring Commvault with FlashBlade: Best Practices](#).

Conclusion

While each FlashBlade model provides exceptionally fast database restores, each comes with its own performance characteristics and capacities. From a data protection standpoint, this ensures that there is a FlashBlade model to meet the needs of any enterprise using SAP HANA in its mission-critical applications.

The validated configurations in this paper demonstrate:

- Single-node throughput of 3.1–3.15GB/s (~11.2–11.3TB/hour) with optimized Backint and Commvault channel settings
- Aggregated throughput with eight parallel instance backups of 12.6GB/s on FlashBlade//S200 and 6.6GB/s on FlashBlade//S100
- Full backup completion for eight parallel instance backups: 5 minutes on FlashBlade//S200 and 8 minutes and 28 seconds on FlashBlade//S100
- Aggregated throughput with the parallel restore of eight HANA instances: 5.66GB/s on FlashBlade//S200 and 5.53GB/s on FlashBlade//S100

The validation work in this paper demonstrates that combining SAP HANA, Commvault Complete Backup & Recovery, and Everpure FlashBlade//S delivers a backup and recovery platform that can keep pace with large, mission-critical HANA landscapes while maintaining aggressive RPO and RTO targets. By using S3 buckets on FlashBlade as the primary backup target and leveraging SAP Backint integration with Commvault, enterprises can continuously protect multi-terabyte HANA databases with highly parallel backup workflows that sustain throughput in the tens of terabytes per hour across multiple instances.

At the same time, FlashBlade scale-out architecture, always-on data reduction, and support for Commvault Storage Accelerator provide a straightforward path to scaling performance and capacity. As requirements grow, additional blades or HANA clients can be added without rearchitecting the environment. Cyber-resilient features such as S3 Object Lock, SafeMode Retention Lock, and asynchronous S3 replication ensure that SAP HANA backups remain immutable, recoverable, and available across sites, even in the face of ransomware or operational errors. Together, these capabilities make FlashBlade and Commvault a robust, scalable, and operationally efficient foundation for SAP HANA data protection that evolves with growing data volumes and increasingly stringent business and regulatory requirements.

[Learn More About FlashBlade](#)

[Visit Our Website](#)

800.379.PURE

