



Hedvig-Commvault-SAP-HANA Integration

Solution Brief

As large volumes of real-time data flood into organizations, they are hard-pressed to put in place a big data strategy to process this deluge of data and to unlock business value in this data. Any organization with an insight-driven business model that is focused on accelerating the time to market for a big data solution should consider leveraging tools with the ability to not only distill and analyze large amounts of data, but simultaneously analyze data and support transaction-oriented applications.

SAP HANA uses in-memory database technology and multicore data processing to deliver high performance for analytical and transactional applications. SAP HANA also allows users to choose between a column-oriented and a row-oriented schema for unstructured and structured data management within the same system. In addition, SAP HANA can be deployed on-premises or in the cloud, with the ability to scale up, as well as scale out.

The rate of data growth determines the strategy for scaling the SAP HANA deployment. While scaling up can be a good starting point to deal with increasing data growth, eventually a tipping point is reached when the memory on the SAP HANA host is no longer feasible to effectively process large volumes of data. A scale-out deployment is implemented with multiple SAP HANA hosts in one SAP HANA cluster using a shared-nothing approach for the underlying data, which is partitioned across these hosts. The data is stored in a shared filesystem accessible to all hosts in the cluster.

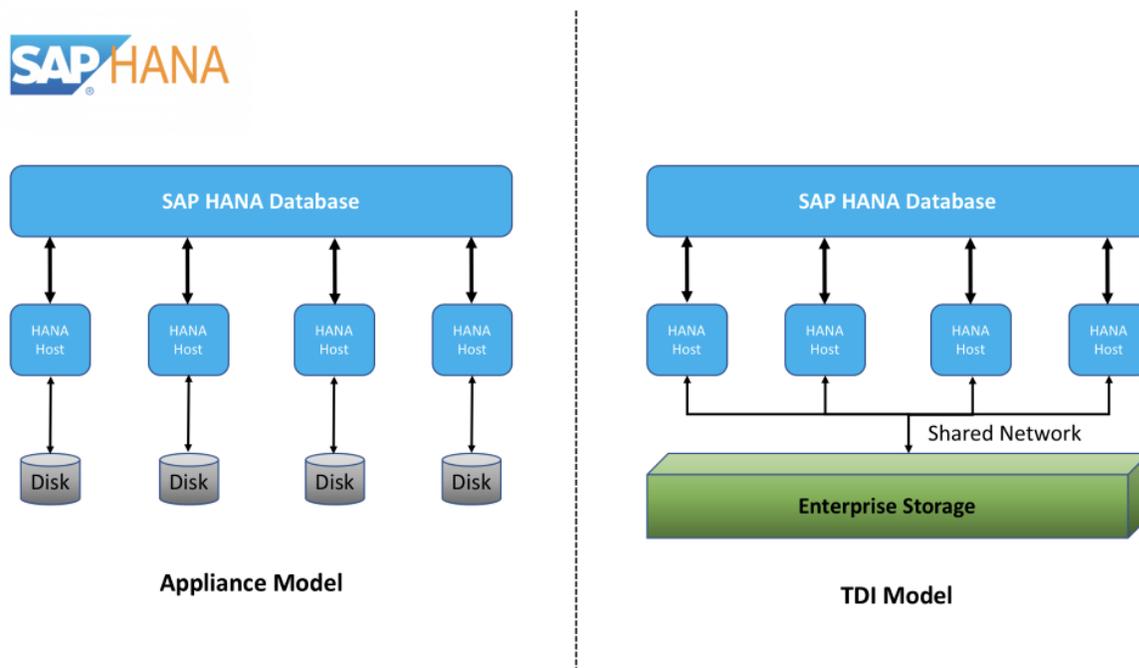
While an SAP HANA deployment can be scaled seamlessly on-demand, it is imperative to complement this with a storage solution that can provide a highly scalable and fault-tolerant shared filesystem capable of handling large volumes of incoming data.

The Hedvig Distributed Storage Platform provides a high-performance and resilient shared filesystem, coupled with the management simplicity required by SAP HANA environments. In addition, Commvault provides a certified SAP HANA data protection solution that is tightly SAP-integrated and is designed for universal protection, thereby providing a comprehensive offering to help maximize the success of SAP deployments, regardless of where they live.

SAP HANA Architecture

SAP HANA is predominantly deployed in two different models – Appliance Model and Tailored Data Center Integration (TDI) Model.

In the *Appliance Model*, organizations purchase SAP HANA appliances that include integrated server, storage, and network components certified by SAP and built by one of the SAP HANA hardware partners. Naturally, these configurations are less flexible in terms of scaling due to the fixed server, storage, and network components.



The *TDI Model* offers more flexibility by enabling organizations to choose their own server, storage, and network components. In this model, storage and network components are shared across multiple SAP HANA hosts. Organizations can leverage their existing enterprise storage infrastructure for SAP HANA workloads, thereby reducing the overall risk and cost associated with the SAP HANA implementation.

With the TDI model, SAP HANA workloads can be seamlessly transitioned to the Hedvig Distributed Storage Platform and take advantage of the advanced enterprise storage capabilities that Hedvig offers. In addition, SAP HANA workloads can achieve native high availability across multiple zones/regions/clouds to provide near-zero recovery point objectives and recovery time objectives.

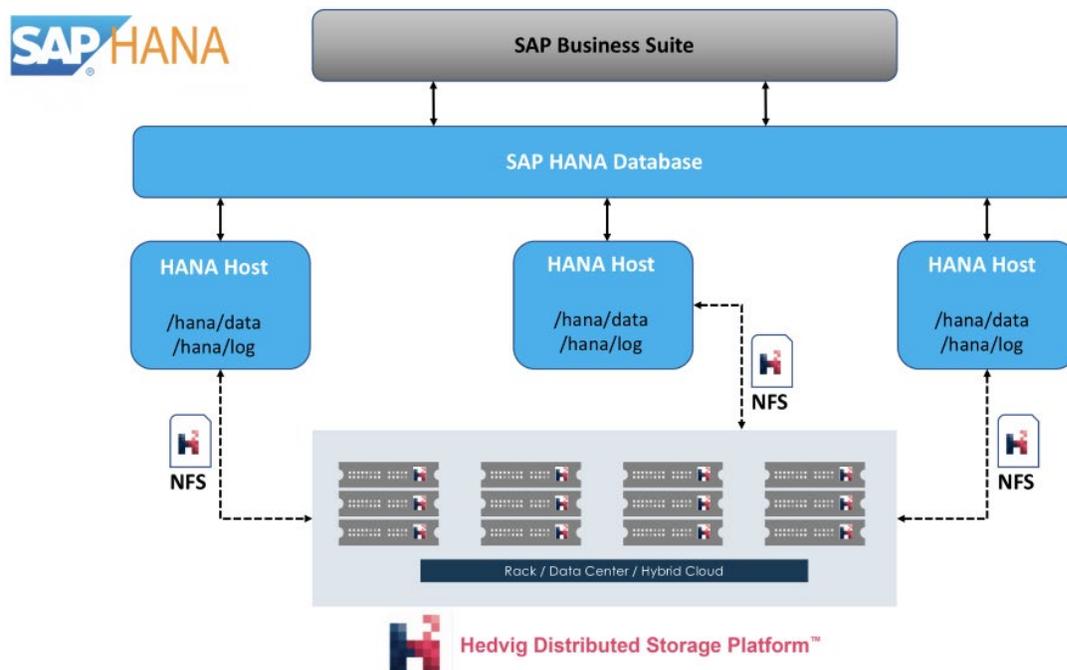
Persistent Storage for SAP HANA

To ensure that the database can always be restored to its most recent committed state, SAP HANA periodically writes changes to data in the database to disk (*data volume*). Transaction logs are also saved regularly to disk (*log volume*).

SAP HANA persists in-memory data by using savepoints. During a savepoint operation, the SAP HANA database flushes all changed data from memory to the data volumes. The data belonging to a savepoint represents a consistent state of the data on disk and remains so until the next savepoint operation has completed. Data volumes are typically accessed using random I/O patterns.

Redo log entries are written to log volumes for all persistent data changes. In the event of a database restart (for example, after a crash), data from the last completed savepoint can be read from the data volumes, and the redo log entries written to the log volumes, since the last savepoint can be replayed. Log volumes are typically accessed using sequential I/O patterns.

It is vital to ensure that there is enough space on the disk to save data and logs. A disk-full event will render the database non-functional.



The Hedvig Distributed Storage Platform provides a modern solution with all the capabilities required to support SAP HANA workloads. Hedvig virtualizes and aggregates flash and spinning disk in a server cluster or cloud, presenting it as a single, elastic storage system.

Hedvig NFS virtual disks, provisioned with the pin-to-flash option, coupled with deduplication, compression, and client-caching, provide a high-performance and low latency tier ideal for data and log volumes. NFS virtual disks backing data and log volumes can be easily expanded on-demand and are resilient to node, rack, and event data center failures, thereby ensuring that mission-critical SAP HANA environments are protected and highly available at all times.

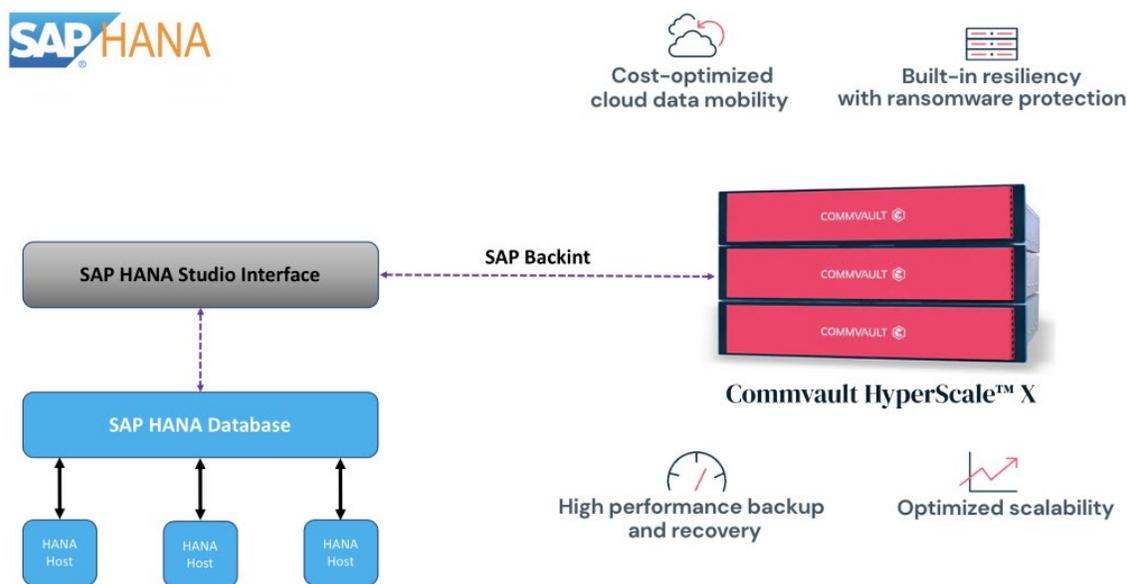
With Hedvig taking responsibility for data replication, it obviates the need to implement SAP HANA System Replication, which creates an exact copy of a primary server by copying data and transaction logs to a secondary server to be used in the event of a failure.

Backup and Recovery for SAP HANA

While Hedvig provides native high availability for SAP HANA workloads, there could still be scenarios where data backups are warranted. SAP strongly recommends implementing backup strategies for SAP HANA environments handling mission-critical data, at the following times:

- After the initial load
- At regular intervals
- Before the database software is upgraded to a new version
- After any situation that causes log writing to be interrupted

[Commvault Complete™ Backup & Recovery](#) helps organizations build, protect, and operate SAP HANA environments that run in private, public, or hybrid cloud environments. Commvault’s tight integration with SAP HANA provides impact-free recovery copies, which are combined with other backup strategies to meet even the most stringent recovery and disaster recovery (DR) SLAs. Advanced copy management also helps SAP customers with other important functions, such as dev/test, migration, long-term retention, and data governance demands.



Commvault uses the SAP HANA Backint interface program to back up data directly to the attached media. In case of a data loss due to hardware failure, natural disaster, accidental deletion, or corruption of data, you can restore the backed-up data and the log files directly from the media using the SAP HANA Backint interface program. This reduces complexity, human intervention, and the risk of data loss.

SAP HANA Studio is the most common tool used to back up and recover HANA databases because it is preinstalled on SAP HANA. It is used by administrators to administer activities, such as to start and stop services, monitor the system, configure system settings, and manage users and authorizations. SAP HANA Studio is a client environment that can be used to access the SAP HANA database located in the same environment or at a remote location.

Benefits of Hedvig + Commvault for SAP HANA

Unified solution – Eliminate silos with a single unified solution to manage all aspects of SAP HANA workloads, from persistent storage for SAP HANA data and log volumes to tight integration with SAP HANA Studio and SAP HANA Cockpit through the SAP HANA Backint API.

Advanced enterprise storage services – Support production storage operations and enterprise SLAs with a rich set of storage capabilities, such as deduplication, compression, encryption, snapshots, clones, replication, auto-tiering, multi-tenancy, and self-healing.

Unmatched scale with performance optimization – Optimize performance for SAP HANA persistent volumes with Hedvig's pin-to-flash and client-side caching capabilities. Drive backup window and application impact of critical SAP HANA deployments to zero with Commvault's advanced snapshot technology.

Cloud-like simplicity – Implement intuitive workflows to streamline and automate storage provisioning with Hedvig's native HTML5 interface, thereby bringing provisioning simplicity of public clouds, such as AWS, to any data center. Refresh SAP landscape systems running on SAP HANA via Commvault's easy-to-use WEBGUI and integration with SAP's Landscape Management tool.

Hedvig Inc. believes the information in this publication is accurate as of its publication date. The information is subject to change without notice. The information in this publication is provided as is. Hedvig Inc. makes no representations or warranties of any kind with respect to the information in this publication and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any Hedvig Inc. software described in this publication requires an applicable software license. All trademarks are the property of their respective owners. Revision date: 091120.

Software-defined AES-256, FIPS compliant encryption of data in flight and at rest.