

SAP Data Warehouse Cloud

By Lars van der Goes
& Dirk Kemper



expertum

White Paper

Table of contents

What is SAP Data Warehouse Cloud?	3
A collaborative environment: the power of Spaces	6
End-to-end integration	8
The Data Marketplace	11
One Catalog Vision – The Enterprise Business Catalog	13
Speed, elasticity and security	14
Access & Authorizations	16
BW Bridge	18
Limitations	20
Architecture options	21
Pricing and licensing	23

What is SAP Data Warehouse Cloud?

SAP Data Warehouse Cloud was conceived by SAP as a product to close a specific gap in its portfolio, namely that of a data warehouse-related solution in the Cloud domain.

As part of the HANA Cloud Services, consisting of SAP HANA Cloud, SAP Analytics Cloud and SAP Data Warehouse Cloud, SAP Data Warehouse Cloud fulfils the 'Quality'-aspect of the data value equation that SAP utilizes to illustrate the direction of its Cloud-Analytics portfolio (Figure 1).

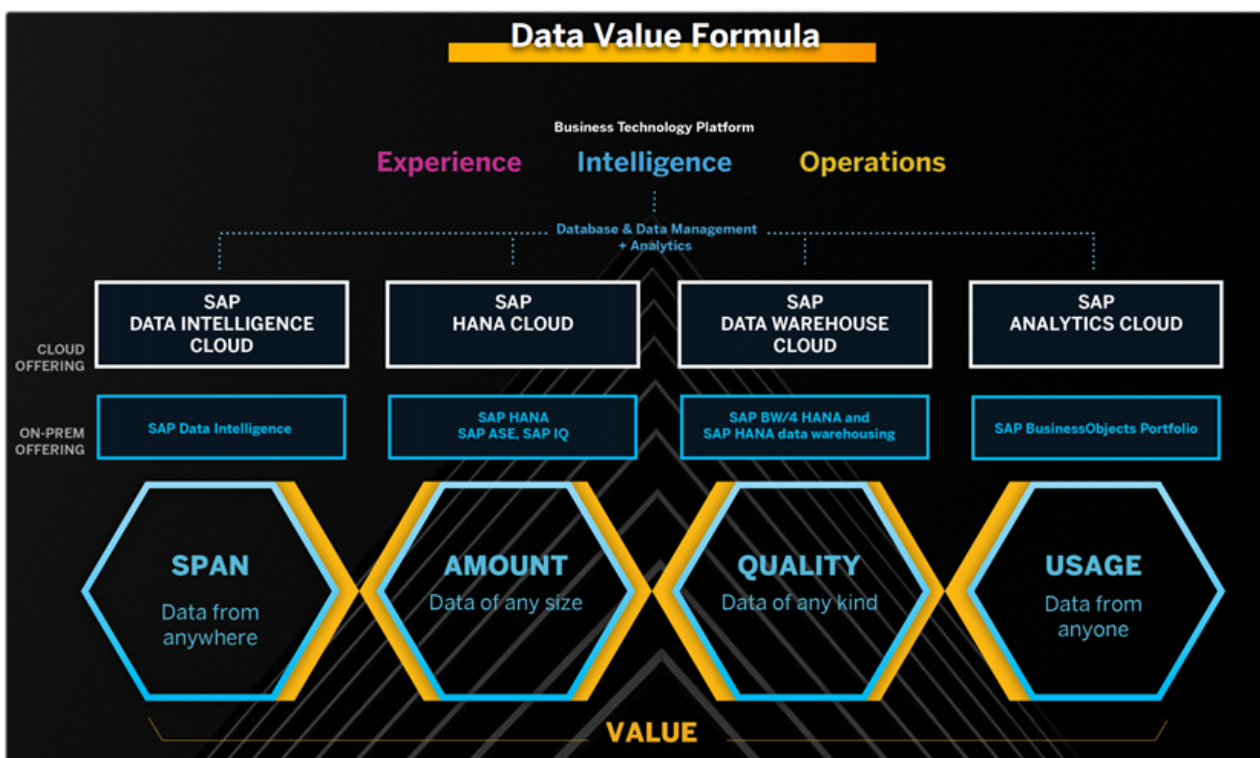


Figure 1 - SAP's Analytics solutions in the context of the Data Value Formula

Data Warehouse Cloud was not initially designed as a replacement of the existing on-premise BI solutions (SAP BW/4HANA and HANA for SQL data warehousing), but rather as a complementary solution in the Cloud domain. SAP published a BW/SAP Data Warehouse Cloud overview-one pager to illustrate this principle, in which both the On-Premise and Cloud preferences (best practices) are illustrated.

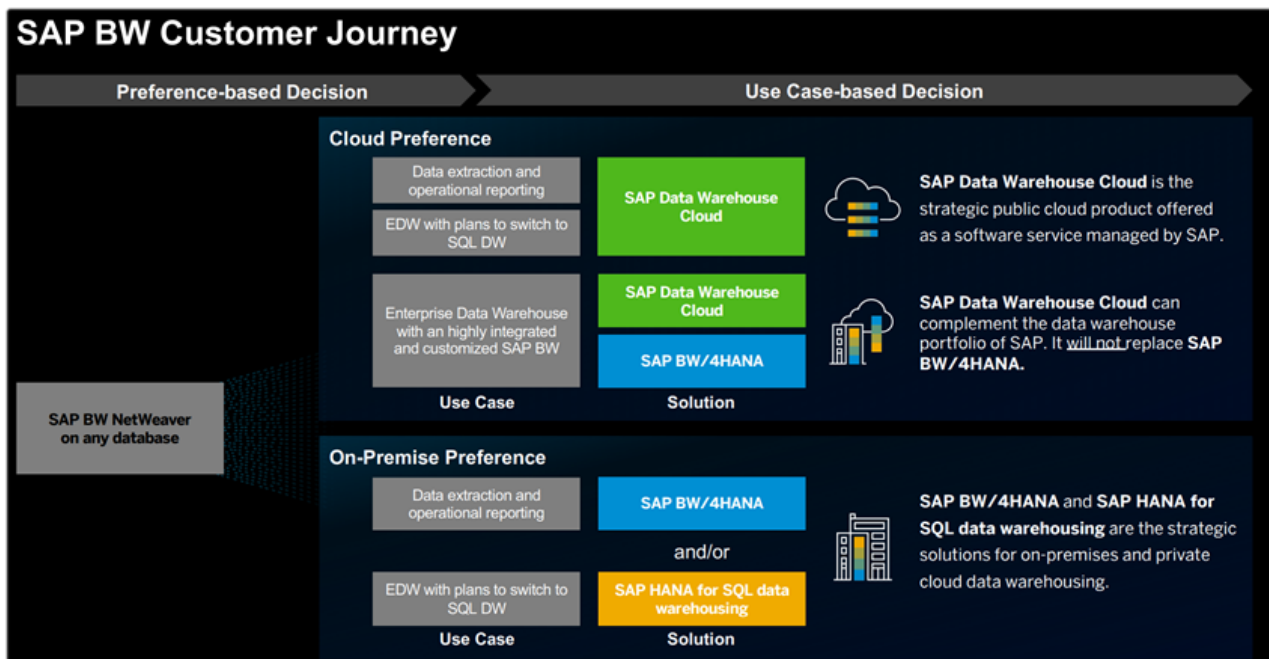


Figure 2 - SAP's BW/SAP Data Warehouse Cloud overview one pager

This initial positioning would be adjusted however, as SAP announced that SAP Data Warehouse Cloud would become the strategic successor of SAP BW/4HANA. This strategic move was made in conjunction with the reveal of SAP BW Bridge; more on this SAP Data Warehouse Cloud feature can be read later in this white paper. This does not mean however, that SAP Data Warehouse Cloud is replacing SAP BW/4HANA. SAP BW 7.5 will remain supported until at least 2027 in terms of mainstream maintenance (2030 for extended maintenance), whereas for SAP BW/4HANA this support extends to at least 2040. SAP BW/4HANA and SAP HANA for SQL data warehousing remain SAP's strategic solutions for on-premises and private cloud data warehousing, while SAP Data Warehouse Cloud will be the strategic public cloud product for Data Warehousing moving forward (although the latest innovations will come to the Cloud first). At Expertum, we believe that SAP Data Warehouse Cloud is now primarily intended to be SAP's answer to competitors such as Snowflake, Google's BigQuery and Amazon's Redshift.

While conceiving SAP Data Warehouse Cloud as a new solution, SAP evaluated the most recurring challenges faced by both IT and Business in the data-related parts of the IT landscape. The problems they identified and offset with the solution's cornerstones, are the following:

- Limited data access (for Line of Business/information users), tackled by creating a collaborative environment.
- Lengthy time to value, tackled by being end-to-end integrated.
- Lack of real-time data, tackled by being fast and responsive.
- High entry costs (for dedicated BI/analytics solutions), tackled by being elastic in capacity.

SAP thus designed SAP Data Warehouse Cloud to address these problems while simultaneously filling the strategic gap in their solution portfolio. In line with this direction, SAP's BI Cloud landscape has been moving to facilitate a broader group of users, ranging from executives and decision makers to planners, general IT, business & data analysts and finally, information workers. SAP defines the latter group as the broad group of company employees that interact with data coming from (SAP) BI solutions.

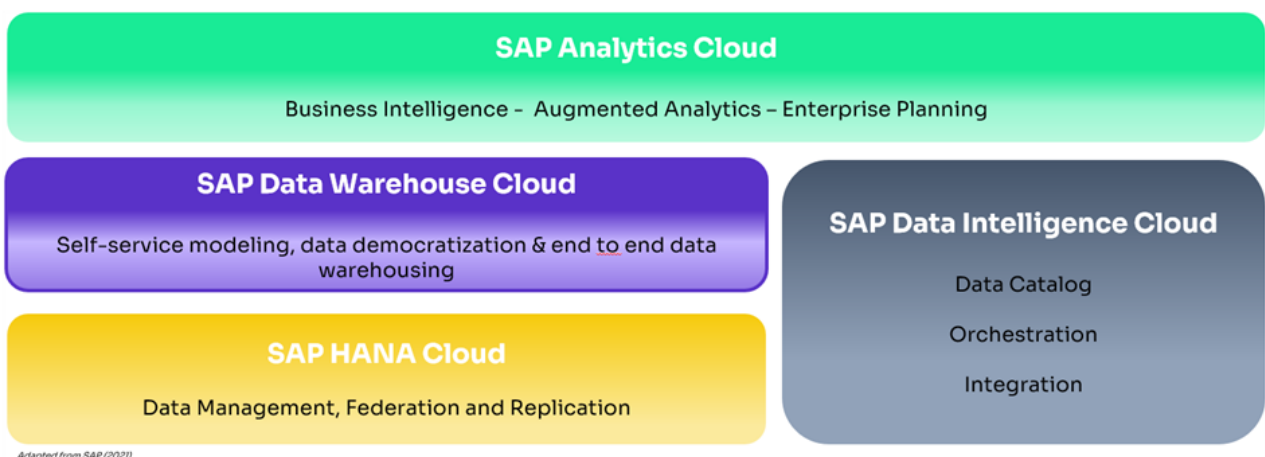


Figure 3 - SAP's BW/SAP Data Warehouse Cloud overview one pager

A collaborative environment: the power of Spaces

Data Warehouse Cloud’s terminology and its familiar SAC-like user interface offer an equally accessible canvas to IT and business users alike. If an organisation uses multiple current-generation SAP solutions (e.g. SAP S/4HANA, SAP C/4HANA, SAP Hybris, etc.), users will be instantly familiar with the layout of SAP Data Warehouse Cloud. Administration menu options, such as Security, User & Role management or Audit options, can be assigned to IT users without visually limiting business users.

The SAP Data Warehouse Cloud concept of Spaces allows IT and Lines of Business (LoB) users/information workers to operate within the same entity, while both parties can view them from different angles. For Business users, Spaces allow them to wrangle, enrich, visualize and reuse data (for example by blending and creating new or different business semantics) within a flexible environment. IT can think of Spaces as a database schema: they manage and allocate (provision) data, where each Space is isolated (as a modelling environment), but data can be shared across them. This allows for security (through versioning and administration within the solution) without scattering different versions of truth across the IT landscape. In turn, organizations are prevented from losing sight of their single source of truth and discourage the creation or maintenance of Shadow-IT initiatives.

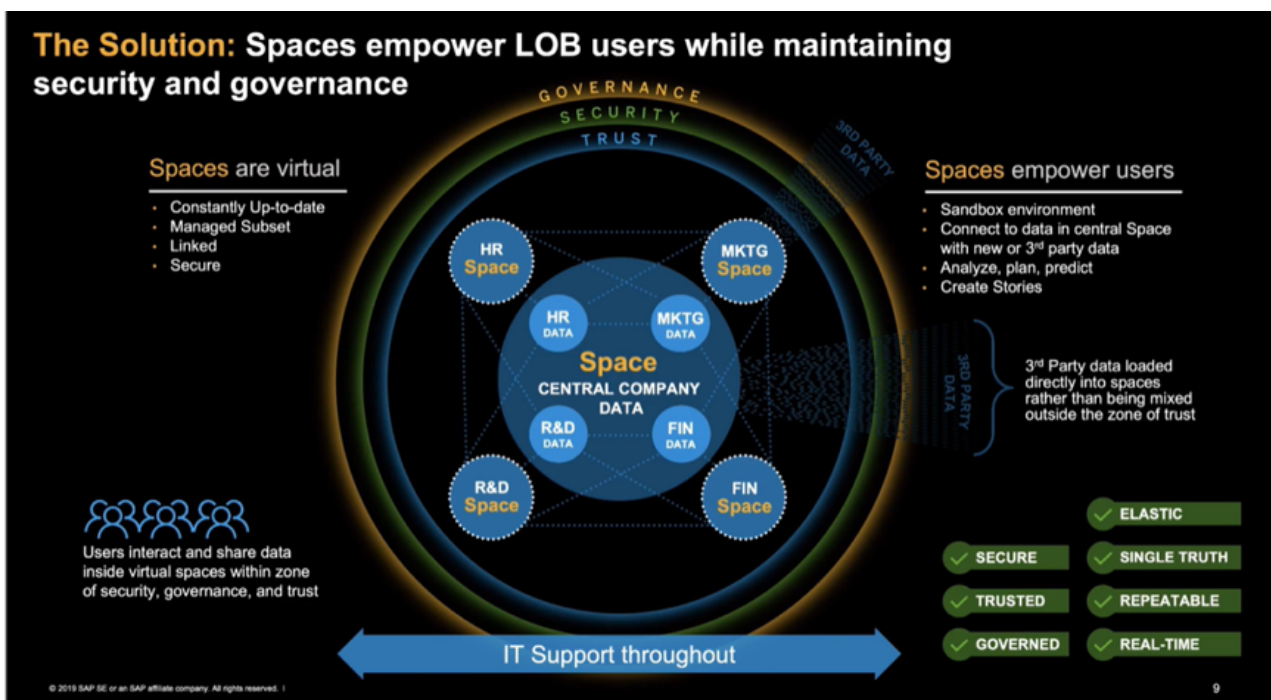


Figure 4 - The concept of Spaces explained

Spaces also create a physical separation of data(sets), in which both disk and memory capacity are allocated and limited per Space. Spaces can either be created as HANA (Cloud) Spaces, Disk-based Spaces or Data Lake Spaces, the latter of which does not rely on in-memory processing and needs to be licensed through dedicated data lake-storage capacity.

Through this principle, SAP Data Warehouse Cloud prohibits Spaces from using more than their assigned capacity, keeping data usage within the license agreement. This new way of data management is different from both SAP HANA and SAP BW/4HANA, in that it provides an effective method for controlling where and in what amounts data is being stored.

End-to-end integration

SAP proposes SAP Data Warehouse Cloud as end-to-end, in the sense that it is able to integrate different sources (SAP and non-SAP), store the data in accordance with the general tiering principles in the underlying HANA-database, provide IT and business users with business-level modelling options including data-lifecycle management and finally, allow the data to be used in analytic solutions either through embedded SAC or third-party applications such as Power BI or Tableau.

From this perspective, SAP Data Warehouse Cloud can be holistically seen as a data mart, with its multitude of integration options as one of its main strengths.

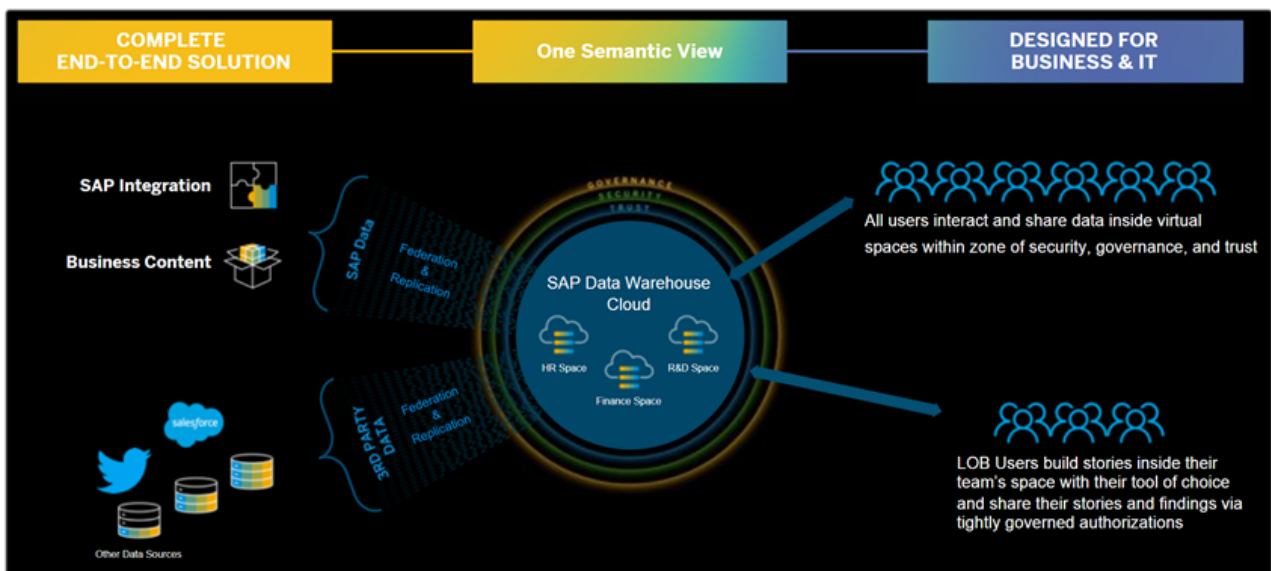


Figure 5 - An illustrated example of an end-to-end dataflow through SAP Data Warehouse Cloud

Thus, SAP Data Warehouse Cloud's success comes or goes with the degree to which it can be integrated into customer IT landscapes. The packaged API framework and tight integration with SAP's Cloud portfolio via both data streaming and data replication connections form the backbone for integration opportunities in SAP Data Warehouse Cloud. For general data extraction, SAP Data Warehouse Cloud offers a generic database connector as well as an ODBC interface. For BW, the technology used in SAP Data Warehouse Cloud to facilitate direct integration is similar to Data Intelligence's (DI) Data Flow functionality (as opposed to the more commonly used ODP framework). Smart Data Access (SDA) as well as its extension Smart Data Integration (SDI) for federated data management and data consolidation, can also be used in conjunction with SAP Data Warehouse Cloud and come packed with the solution.

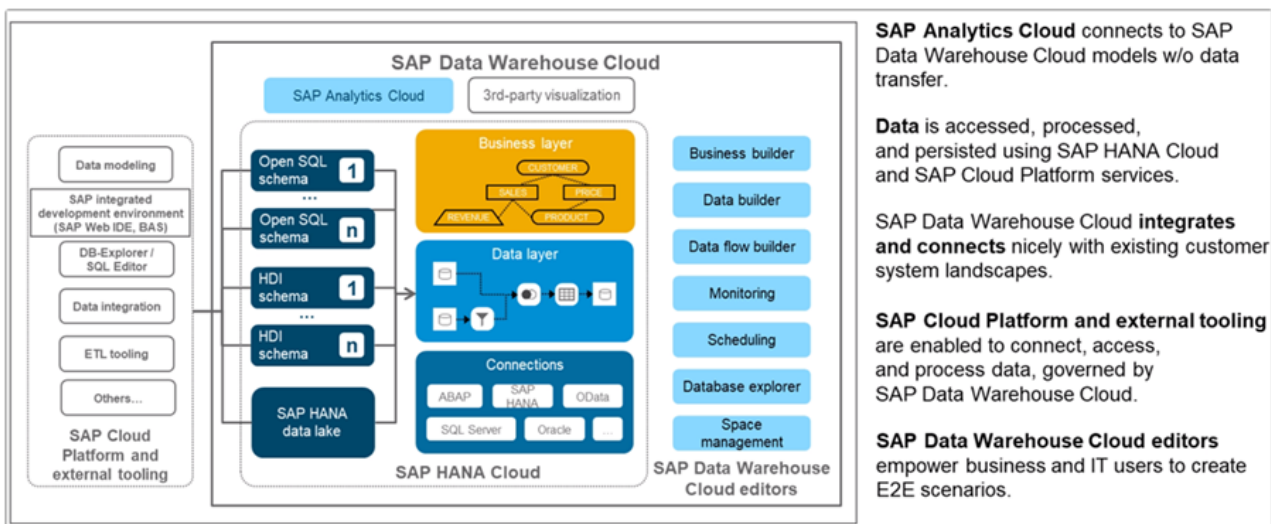


Figure 6 - SAP Data Warehouse Cloud SAP-native integration

The most prominent peer solution to SAP Data Warehouse Cloud, SAP Analytics Cloud, is able to either use live data connections (in which data does not get transferred between systems) or direct import connections in conjunction with SAP Data Warehouse Cloud. As stated before, SAP Data Warehouse Cloud also comes with embedded SAC as part of the solution (Story-functionality only) and five licenses for the standalone SAC product to get customers started.

Another interesting feature is the SAP DB-Explorer; an integrated solution that leverages (open)SQL statements and provides database users in SAP Data Warehouse Cloud with an Eclipse-like data browser for connected database systems (e.g. HANA or HANA Data Lake). For example, it can leverage SQL statements and allow customers to connect (to), access and process data which is governed by SAP Data Warehouse Cloud (no data transport necessary; streaming of data for consumption).

Finally, SAP Data Warehouse Cloud integration can be conducted through SAP Data Services software, SAP Data Intelligence and certified partner tooling that can interact with HANA (Cloud), ODBC or JDBC. The idea here is that existing ETL landscapes can be (partially) reused when implementing SAP Data Warehouse Cloud.

The Data Marketplace

SAP Data Warehouse Cloud's Data Marketplace is a one-stop shop for external integration and harmonization scenarios for and between SAP Data Warehouse Cloud users. This feature within SAP Data Warehouse Cloud supports various business scenarios and outside-in use cases, allowing for quick integration of specialised data that does not require extensive planning or building effort. The general principle is that users of SAP Data Warehouse Cloud can act as a Data Provider within the Marketplace, whether they are using the solution as a data warehouse or just as a platform for sharing their data in the Marketplace against a reduced (non-productive) license fee. Data Providers share their data in the Data Sharing Cockpit in the form of a Data Product, which can be leveraged by Data Consumers in Spaces through the Marketplace. A high-level overview of this process is shown in figure 8.

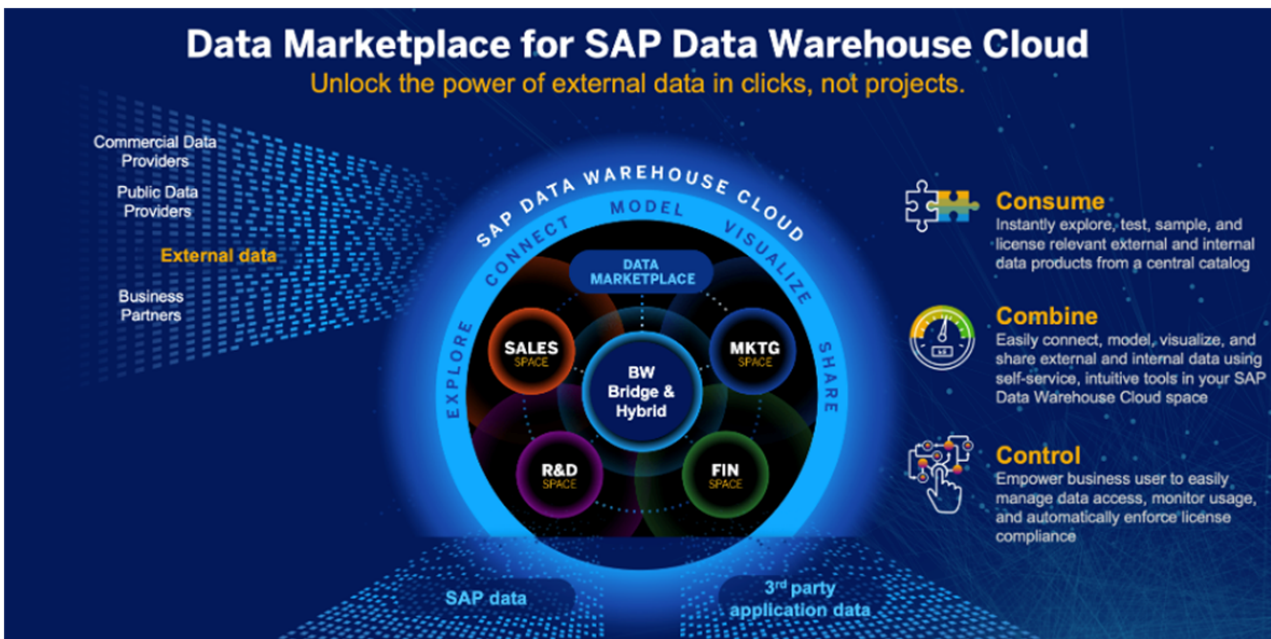


Figure 7 - An overview of SAP Data Warehouse Cloud Marketplace

The Data Marketplace hosts both SAP and non-SAP data providers such as PWC, EY and Datazeit, among others. SAP's vision is that by offering a central place where users can retrieve these plug-and-play data products from commercial and public data providers as well as business partners, the need for integration projects and API configurations to combine internal (organizational) data with external data will be a thing of the past. To smoothen things out, the Data Marketplace also allows users to share their Space data without actual data distribution and security risks, as it can provide insights into the final dataset based on the authorizations of the utilized data product.

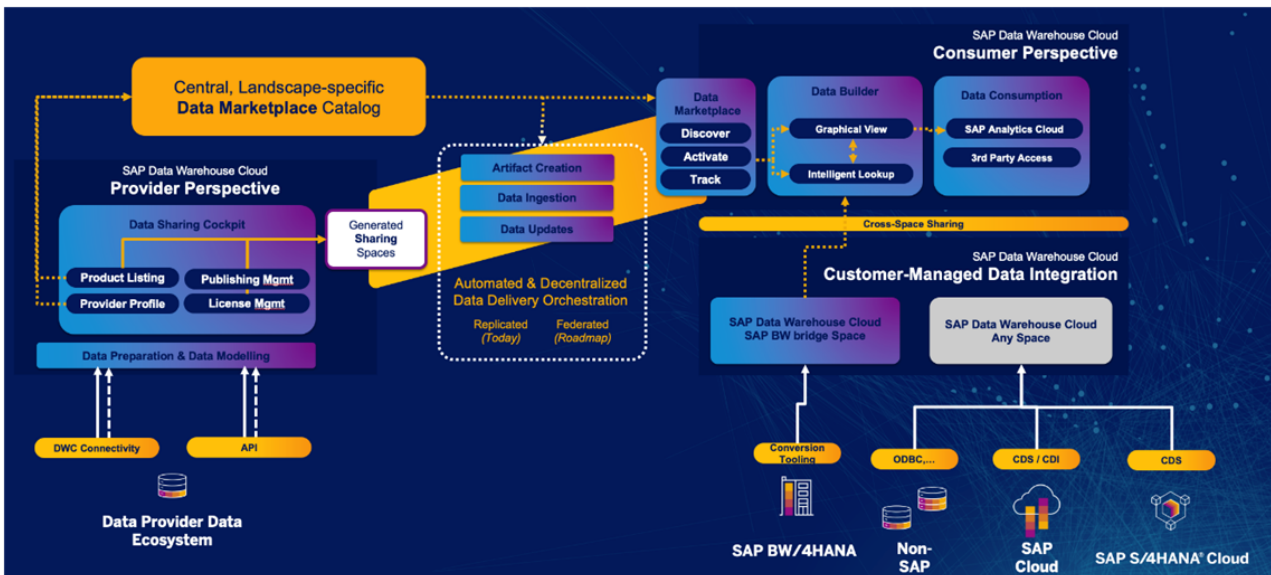


Figure 8 – How the Data Marketplace works

One Catalog Vision - The Enterprise Business Catalog

SAP plans to launch the Enterprise Business Data Catalog, a central point of access for analytical assets within the IT landscape, in the second half of 2022. It is presented as a unified enterprise catalog; this single point of access facilitates the management of data both inside and outside SAP's Cloud Solutions to facilitate transparency and collaboration. The current scope is planned as displayed in the image below.

The vision behind the Catalog is a solution that consolidates SAP's existing cataloging solutions such as SAP Analytics Hub and SAP Data Intelligence Data Catalog. As part of this Catalog, SAP is also working on a Content Transfer solution to represent all S/4HANA objects natively in SAP Data Warehouse Cloud. These objects are subsequently also available in the Enterprise Business Catalog through SAP Data Warehouse Cloud.

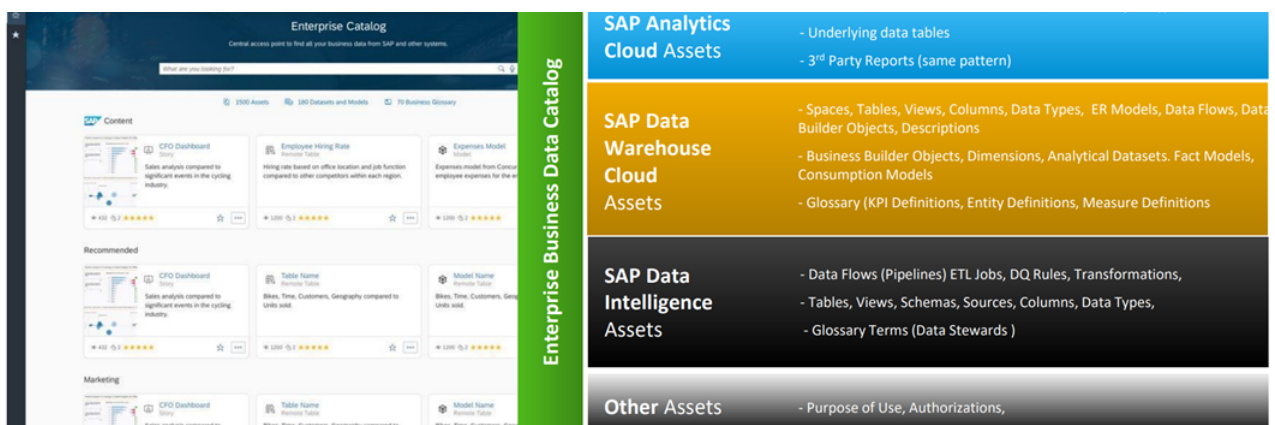


Figure 9 – The proposed scope of the Enterprise Business Catalog

Speed, elasticity and security

Being fast is a major selling-point for SAP Data Warehouse Cloud. With the HANA database at its core, SAP positions SAP Data Warehouse Cloud to be a fast offering running on the HANA Cloud Engine. An additional advantage from the underlying HANA Cloud database, is that Machine Learning and Predictive functionalities are quickly accessible via enablement of APL and PAL content within the solution.

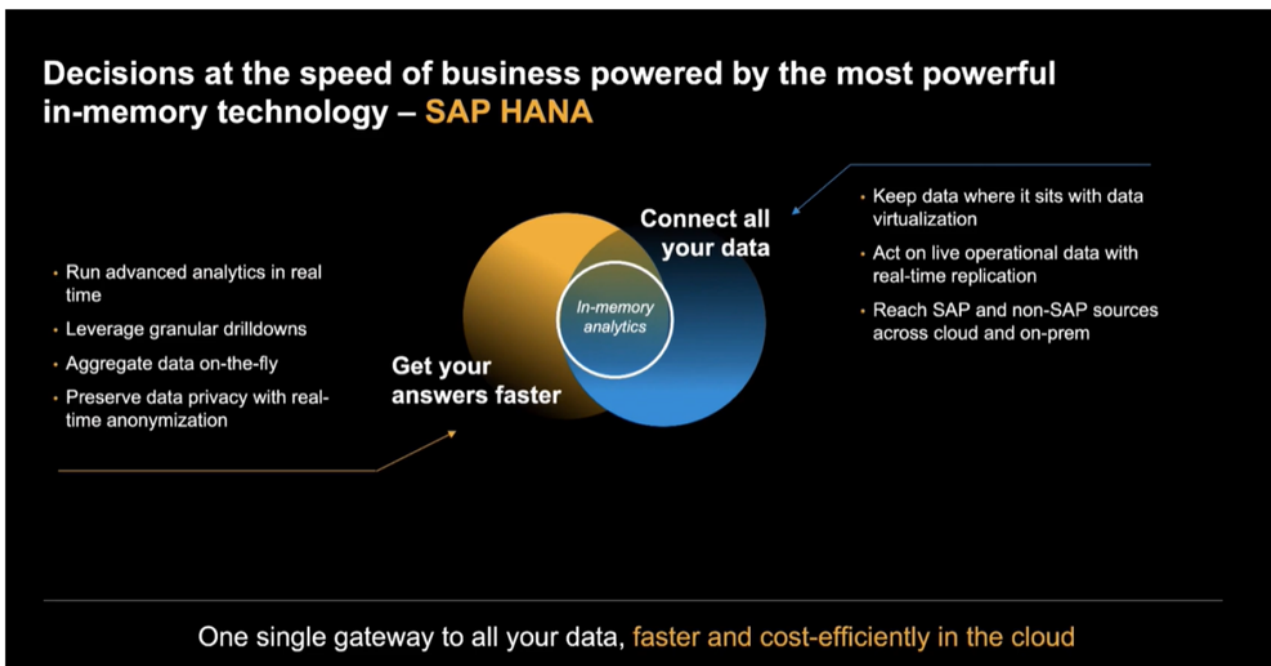


Figure 10 - HANA as SAP Data Warehouse Cloud's speeding engine

Another one of SAP Data Warehouse Cloud's biggest selling points, is its scalability. SAP promises a transparent, clear overview of both storage and memory capacity and its related cost (more on licensing later). This idea goes hand-in-hand with the data tiering concept (from in-memory for hot data to HDFS-storage for cold data) in SAP Data Warehouse Cloud and translates to a solution that can be tailored to an organisation's information needs. In practice, it is unlikely that an organization's information needs will ever call for lower data volumes, with the result that a downscaling of system size is just as unlikely.

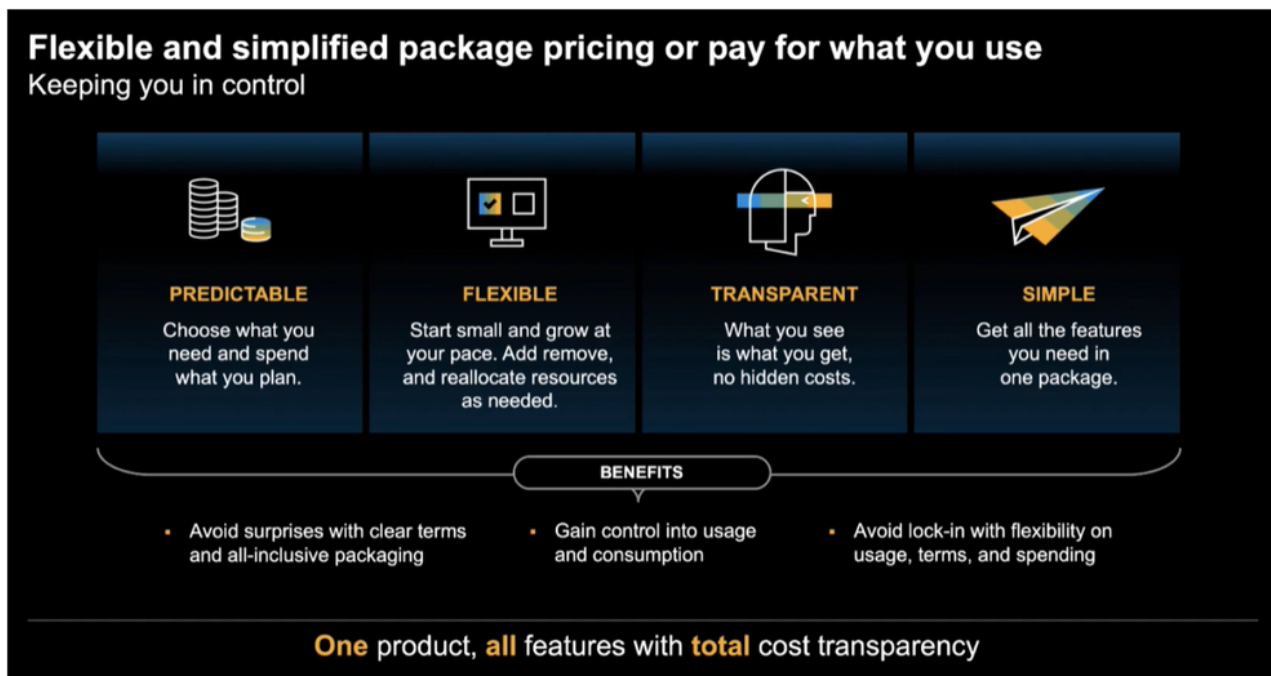


Figure 11 - SAP Data Warehouse Cloud's flexibility concept

Access & Authorizations

In SAP Data Warehouse Cloud, user management and security are governed by IT users. While SAP's Cloud Identity and SSO are the standard for authentication across SAP systems, SAP Data Warehouse Cloud's internal security is arranged via two 'Building Blocks'; Data Access Control (which is a separate menu tab in the system) and authorization scenarios.

Data Access Control allows for granular control on row-level data privileges and is applied on objects in the data layer of SAP Data Warehouse Cloud. Since Data Access Control is an IT-privilege, it cannot be overruled in the business layer of SAP Data Warehouse Cloud. Authorization Scenarios can further restrict upon Data Access Controls. They define the context in which data is consumed and Data Access Control is applied and act as a 'gatekeeper' for consumption. An overview of a 'Building Blocks'-scenario, based on a geographical dimension, is shown below.

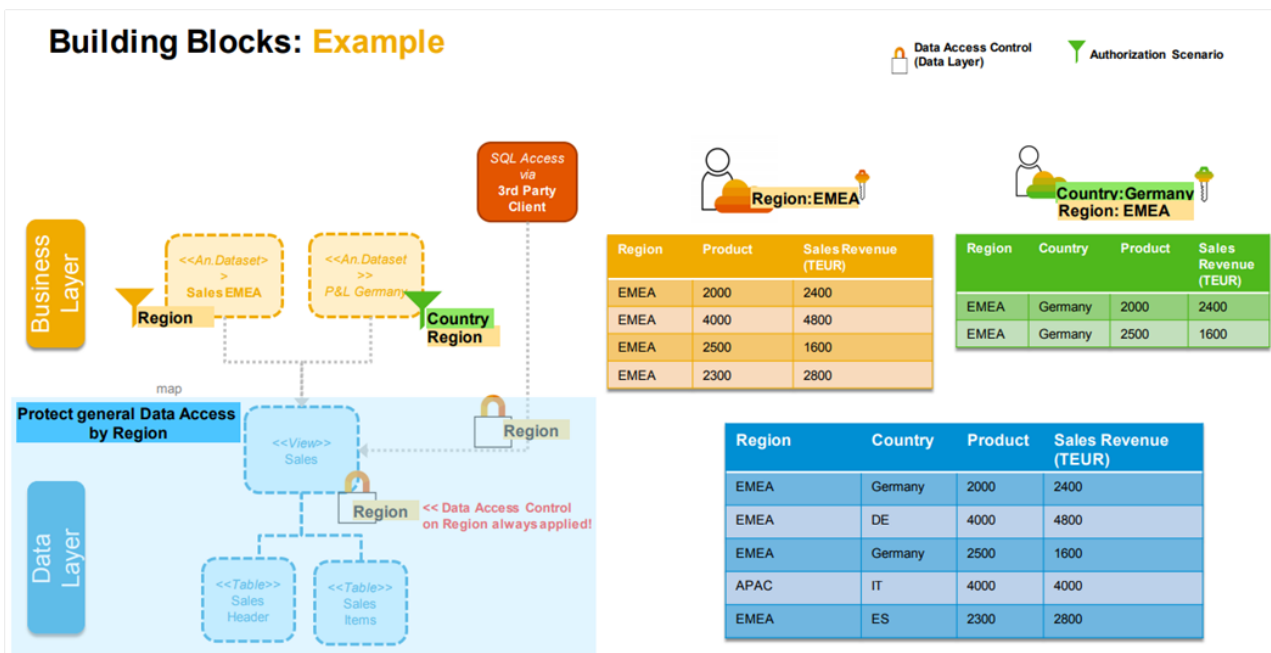


Figure 12 - A SAP Data Warehouse Cloud authorization example

SAP Data Warehouse Cloud is also able to leverage existing authorizations from BW/4HANA. SAP currently offers two main possibilities to realize this:

- Accessing the data in the source system in a federated way where a technical user (with a well-defined set of authorizations) can be used to access the data from a given source as it is today.
- Leverage existing authorizations from the source by replicating them to SAP Data Warehouse Cloud, generate corresponding Authorizations (in case of SAP Data Warehouse Cloud 'Data Access Controls') and apply them in SAP Data Warehouse Cloud.

Both options leverage SAP Data Warehouse Cloud's Data Access Controls, as well as dedicated SQL scripts to generate secure views on BW objects such as InfoProviders. Thus, remote authorizations for SAP Data Warehouse Cloud from SAP BW/4HANA should help customers implement authorization rules at row-level security, by leveraging investments already made in SAP BW/4HANA. More source systems (i.e. SAP S/4HANA, SAP NetWeaver BW 7.5) have been named by SAP as future candidates to support this functionality.

BW Bridge

SAP recently announced BW Bridge, a migration solution with the purpose of translating BW content to SAP Data Warehouse Cloud. Essentially part of SAP's 'RISE'-direction, this solution is planned for release in Q1 2022 and allows customers to 'migrate' their BW content to the Cloud, where a BW-compatible engine running within SAP Data Warehouse Cloud will allow users to move their BW objects (InfoObjects, InfoProviders, transformations) directly into Data Warehouse Cloud. SAP aims to solve three of the most frequently defined gaps as experienced by customers:

- How to leverage existing SAP BW assets and skills?
- How to retain instant access to SAP data while expanding data reach?
- How to empower users and scale insights across the enterprise?

The idea hence is to reuse the investments related to BW for the future, while safeguarding and maintaining the governance (authorization) concept build over potentially years of BW developments.

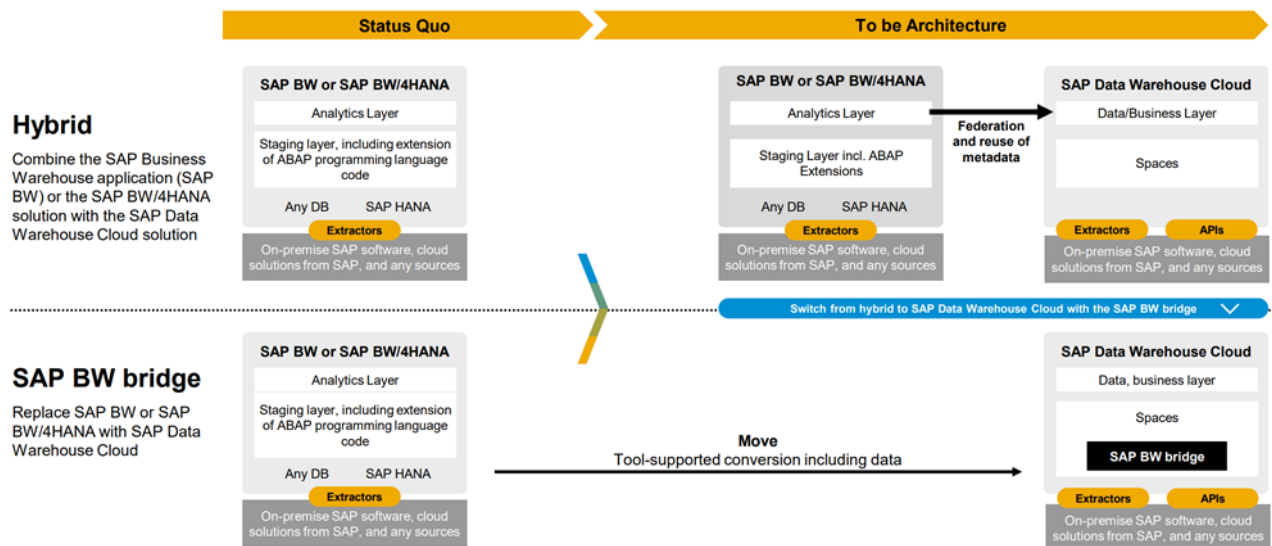


Figure 13 – The main BW Bridge scenarios

BW Bridge's main function opens up a tool-supported move which physically places all existing BW objects in SAP Data Warehouse Cloud in their existing format.

Subsequently, the BW Bridge conversion tool will translate the BW objects to their SAP Data Warehouse Cloud artifact equivalents (more information in the graph below). Customers can either move to this setup with a new or an existing SAP Data Warehouse Cloud implementation, the latter of which could already be in place through a hybrid setup as indicated in the top half of figure 13. The result is a fully-fledged ABAP-based BW instance in SAP Data Warehouse Cloud, which supports BW cornerstone features such as DTPs, Transformations and Process Chains, but omits BW reporting functionalities. BW Bridge also supports the usual BW(4) objects (such as ADSOs and Composite Providers) and will have a dedicated Space in SAP Data Warehouse Cloud, running in parallel to the other Spaces in the solution. Through this method, customers can use all existing data and data flows they could previously, without an on-premise installation. BW Bridge will also have its own lifecycle management for DEV – PRD. BW Bridge’s tool supported move will be available in Q1 2022 for BW/4HANA, 7.5 and 7.4 customers, whereas support for the even older 7.3 version is scheduled for release later in the year.

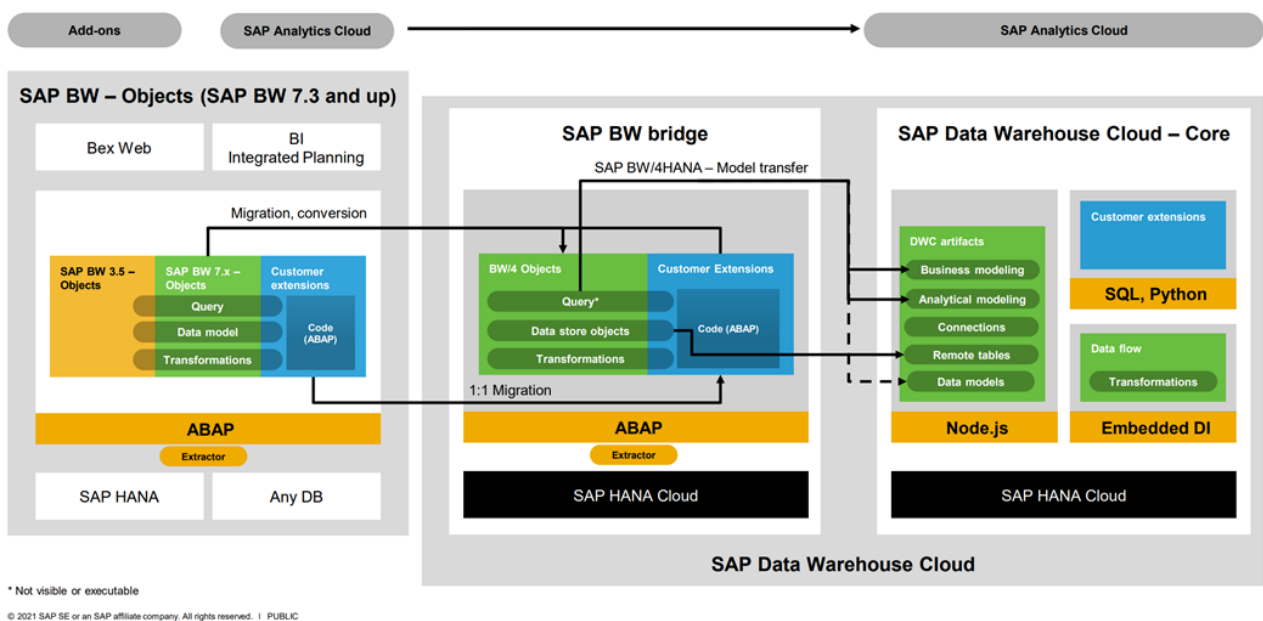


Figure 14 – Object conversion through BW Bridge

Limitations

If we take a closer look at BW Bridge, we can see that conversion options only work with BW versions of 7.3 and up. Although migration paths for this 2010-version are possible, the rather long journey would make BW Bridge a costly option for customers with an older SAP BW system.

Logically, this also means that legacy SAP BW objects such as DSOs and InfoCubes are not supported. Queries are also not visible or executable, but will have to be converted as part of the SAP BW data models before they can be used in SAP Data Warehouse Cloud's Business and Analytical modelling. Another question that pops up with these migration options, is on how older objects (pre-SAP BW/4HANA) will be handled by BW Bridge, which is currently still unclear as the solution is yet to be fully released.

In relation to other peer systems; SAP Business Planning and Consolidation (BPC) and BI-Integrated Planning will also not be supported, as SAP Analytics Cloud in combination with SAP Data Warehouse Cloud as a foundation and Group Reporting will be the respective alternatives for these use cases.

Architecture options

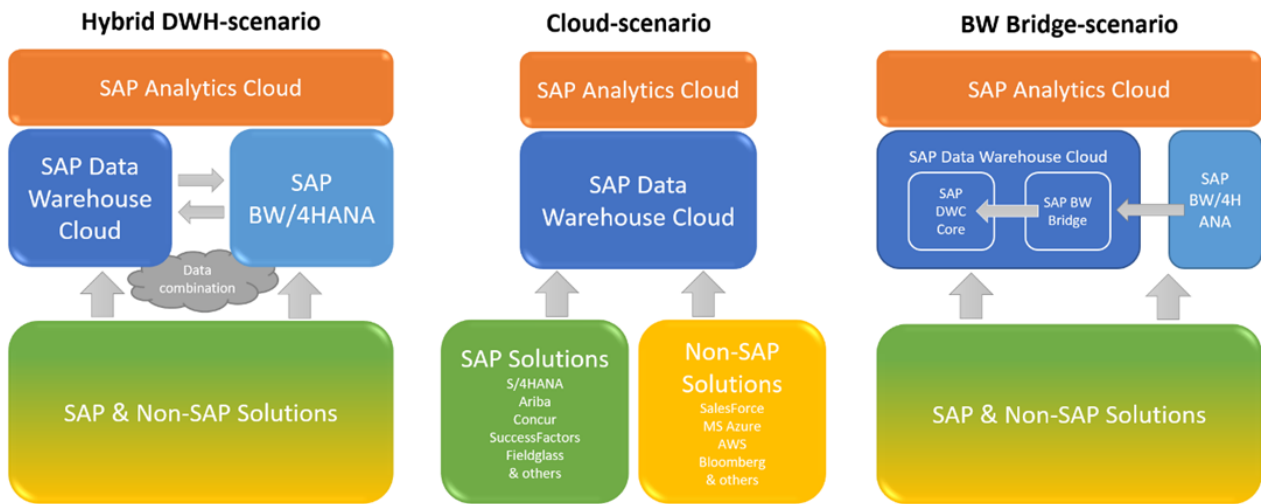


Figure 15 – Architecture options

Now that we have explored the key components and functionalities of SAP Data Warehouse Cloud, let's have a look at the architecture options.

The first scenario that SAP proposed after launching SAP Data Warehouse Cloud, the Hybrid data warehouse-scenario, was in line with the solution's initial role as a complementary Cloud data warehouse. Here, SAP Data Warehouse Cloud would act as an extension of SAP BW/4HANA or HANA for SQL data warehousing. With regard to the latter, it is possible to approach SQL modelling through a hybrid approach (enhancing models between the two systems through interacting SQL code) or integrated approach (merging HANA for SQL data warehousing and SAP Data Warehouse Cloud on a single virtual tenant). The most important question with this scenario is whether it carries enough merit for organisations to commit to the required TCO of running two data warehouses in parallel. As SAP Data Warehouse Cloud matured, SAP itself advised not to opt for a dual-data warehouse setup, but to instead consider utilizing the BW Bridge feature (encompassing one of the other scenarios) instead.

The second scenario involving SAP Data Warehouse Cloud, the Cloud-scenario, would entail using the solution as the single, main data warehouse in an IT landscape. While SAP Data Warehouse Cloud is currently not as feature-rich when compared to SAP's on-premise solutions such as SAP BW/4HANA or HANA for SQL data warehousing, there are cases that could benefit from a 'greenfield' Cloud data warehouse implementation. Smaller organizations that have not had a data warehouse in the past, or Cloud-native firms that are looking for a new way to govern their increasing data volumes for example. Smaller SAP customers with growing businesses and no established, mature business intelligence infrastructure could also consider SAP Data Warehouse Cloud as a next step.

Finally, with the launch of BW Bridge, the third scenario for SAP Data Warehouse Cloud can be considered; the BW Bridge-scenario. This scenario is only viable for customers with an existing SAP BW(/4HANA), arguably the biggest part of SAP's target audience for SAP Data Warehouse Cloud, with an opportunity to (gradually) make the transition to Cloud-based data warehousing, without immediately having to write-off their existing data warehousing investments in the Cloud environment. In the previous section we already discussed the added value of BW Bridge in closing the gap between an existing SAP BW(/4HANA) system and SAP Data Warehouse Cloud. BW Bridge essentially creates an ABAP-based engine (or stack) within SAP Data Warehouse Cloud, allowing it to utilize all existing SAP BW features within the new Cloud data warehouse. We believe that SAP's proclamation that states that SAP Data Warehouse Cloud will be BW/4HANA's successor, makes this scenario the most important one to consider.

Pricing and licensing

On paper, SAP Data Warehouse Cloud's scalability is one of its main selling points. SAP provides a transparent, clear overview of capacity and its related cost for all components (Storage, RAM/memory, Data Lake storage and BW Bridge capacity). In practice, we still see combination limits between Storage, RAM and Data Lake and BW Bridge capacities. 4 TB (somewhat on the lower end for most DWH users) is the current storage limit, with an optional data lake storage for federated access.

BW Bridge is presented as part of SAP Data Warehouse Cloud, and is hence also billed in this fashion. The functionality itself is free of charge, but customers will have to license their data footprint (storage and computing capacity). While also available, data lake storage adds significantly to the license costs. An alternative to this would be to allow SAP Data Warehouse Cloud to access separate data lakes, built on commodity hardware for example, via Open SQL. For more details, custom scenarios or licensing options, please do not hesitate to contact us!

How can we help you?

Do you have a question about SAP Data Warehouse Cloud? Get in touch and we'll get back to you as soon as possible!

inspire@expertum.net

expertum.net