

Application Architecture SAP Datasphere

Status	Approved
Owner	SHEPHERD-ext, Robert
Stakeholders	PETTIFORD-ext, owen DAHN-ext, Werner WANAMAKER, Craig Guillaume Muller
LeanIX Link	SAP Datasphere

Introduction

SAP Datasphere (DSP) is a cloud data warehousing solution used by Syensqo to extract, load and transform data from the SyWay transaction data processing systems. It is a component of BTP. The data is made available for reporting and analytics in SAC and distribution to other systems such as PaPM, Sustainability Control Tower and MS Fabric.

At time of endorsement (05/12/2025), it is recognised that Datasphere will be a component of SAP's enterprise data solution, Business Data Cloud (BDC). However, BDC is currently out of scope for SyWay. BDC will be reviewed in 2026 and, if adopted, some elements of this document may change.

Purpose

The purpose of this document is as source of information about Datasphere which will facilitate the support of SyWay implementation. It provides additional information to supplement the over-arching application architecture document for SAP BTP.

Scope

This document applies to Datasphere only.

Linked Documents

Information from other documents is referenced throughout this document. This is to avoid duplication and prevent contradictions. The key documents referenced are:

[Application Architecture SAP BTP](#)

[SAP Analytics Approach](#)

[SAP Analytics and Reporting Standards](#)

[Security Approach](#)

[Security Approach for Analytics](#)

[Application Architecture SAP Analytics Cloud](#)

[Network and Infrastructure Architecture DD-TEC-070](#)

[DD-TEC-170 Transport Management for Release 4](#)

[CD-SOL-020 Reporting Approach](#)

Key Decisions and Requirements

The table below provides details of Datasphere specific architectural decisions. (Decisions made e.g. for the BTP platform are not restated.)

Decision	Rationale
The project will utilise seamless planning whereby SAC planning stores its data in DSP	This is SAP's strategic direction. Seamless planning offers far better functionality when integrating planning data for reporting and reference data for planning.
DSP and SAC will be deployed in the same data centre.	SAP do not support seamless planning across different data centres.
DSP will only connect to a single SAC tenant	SAP limitation
All SAP data fed to MS Fabric will be via DSP	There are high licensing costs associated with extracting data through alternative approaches. Using a single extraction mechanism will improve data consistency across platforms and reduce the performance impact on S/4 of extracting data.

Datasphere will be used to consolidate SAP S/4 data from the regional systems	This will provide a unified dataset to support cross-region reporting. This is expected to be of high importance as some GBUs span regions.
SAP Business content (delivered models) will be used as a reference for model design	This will lead to faster implementation.
No CUI data will be loaded into DSP	Data must be secured using NextLabs, and NextLabs does not work with Datasphere. N.B. As of 12/11/2025, it is understood that Datasphere is on the NextLabs roadmap with an announcement regarding availability due in 2026
PaPM Will read data from DSP and write calculations back to DSP	This is the standard SAP approach for PaPM with the 'bring your own database' operating model
Report cataloguing using Colibra is not in scope	The cost/benefit evaluation does not justify this application

Application Architecture

Application Architecture Overview

Datasphere Details

Customer Number	3008440
Cloud Provider	MS Azure
Cloud Region	Netherlands
Service model	Software as a Service
Licence	SAP Cloud Platform Enterprise Agreement (CPEA)
Deployment model	We are using the Public model
Database	HANA Cloud


Application Architecture Components

Datasphere Internal Application Components

The SAP help documentation provides a [Datasphere Overview](#) which explains the system architecture.

The following diagram and table represent the application components and how SyWay is utilizing them.

SAP Datasphere



Provide your users an end-to-end view of their data landscape that is trusted, secure, and actionable

Datasphere's Architectural Components	Description	SyWay Usage
SAP Analytics Cloud Native Integration	How Datasphere and SAC interact	Heavily used
API External Consumption	Mainly used to allow 3rd party tools, e.g. Power BI to access Datasphere	Not used. SAC is the front end tool of choice.
Catalog	An internal cataloguing capability providing a glossary of terms and documentation	Not planned for usage unless / until SAP can surface this information to SAC
Administration and Security Services	The mechanism by which Datasphere ensures the right data is seen by the right people	Heavily used
Space Management	The critical component for the structuring of Datasphere and the securing of data within it.	Heavily used
Data Modelling Services	Developer tool for creation of data models.	Heavily used, is basis for all build
Business Modelling Services	Vestigial application to support business owned modelling left over from Data Warehouse Cloud	Not used
Data Lake	A dedicated, on-read schema-flexible storage area in SAP HANA Cloud for raw and archived data repository Optimized for ingesting and storing large volumes of raw data and acts as the "landing" zone for unstructured data before any modelling or transformation takes place.	Not used. The SyWay default for unstructured data is MS Fabric
SAP HANA Cloud Disk / In Memory	Ingested data is stored in SAP HANA database tables. The SAP default of 'on disk' will be taken, though in-memory storage is available and will be utilised where on-disk performance does not meet expectations.	Heavily used
Replication and Data Flow	Replication flows are the main tool for ingesting data into Datasphere from S/4 HANA. Data flows are used to bring data into Datasphere from 3rd party applications. For outbound data processing see 'Premium outbound integration'. *	Heavily used

Premium outbound integration	A lean, high-performance data pipeline from SAP to external object stores. It emphasizes speed, cost-efficiency, and governance alignment	Used for data extraction in all SyWay use cases
BW Bridge	Enables data legacy BW systems to be incorporated into the cloud system	Not used
Data Marketplace	Mechanism to share data between Datasphere tenants and customers	Not used
Semantic Onboarding	Means of sharing pre-defined Datasphere models	Used to install business content as a guide for SyWay development

* For connectivity, replication flows use the cloud connected and data flows use the DPA agent. These are both described in detail in the [Network and Infrastructure Architecture DD-TEC-070](#) document.

Application Security

Authentication

This is described in the [Application Architecture SAP BTP](#) document.

Authorisation

This is described in the [Application Architecture SAP BTP](#) document. Additional information can be found in the [Security Approach](#) and [Security Approach For Analytics](#) documents, with additional information describing the implementation details available in the [SAP Analytics and Reporting Standards](#).

Communication Security

This is described in the [Application Architecture SAP BTP](#) document.

Data Security

Platform level data security is described in the [Application Architecture SAP BTP](#) document.

At an application level, data security is part and parcel of the authorisation approach. As with the authorisations, additional information can be found in the [Security Approach](#) and [Security Approach For Analytics](#) documents, with additional information describing the implementation details available in the [SAP Analytics and Reporting Standards](#).

Special mention should be given to the fact that data security afforded by NextLabs is respected in Datasphere. This is achieved by extracting data from S/4 using a service user with zero sensitive data access. As stated above: Whilst as of 12/11/2025 NextLabs does not work with Datasphere, it is understood that Datasphere is on the NextLabs roadmap with an announcement regarding availability due in 2026.

Other Controls

This is described in the [Application Architecture SAP BTP](#) document.

System Landscape

The system landscape is described at a high level in the [Application Architecture SAP BTP](#) document and in more detail in the [SAP Analytics Approach](#) document.

Operation Architecture

Transport Management

Please refer to [DD-TEC-170 Transport Management for Release 4](#)


Monitoring

This is described at a high in the [Application Architecture SAP BTP](#) document. The following sections give additional detail which will be useful when supporting the application.

Data Loads In Datasphere

Many data loads will be delta replication flows which run on a frequency defined per load (usually hourly). Full data loads will be triggered using Task Chains in DSP and tasks in SAC. Currently these SAC tasks are not integrated with the Task Chains in DSP.

Data loading can be monitored using the Data Integration Monitor



Further Investigation Into Failures Of Data Loads From S/4

If a load is seen to have failed in Datasphere, further investigation can be done in the relevant source system. There are two main jobs responsible for moving data from the source system to Datasphere via the Cloud Connector:

- Observer job (/1DH/OBSERVE_LOGTAB) When new data is posted in the base table, the Observer job pushes it from the master logging table to the subscriber logging table.
- Transfer job (/1DH/PUSH_CDS_DELTA) The Transfer job then moves this data into the buffer table, from there the replication flow picks it up and pushes it to the target system.

Transactions used for monitoring in S/4


- DHCDMON Monitor delta capture process
- DHRDBMON View buffer tables properties and operations
 - Maximum buffer records
 - Current number of records
 - Package size
 - Packages ready for transfer

The buffer table is important because:

- It splits large datasets into smaller, manageable data packages.
- If a package fails, it can be resent, making replication more resilient and reliable.
- Once a package is successfully written to the target, it's committed and deleted from the buffer to free up space.
- It also helps in analysing performance throughput and identifying potential bottlenecks.

System Monitoring

There is a standard system monitor function used to identify data storage, out of memory, CPU capacity issues, etc:



It is also possible to perform additional monitoring functions E.g.

- Enhanced replication flow analysis
 - From \$TEC schema, import the REPLICATIONFLOW_RUN_DETAILS
 - Get all TASK Related Data from DWG_GLOBAL schema and view TASK_LOCKS_V_EXT
 - It is possible to build a model on top of these two tables and view the metadata related to Replication Flows. This allows tracking of key details like execution time, status, and any errors.
- Early watch reports (<https://me.sap.com/ewa/workspace>)
- Integration with Application Lifecycle Monitoring (ALM) where we can review the system loads
- Database Explorer for performance analysis
 - Switch on expensive statement tracing in DSP monitor and find the statement in tables M_EXPENSIVE_STATEMENTS or M_REMOTE_STATEMENTS as appropriate. The captured SQL can be analyzed in PlanViz.
- Reviewing DPA logs in DSP.

SAC reporting on DSP operations

There are SAP Datasphere monitoring views which help you monitor data integration tasks in a more flexible way. They are built on the V_EXT view s, and are enriched with further information as preparation for consumption in an SAP Analytics Cloud story.

Sizing

The estimates in the original [CD-SOL-020 Reporting Approach](#), chapter 8, are referenced and summarised below:

Component	Size	CU(Month)	Comments
Compute blocks	512 GB	13,315	
Storage	1,344 GB	245	
Catalog Storage	0.5 GB	0	
Data Integration	7,200	5,488	trade off with using DPA
Premium Outbound Integration	40 GB	1,000	
BW Bridge			Not considered
Data Lake			Use MS Azure
DPA server			

The DB growth rate has been estimated to be 90GB per annum.

High Availability

Datasphere uses the standard default system availability SLA for **Public Cloud Services** at SAP, which is **99.7%**

Disaster Recovery

SAP Datasphere: Backup/restore follows the **SAP HANA Cloud resiliency layer**; recovery is handled by SAP for disasters within SAP's control.

Backup/Restore

Datasphere performs a back-up of tenants every 15 minutes (RPO 15 minutes). There is also no guaranteed RTO for Datasphere but it is leveraging the SAP HANA Cloud service resiliency layer. Please see note [3574161 - SAP Datasphere Tenant Backup](#)

Maintenance Plan - Release Management

The key features of Datasphere's maintenance plan are:

- **Continuous delivery** in the background: small fixes and security updates can be deployed anytime.
- Major functionality is bundled into **Quarterly Release Cycle (QRC)** updates.
- Customers can choose if they want to adopt **QRC updates immediately** or delay them (to test changes first).
- Updates include **new features, fixes, and security patches**, and they're applied automatically by SAP in the background.
- No customer-side installation or downtime planning is needed.

Shared Responsibility Model

As Datasphere is a SAAS service, there is shared responsibility between SAP and Syensqo. The details of this responsibility sharing are set out in the document '[Hyperscalers: Securing SAP Environments](#)'.

As of 12 Nov 2025 the details are as follows:

Party	Service	Responsibility
Syensqo	Customization & Configuration	Customers must configure and customize the application per their business requirements
	Management of identity and access	Customers must manage the complete identity lifecycle, including onboarding and offboarding users, creating and assigning roles, forming user groups, granting and restricting privilege access, and similar functions for their application
	Data Integrity Requirements	Customers must define proper data classification, storage, and deletion requirements. Although SAP will execute processes on data, defining data requirements is a big part of the customer's responsibility. Protection for data at rest will be assigned by SAP based on the data classification
	Application Audit logs	Customers are responsible for capturing, monitoring, and analysing the application audit logs
	Application compliance	Customers are responsible for industry-specific certification and compliance for data used by or within the application.
SAP	Deploying and configuring Resources	SAP is responsible for deploying and configuring VMs, databases, container images, and the VM operating system.
	Securing VM and images	SAP is responsible for securing and patching operating systems and container images, as well as hardening configurable items on servers and databases
	Logical separation	SAP is responsible for logically segregating applications and data within various environments and between various tenants and customers
	Protecting data	SAP is responsible for implementing data protection, backup, and restoration, based on the data classification. The data retention policy is defined by customer but can be executed by SAP
	Monitoring and incident reporting	SAP logs all the security and infrastructure events. Logs will be aggregated in a system information and event management (SIEM) tool, and an alert will be generated based on the predetermined trigger. SAP will also monitor for incidents and will follow SAP's incident response plan as and when needed.
	Audit and compliance	SAP is responsible for maintaining and providing certification and compliance for the application and related infrastructure.
	Change management	SAP is responsible for managing the maintenance window and other administrative tasks regarding change management
	Availability	SAP is responsible for deploying and maintaining the availability and meeting the SLA
	IaaS	SAP maintains responsibility for the IaaS that the hyperscaler provides on SAP's behalf, and for ensuring each hyperscaler performs as per the contractual agreement
Hyperscaler	Physical security	The hyperscaler is responsible for the physical data centre and the safety and security of people in the data centre. This includes the responsibility for background checks of the people who work in the data center and in connection with other hyperscaler- provided services

	Resiliency	The hyperscaler is responsible for providing the capability of a resilient network and infrastructure across multiple regions and availability zones.
	Physical infrastructure	The hyperscaler is responsible for providing a secure network and infrastructure, including hypervisors
	Audit and compliance	The hyperscaler is responsible for IaaS compliance with industry standards.

Area	Activities
Application security	Application security is the heart of the overall security strategy. Application development at SAP follows the secure development lifecycle. The process starts with planning and assessment, which includes a very important security measure: threat modelling. SAP uses the well-known STRIDE threat modelling technique from Microsoft. Developers follow the secure coding guidelines during the development process. The developed code is reviewed under the "Secure code review" step as a part of the process. Next, a static vulnerability scan is performed on any code developed in-house. Any vulnerability found during the review or scan is mitigated – or documented, if not mitigated – before the release. Software is next scanned for open source vulnerabilities, if any open source libraries or components are used. Dynamic application security testing is performed after software is fully developed and compiled. The last step in the application security is unit testing of the security-related functionality to address issues like invalid input parameters. Once the software is developed and the application is deployed in production, vulnerability scanning is performed at regular intervals and after each new release. Vulnerabilities found during the scanning are managed based on their Common Vulnerabilities and Exposures (CVE) score. SAP does not report or disclose vulnerabilities, but a Service Organization Control 2 (SOC 2) audit report lists any unmitigated vulnerabilities. The SOC 2 report can be obtained from SAP.
Data Security	The customer defines the data protection, retention, backup, and deletion requirements. SAP is responsible for making sure that tenant data is logically segregated. SAP also makes sure that data is segregated between nonproduction and production environments. Encryption As per the SAP security policy, data in transit and data at rest should always be encrypted. Any communication between the hyperscaler and client uses Transport Layer Security (TLS) with HTTPS. Data at rest is encrypted using disk encryption to prevent data exposure in case of a physical theft of the drive. Other encryption methods, such as volume, backup, or in-application encryption, are used based on the technical, functional, and business requirements of the application and customer. Encryption Key Management SAP does not utilize default keys provided by hyperscalers. SAP is responsible for creating, rotating, and deleting the encryption keys. SAP also manages access to the key. One of the "key" differences between an application hosted by SAP versus third-party hyperscalers is the key storage. When an SAP application is hosted by a third-party hyperscaler, the key is stored with the hyperscaler using the hardware security module (HSM) or other secret management storage that the hyperscaler provides. This key storage or HSM is always FIPS 140-2 compliant. Any access to this storage is logged and audited by SAP. The encryption key is always managed by SAP, regardless of where the key is stored. Retention, Deletion, and Backup Data retention with most SAP applications is automated and customer driven. Customers can create policies or rules in the application stating how long the data should be retained based on their requirements. Data will be deleted at the end of the retention period. Customers can also delete their data at any time they have access. Data backup and deletion processes and schedules are not impacted by the migration to hyperscaler. These processes remain unchanged. It is important to note that SAP and hyperscalers will maintain compliance with laws and requirements around personal data, such as EU access, the General Data Protection Regulation, and other industry and geographic regulations.
Infrastructure and Network Security	SAP creates virtual resources using cloud APIs and is responsible for everything between and including virtual resources and the application. SAP will deploy and manage everything from the virtual machine up. This means that SAP has responsibility for managing infrastructure, creating and managing various virtual private clouds, and creating and managing security groups and firewalls. SAP is also responsible for managing and patching the operating system and middleware. SAP will regularly scan the environment for operating system and middleware vulnerabilities. SAP will deploy patches to operating systems and middleware based on the vendors' specifications. SAP's architecture blueprint dictates that database servers and application servers are isolated from each other and from the public-facing Web server. DB server and application servers are hosted within a private subnet, while Web servers are in the public subnets behind the Web application firewall (WAF) and security groups. SAP's strategy is to provide database clusters. High availability will not change as a result of migration to a hyperscaler. Hyperscalers are responsible for providing overall network and infrastructure protection against DDoS and network- or infrastructure-based attacks to the data centres, but it is SAP's responsibility to provide anti-DDoS, IPS/IDS, WAF, and network monitoring of the resources created by SAP. It is SAP's responsibility to perform regular penetration testing, and SAP will work with the hyperscaler for network penetration testing. The physical security of the data centres and vetting of the workforce who are working in and around data centres are responsibilities of the hyperscaler.
Logging, Monitoring, and Incident Response	The customer has full access to application and audit logs. SAP is responsible for collecting, storing, and analysing infrastructure and security logs. SAP manages the threat triggers and generates alerts from the logs. SAP does not share infrastructure and security logs with customers. SAP aggregates the logs into the SIEM tool and automates the process of analysing and generating alerts. Monitoring various logs and generating alerts when there is a deviation from the baseline is a very time-consuming but essential part of the security – and SAP handles that for you, so you can focus on your customers. The team of seasoned SAP professionals perform infrastructure monitoring, database monitoring, security incident management, secure admin access, regular backups, security scanning and remediation 24x7 to secure the environment for customers. Hyperscaler landscapes pose unique challenges, and SAP's security incident response team works closely together with GCS multi-cloud security operations to continuously improve security incident response process and automation for SAP's multi-cloud landscape. Although SAP does not notify customers of every incident, we will provide breach notification report and root cause analysis to customers for any incident that is classified as a personal data breach.

Identity and Access Management	The customer is responsible for identity and access management (IAM). SAP provides single sign-on and other IAM-related services as needed. SAP offers solutions that can manage the complete identity lifecycle, integrate on-premise and cloud solutions, work with multi-factor authentication, and simplify the access management process for you. The customer has complete control over who can access the data and to what extent. Most important, the customer has the ability to provide admin or privileged access to the application. This access should be granted only as needed and must be monitored. SAP has access to cloud accounts as well as privileged access to the application and SAP environment within the hyperscaler environment. SAP employees or partners do not have any access to customer's data or information.
Connectivity to Cloud	Azure ExpressRoute allows you to extend your corporate or personal network into the Microsoft cloud over a private connection. Azure ExpressRoute provides Layer 3 connectivity between your site and Microsoft cloud. Azure ExpressRoute provides redundancy for the network connection as well as a guaranteed uptime SLA for connectivity.

Additional information can be found at [SAP's Cloud Services Service Level Agreement](#), specifically the document 'Service Level Agreement for Private Cloud Edition Services and Tailored Option Services/'

Exceptions

See also

[SAP Analytics Approach](#)

File	Modified
PDF File Approval by Frank Bolata 2025-12-17.pdf	Dec 18, 2025 by WENNINGER-ext, Sascha
PDF File Stakeholder endorsement - Owen Pettiford.pdf	Dec 05, 2025 by WENNINGER-ext, Sascha
PDF File Stakeholder endorsement - Werner Dahn.pdf	Dec 05, 2025 by WENNINGER-ext, Sascha
File DSP Integration draw.io diagram	Nov 06, 2025 by SHEPHERD-ext, Robert
File ~DSP Integration.tmp draw.io Draft	Nov 06, 2025 by SHEPHERD-ext, Robert
File drawio-backup-DSP Integration-rev-4 draw.io diagram backup	Sept 22, 2025 by BARROW-ext, ian
File PaPM draw.io diagram	Sept 11, 2025 by BARROW-ext, ian
File ~PaPM.tmp draw.io Draft	Sept 11, 2025 by BARROW-ext, ian
File Cloud_connector_DSP draw.io diagram	Aug 21, 2025 by BARROW-ext, ian
File ~Cloud_connector_DSP.tmp draw.io Draft	Aug 21, 2025 by BARROW-ext, ian

[Download All](#)

Change log

Version	Published	Changed By	Comment
CURRENT (v. 129)	Feb 04, 2026 14:54	SHEPHERD-ext, Robert	
v. 128	Dec 05, 2025 11:35	WENNINGER-ext, Sascha	added stakeholders
v. 127	Dec 05, 2025 09:20	SHEPHERD-ext, Robert	
v. 126	Nov 26, 2025 09:19	BARROW-ext, ian	
v. 125	Nov 13, 2025 16:38	WENNINGER-ext, Sascha	
v. 124	Nov 13, 2025 15:25	SHEPHERD-ext, Robert	

v. 123	Nov 13, 2025 15:19	SHEPHERD-ext, Robert
v. 122	Nov 13, 2025 15:09	SHEPHERD-ext, Robert
v. 121	Nov 13, 2025 15:04	SHEPHERD-ext, Robert
v. 120	Nov 13, 2025 15:01	SHEPHERD-ext, Robert

[Go to Page History](#)

Workflow history

This view shows the 5 most recent entries. The complete workflow log is available from the 'Document Activity' menu item.

Feb 04, 2026	Actor	Type	Activity	Version
Approved	 SHEPHERD-ext, Robert	Edit	updated the page at 2:54 pm	
Dec 18, 2025				
	WENNINGER-ext, Sascha	State	changed state to Approved at 2:52 am	v128
Pending SteerCo Review	WENNINGER-ext, Sascha	State	gave <i>Final Approval</i> approval at 2:52 am Approved by Frank Bolata. Email attached	
Dec 05, 2025				
	 PETTIFORD-ext, owen	State	changed expiry date to '19 Dec, 2025 02:23 pm' at 2:23 pm	
		State	changed state to Pending SteerCo Review at 2:23 pm	v128
Pending Stakeholder Review	 PETTIFORD-ext, owen	State	gave <i>Stakeholder Review</i> approval at 2:23 pm	