

SAP Business Data Cloud

‘Combines familiar functionalities we know’



expertum
Whitepaper



SAP Business Data Cloud

White Paper



Table of contents

Preface.....	3
What is SAP Business Data Cloud?	4
Solution components	4
Architecture.....	6
SAP Business Data Cloud Connect.....	8
SAP Snowflake	9
Guidelines and limitations for zero-copy data integration	9
SAP BW Private Cloud Edition (PCE).....	9
Pricing & Licensing.....	11
SAP Datasphere	12
Positioning within SAP's North Star Architecture	14
SAP's Strategic Acquisitions	15
The Way Forward	16
Practical Scenarios.....	16
Appendix: End of life dates SAP Netweaver solutions	18

*Rights and credit for most visuals and/or schematics in this document belong to SAP SE.

Preface

Ever since SAP announced SAP Analytics Cloud as their flagship solution for reporting and dashboarding almost a decade ago, their goal has been to create a seamless, semantically-rich data foundation underneath in order to optimize the flow of data from back-end to business. The idea was to create a transparent, governed infrastructure that could consistently translate data to actionable business information and support the enablement of ever more intelligent Artificial Intelligence (AI). With the publication of the AI-native North Star Architecture and the acquisition of several cornerstone software developers earlier this year, SAP is taking great strides towards realizing this vision.

In this updated and revised White Paper, we will explain how SAP Business Data Cloud, with SAP Datasphere at its core, offers customers who have other SAP Line of Business solutions in their landscape the most effective option for a data and analytics platform through strong SAP-to-SAP integration and the preservation of semantic context. In the fiercely competitive data and analytics space, SAP's greatest asset is its semantically rich enterprise data. We will dive into the solution's differentiating functionalities such as Data Products and Delta Sharing, both of which leverage aforementioned enterprise data, and will touch upon the various integration options and on how to sustainably make the shift from SAP Business Warehouse (BW) to this next-generation Cloud solution.

What is SAP Business Data Cloud?

In the most blatant terms, SAP Business Data Cloud (BDC) is a combination of SAP Datasphere, SAP Analytics Cloud and an SAP-specific version of Databricks, where each component can be separately licensed in accordance with a SaaS model. This means that for future SAP customers, it is more enticing to look at BDC as a potential cloud data warehousing solution as opposed to just Datasphere. Existing Datasphere customers can repurpose their tenant towards BDC via SAP's tenant rewiring process. For now, let's take a closer look at what Business Data Cloud actually is.




Figure 1 – SAP's vision on bringing applications, data and AI together

Solution components

Years ago, SAP stated its vision was to create a unified data model that would function as a foundation across its solutions. One important step towards this goal was to integrate their data warehousing, reporting and planning solutions in a single technology stack. With Business Data Cloud, SAP moves to realise this ambition and adds a modern AI component to it. The main design goals behind BDC are to:

- Uncover the hidden potential in an organization's business by unlocking seamless access to crucial insights;
- Boost confidence in an organization's data quality and integrity to empower data-driven decisions;
- Harness the power of fragmented, unstructured data sources and turn them into valuable business insights. This is done by combining SAP and non-SAP in a unified layer where business semantics are maintained, extending SAP's Business Data Fabric concept.

A key feature of SAP Business Data Cloud are the **Intelligent Applications** (formerly known as Insight Apps); out-of-the-box objects that create business insights through a combination of an SAC Story and a dedicated Datasphere Space. Intelligent Applications consist of multiple interconnected objects, artifacts, and tasks that are executed and created during their installation. From top to bottom, Intelligent Applications consist of SAC Stories (dashboards), data models, data products (more on those later) and data sources.

It should be noted that Intelligent Applications are *licensed separately* from the main BDC license; customers do not receive them as a form of business content. The Intelligent Applications that have been released up until today only contain data coming from either S/4HANA or SuccessFactors. Customers can also create their own Intelligent Applications (a bit more elaboration on that in the section below).




Figure 2 – How Intelligent Applications are built up

An important building block of Intelligent Applications are **Data Products**, arguably the core differentiating object within BDC. Data Products contain predefined table structures and interfaces that are directly available for consumption and include their respective business semantics. They contain data sources from SAP LoB solutions such as S/4HANA and SuccessFactors. The integration, cleansing and transformation of Data Products is managed by SAP in the background (on a Spark Engine). These Data Products (of type 'Delta Share') can subsequently be shared via **Delta Sharing** towards (SAP) Databricks, Google BiqQuery, (SAP) Snowflake and in the near future, Microsoft Fabric (GA slated for Q1 2027). Delta Sharing uses the Open Resource Discovery protocol (ORD). This allows for a zero-copy approach (no movement of data) that is centrally governable and highly scalable (see image below).

Of course, there is no such thing as *true* zero copy; delta sharing works by copying Apache Parquet files over the network between BDC and the partnering solution. So, while it is true that no data is being replicated and no ETL-pipeline is required, data is being queried over the network (see the bidirectional arrow in the figure below). This approach is comparable to federated queries (Google BigQuery) or External Tables (Snowflake), for example.




Figure 3 – Delta Sharing schematically explained

Data Products come in either *standard* versions, delivered out of the box by SAP, or *custom* versions, products that can be created by customers or partners. To create Custom Data Products, SAP introduced the **Data Product Studio** earlier this year, which is also able to use non-SAP data. This separate application, which is part of the BDC Cockpit, allows customers to also share Custom Data Products with SAP Databricks (through Delta Share). Similar to how Custom Data Products are built, customers can also create custom Intelligent Applications inside SAP Analytics Cloud (SAC) and Datasphere, respectively. Note that, for these Intelligent Applications, the Data Products you want to use will have to be installed in SAP Datasphere first and you cannot include Analytic Models in them (manually). It will be interesting to see if Enterprise Databricks (stand-alone) and Google BigQuery as both mentioned above, will also be able to fully leverage Custom Data Products as created via the Data Product Studio

Now, what about **custom fields** in the source object, for example, S/4HANA CDS views? SAP has stated that by default, Data Products will *not* include any specific 'Z-fields'. However, if a customer onboards a (Custom) Data Product through BDC's Foundation Services, the metadata for the source views is analyzed and subsequently harmonized with the Data Product in SAP Datasphere, enabling it to retrieve the contents of custom fields as of that moment. So long story short; if your custom fields are part of the source object's metadata definition at the time of the respective Data Product's onboarding, you will be able to leverage these custom fields. Note that only *installed* Data Products (and thus, replicated data!) support this feature; custom fields are *not* carried over to Intelligent Applications that are delivered as SAP-managed content.

Architecture

In this section we will briefly discuss the architecture of SAP Business Data Cloud. As the image below illustrates, source systems can either be SAP (for example SAP S/4HANA or SuccessFactors) or non-SAP. The Foundation Services in BDC are used to replicate data from SAP source systems (either SAP Cloud or on-premise solutions or non-SAP solutions) and bundle it with other business application data, transforming it into Data Products.




Figure 4 – SAP Business Data Cloud architecture

SAP Datasphere is at the core of BDC, acting as the main data structure layer on top of the aforementioned Data Products and as the building area for custom Data Products. Apart from its main role as a data warehouse, Datasphere in this context can also be used to extend both Intelligent Applications and Data Products with AI applications. Finally, it also acts as a central hub to manage data access controls and roles. **SAP BW** also plays a role in the form of either BW 7.5 Private Cloud Edition (supported until 2030) or BW/4HANA Private Cloud Edition, where BW objects from these specific solutions can be onboarded as Data Products in BDC.

SAP Databricks offers machine learning, AI and Pro-Code functionalities managed by Databrick's Unity Catalog. In practice, this will mean that you can enrich Data Products with these additions and push them back to BDC's catalog, exposing them for consumption by connected systems. Having the option to share delta directly from Delta Tables without replication offers a great way to integrate SAP datasets with Databricks. This also means that the **Premium Outbound** functionality (in SAP Datasphere) via replication flows is no longer required if you work with SAP Databricks (and for the stand-alone Databricks only through federated flows). Through these options, customers can now employ data scientists and data engineers in new ways within a full-SAP data environment. Do note that this SAP-tailored version of Databricks does not include the full Databricks architecture: it must also be activated as a separate component within BDC and is a compute engine only; one that does not come with any specific data storage.

SAP Databricks will have to be onboarded in the Business Data Cloud formation inside the customer's tenant (via the SAP For Me-account), before it can be used. After this has been done, Data Products can be shared to SAP Databricks from BDC's Foundation Services. In the case of Custom Data Products; these can be build based on local data (which currently *has to come* from the HDLFS Space in Datasphere) and can then be shared with SAP Databricks via Delta Sharing. Finally, customers can then use either SQL or Python Notebooks to work with the data in SAP Databricks.

When it comes to **S/4HANA Private Cloud**, customers should note that the SAP Cloud Connector will be required to integrate it with BDC. Moreover, there are limitations to the ABAP integrations between the two systems. In BDC's current (2026.19) release S/4HANA Private Cloud, S/4HANA Public Cloud and SAP SuccessFactors are supported within a BDC Formation. More information can be found here: <https://me.sap.com/notes/3500131>.

If we look at the front-end of BDC, it consists of **SAP Analytics Cloud** (both the 'regular' functionalities as well as Planning) and the new Intelligent Applications. The latter have been briefly touched upon in the previous section, but it is important to note that they are delivered as SaaS objects that can focus on different scenarios such as prediction or forecasting, depending on the use case at hand.

Finally, customers who do not yet have S/4HANA and are still on **SAP ECC** but *are* contemplating a connection to BDC/Datasphere should know that this can be realized through replication flows. Note that if data needs to be loaded into Datasphere's Object Store (in HDLFS-format), customers should use the ABAP adapter over the classic ECC adapter. However, our advice would always be to consider when (not *if*) you will migrate your legacy ERP system to S/4HANA in order to guarantee the prospect of a healthy future data landscape.




Figure 5 – A more comprehensive overview of SAP Business Data Cloud

It is our expectation that customers will be able to license the separate solution blocks of BDC as they see fit for their landscape. Note that, while the BDC cockpit can manage multiple tenants of different solutions, each BDC formation can only ‘host’ a single DSP and SAC tenant.

SAP Business Data Cloud Connect

Last year, SAP and Google announced the **SAP Business Data Cloud Connect for Google BigQuery**. This marked a major step in SAP’s strategy, given that the zero-copy mechanism implemented here was first announced in conjunction with BDC itself, but at the time exclusively with Databricks. The BDC Connect for Google (in which Google Cloud was also the first hyperscaler to support BDC Connect) is bi-directional, meaning that customers can carry over semantically rich data from their SAP ecosystem to and from BigQuery. SAP and Google are quick to mention the use cases for AI and autonomous agents here, but I am a little skeptical on the efficiency of zero-copy (or virtualization/federation) when it comes to the volumes of data required for these applications. Nevertheless, this is good news for customers who (also) have Google solutions as part of their data landscape, as long as we keep in mind that BDC Connect *only* works through the delta sharing technique and that there is no physical data replication option.

Similar to BDC Connect for Google, Databricks also presented the general availability of the **SAP Business Data Cloud Connect for Databricks**. Offering the same, bi-directional zero-copy sharing of data(products) (specifically via Delta Sharing in the case of Databricks), the BDC Connect for Databricks is especially useful for customers that are considering BDC (or Datasphere), but already have a separate Databricks instance in their landscape. For these customers, BDC Connect offers a relatively simple and secure (mTLS and OAuth) method to get SAP business data to Databricks without losing SAP’s business semantics. One question we still have here is whether BDC Connect works across hyperscalers, or if it has the same limitation to the same Cloud similar to the SAP BDC connector for Databricks (the latter of which is currently only available for AWS and GCP) in terms of either availability and/or concurrent network costs.

With regards to BDC Connect towards Databricks, we also have current concerns on the (probable) lack of storage optimization of data on HDLFS (Apache Parquet files) when using Delta Sharing. In stand-alone Databricks, this process can be optimized (e.g. with liquid clustering), but it is unclear whether or not this will work with the data stored in the BDC Object Store, as it is read-only. This in turn could cause performance issues with large data volumes that should (ideally) be partitioned. We will report on our future findings at Expertum and check back in here when there is more clarity on these topics.

Finally, BDC Connect is currently also available for bi-directional data sharing with **Snowflake** (Enterprise) and is slated to be released for **Microsoft Fabric** (OneLake) in Q1 of 2027 (GA).

SAP Snowflake

SAP Snowflake gets its own sub-header because it is not technically part of BDC's architecture; only in terms of procurement and packaging. In essence, it simply *is* Snowflake Enterprise that a customer licenses via SAP. The customer gets the full Snowflake experience; this includes access to Snowflake's proprietary Horizon Catalog as well as Apache Polaris, and of course support for the Apache Iceberg table format. SAP Snowflake comes packaged with the bidirectional, zero-copy integration with BDC I just mentioned. In practice, this means that customer data enrichment and AI models (on Snowflake) can be performed/maintained closer to where their (SAP) data resides. The goal here is for customers to be able to access SAP's business semantics, lineage and governance directly in their SAP Snowflake environment. All commercial links (selling, billing and support) is done through SAP itself. Hence, SAP Snowflake is also managed and provisioned via the customer's SAP For Me environment.

Guidelines and limitations for zero-copy data integration

Of course, the whole delta-sharing principle from BDC towards solutions like Databricks, Snowflake, BigQuery and Microsoft Fabric comes with some nuances. First of all, SAP states that SAP BDC Connect may not be used for data replication as stated in the SAP Business Data Cloud Supplement. More specific; it states that 'Third Party Integrations' are only permitted to *temporarily* store Data Products for *performance optimization* purposes, so not for *distributing* Data Products to subsequent systems (e.g. MS Power BI, Tableau or an external Data Lake). However, once data is *enriched* (e.g. through joins, unions or AI processing) this data is no longer considered an SAP BDC Data Product and thus is no longer subject to the rules laid out in the BDC Supplement. This 'resultset' may then be stored within a target partner system and used as a native data object. For more information, always ensure that you consult the required papers (contracts, supplements) yourself before taking actions with your SAP data that are not clearly described in the standard BDC documentation.

SAP BW Private Cloud Edition (PCE)

With the arrival of BDC, SAP (for all intents and purposes) discontinued SAP BW Bridge as a solution going forward. Across the board, SAP BW customers are stimulated to upgrade/update their solutions to *at minimum* SAP BW 7.5 on HANA. Customers that currently have SAP BW 7.5 on HANA or SAP BW/4 HANA (both on-premise options) can migrate to their respective PCE counterparts. Note that 'BW PCE' is sometimes also referred to as 'SAP RISE BW', 'BW BDC', 'HEC BW' or 'BW Tailored'. Additionally, an even more tailored option is the Private Cloud Tailored Option (PTO). This path allows customers to go outside the T-shirt sizing of PCE and request specific configurations such as CPU, memory, disk, recovery and others. Unsurprisingly, this is also the most costly way forward.

Both BW 7.5 and BW/4HANA PCE cover 100% of their respective on-premise variants functionalities (within their release/SP-version). The **Data Product Generator (DPG)** is a BDC tool that derives data products from SAP BW(/4) PCE InfoProviders into the HANA object store of SAP Datasphere (in BDC). It has to be paid for through BDC Capacity Units. Customers should also note that the Data Product Generator is available in PCE deployments in BDC only, but that the objects that the customer currently uses in SAP BW(/4) on-premise will also be available in BW(/4) PCE.




Figure 6 – An architectural overview of the workings of the Data Product Generator

Next to the DPG, SAP also introduced the **Query Template Generator (QTG)**. Queries have long been a challenge for customers who were looking to migrate away from SAP BW; with the QTG, they are now able to convert BW Queries into Datasphere-native Analytic Models. The query's metadata is translated into the Analytic Model, including facts and reusable dimensions, whereas the underlying InfoProvider can be migrated via the DPG.




Figure 7 – An architectural overview of the workings of the Query Template Generator

The migration to BW PCE requires an SAP HANA database (meaning SAP BW 7.5 will need to be on HANA in order to qualify for PCE). As of Q4 2025, BW legacy objects (7x) are also supported in BW PCE, meaning there will be no need to convert objects such as InfoCubes or classic DSOs to their new, BW/4 counterpart. However, a note that we should place here is that we currently have not yet seen (at a customer) to *what degree* legacy BW objects are supported (are they read-only or can they be Edited and/or Created in BW PCE?).

What *is* a requirement however, is that SAP requires customers to migrate to a Hyperscaler that is managed by SAP (a 'Lift & Shift'). This activity will most likely have to be supported by SAP or an SAP partner. SAP states that the database or the ABAP application stack underneath SAP BW are not affected by this Hyperscaler migration, but basic testing and sanity checks will need to be performed after the fact. Finally, the **expected throughput time** for a single BW tenant (e.g. a D or T system) to be migrated towards PCE from SAP's side is around a month.

Pricing & Licensing

SAP Business Data Cloud's pricing structure can be considered somewhat complex, as it functions through BDC credits (as opposed to BTP, or now BAIP, credits) and a separate configurator. For example, SAP Managed Data Products will require an additionally monthly fee on top of the base BDC license (and scale based on Full User Equivalents (FUEs) for S/4HANA sources). Aside from that, customers who want to utilize Datasphere's Object Store will need a baseline setup of at least 4150 Capacity Units (CUs), which means that even smaller environments that want to use functionalities such as Delta Sharing will need to reach this threshold in terms of size. For a first impression, [this capacity unit estimator](#) can be used. If you are interested in more tailored advice, please do not hesitate to [contact us](#).

SAP Datasphere

Branded SAP Data Warehouse Cloud at its inception, **SAP Datasphere (DSP)** was initially 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. DSP acts as the core of BDC where data can be replicated, transformed and shared towards (external) targets. The solution is positioned as a Data Warehouse as a Service (DWaaS) in the Public Cloud supported through short release cycles, with a focus on reducing time-to-value.




Figure 8 – An overview of SAP's data warehousing strategy

SAP Datasphere is Walldorf's strategic successor of SAP BW/4HANA. This does *not* mean however, that SAP Datasphere is *replacing* SAP BW/4HANA (see the appendix for specific end-of-life dates). SAP BW 7.5 will remain supported until 2027 in terms of mainstream maintenance (2030 for extended maintenance), whereas for SAP BW/4HANA this support extends to 2040. SAP BW/4HANA has conformed to an annual release cycle since 2022, with SAP publishing a PDF to document the updates and changes made to the solution. 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 Datasphere 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 see SAP Datasphere as SAP's answer to competitors such as Snowflake, Google's BigQuery, Databricks (despite being a partner of SAP) and Amazon's Redshift, in addition to it being a BW/4 successor.

While conceiving SAP Datasphere 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), addressed by creating a collaborative environment via the Space-concept;*
- *Lengthy time to value, addressed by being end-to-end integrated;*
- *Lack of real-time data, addressed by being fast and responsive;*
- *High entry costs (for dedicated BI/analytics solutions), addressed by being elastic in capacity.*

SAP thus designed SAP Datasphere 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 9 – The concept of Spaces in Datasphere visualized




Figure 10 – SAP Datasphere as a Business Data Fabric ecosystem

Positioning within SAP's North Star Architecture

Traditionally, the focus for data and analytics was on consolidating data streams, optimizing storage formats, and ensuring robust reporting capabilities for executive decision-making. SAP operated under the banner of the Intelligent Enterprise, which focused heavily on the automation of highly repetitive tasks and basic predictive analytics. However, in the new reality of the **Autonomous Enterprise** (as presented by SAP during Sapphire 2026), merely storing data and generating retrospective reports is no longer sufficient. Businesses today demand systems that do not just execute predefined, hardcoded tasks, but are inherently capable of reasoning, learning, and adapting to new variables in real-time.

To bridge this gap, SAP introduced its **AI-Native North Star Architecture**. This blueprint is not just a high-level product roadmap; it is a fundamental reimagining of SAP's entire enterprise software foundation. In the context of Walldorf's Data and Analytics portfolio, the North Star Architecture is so crucial because it sets the strategic direction for an ecosystem where AI agents, complex business processes, contextual intelligence, and strict semantic governance converge. Without a unified architectural vision, integrating advanced workflows and AI into deeply entrenched, highly complex legacy ERP environments would inevitably result in fragmented data silos and risky AI hallucinations.

At its core, the North Star Architecture brings multiple distinct layers together: a centralized **platform** layer, a robust data **foundation**, an intelligent **process** layer, and a unified **user experience** driven by autonomous agents (e.g. Joule). The North Star Architecture serves as the guiding framework that aligns every component of SAP's solution portfolio, ensuring that enterprise data is no longer treated as a static asset, but rather as a fuel that powers an ever-changing organizational (IT) infrastructure. By systematically decoupling storage and compute while enforcing strong semantic governance, this architecture prepares SAP solutions for a future where generative AI and intelligent agents are embedded natively into the (business) fabric of everyday operations.




Figure 11 – The four layers of SAP's AI-native North Star Architecture

SAP's primary goal here is to create an **overarching semantic map of a business** through the **SAP Knowledge Graph**. This graph encodes nearly fifty years of SAP ERP engineering, business semantics, and custom code (Z...) into machine-readable relationships, acting as the critical bridge that allows AI models to understand the intricacies of enterprise data without losing business context. Joule Agents also rely on the Knowledge Graph to ground their actions, ensuring no relevant semantics are omitted during processing. Having said that, building an AI-native ecosystem requires more than just clever tools; IT landscapes will have to overcome the 'data tax' associated with constantly moving, replicating, and transforming data. I see this reflected in ever-growing discussions at customers on where (SAP) data is stored, where it is processed and where it is to be consumed.

SAP's Strategic Acquisitions

In the above image illustrating SAP's North Star Architecture layout (figure 8), I have plotted the position of their three acquisitions in the Foundation Layer; right in BDC's ballpark. But what do they actually do?

To ensure that the data foundation underlying the North Star Architecture is properly structured, SAP acquired **Reltio** earlier this year (in May). Reltio's advanced master data management (MDM) capabilities bring Model Context Protocol (MCP) support for event-driven data sharing directly into the SAP Business Data Cloud, essentially creating a powerful AI-supported metadata catalog (sounds familiar...?). This ensures that the foundational entities AI relies upon are clean, harmonized, and form a single source of truth across the landscape. SAP's own Master Data Governance (MDG) suite will also be able to synchronize with Reltio to enforce the 'golden system of record' across the landscape.

Following this, SAP aimed to address the persistent headache of data silos by acquiring **Dremio** (in June), a move that effectively tackles the overhead of ETL processes. Dremio transforms BDC into an **Apache Iceberg**-native enterprise lakehouse and, coupled with the open data catalog Apache Polaris, this acquisition allows SAP and non-SAP data to be queried seamlessly via open standards without costly and historically required data movement. People who have already worked with BDC's **Delta Share functionality** will notice that this is a step towards more agnostic data integration: where Delta Share is tailor-made for Databricks' Unity Catalog and Spark integration, Iceberg allows customers to query and federate data via REST catalogs across multiple engines, such as Flink or Trino. BDC will thus offer both Delta Lake *and* Apache Iceberg table formats going into the future, allowing customers to leave their data in BDC and federate it to wherever it needs to be (either via Delta Share or Iceberg).

While Reltio and Dremio solve the distinct challenges of data governance and physical location, SAP's €1 billion commitment to **Prior Labs** (whose acquisition was completed last month), specifically addresses the Model Services in the Foundation Layer. Standard Large Language Models (LLMs) often struggle with the mathematical precision of relational records; but businesses primarily run on highly structured tables, not just text. Prior Labs offers the solution in that it specializes in **Tabular Foundation Models** (TFMs) (e.g. TabPFN-2.6). These statistical reasoning engines learn directly from tabular data, offering lightning-fast zero-shot predictions and in-context learning. By embedding TFMs alongside its own Relational Pretrained Transformer (RPT-1.5) technology into the North Star Architecture, SAP ensures that its autonomous agents natively comprehend the statistical causality within structured business metrics. Hence, put together, Reltio, Dremio, and Prior Labs enable an important part of the North Star Architecture when it comes to data: harmonizing it, federating its access, and mathematically reasoning over its content.

For more information about the context and workings of these new technologies (in the context of AI), please refer to [this blogpost](#) or [contact us](#).

The Way Forward

So, what about customers who are already invested in the SAP data and analytics ecosystem? Existing customers that use SAP Datasphere or SAP Analytics Cloud will be able to add the Business Data Cloud Cockpit as an additional service within their BTP environment. They can create a formation which includes all three systems (SAC, Datasphere and BDC), provided that they use the same tenant and data center. Linking the systems and jumping from one to another offers a seamless, integrated experience to users. This means SAP Datasphere's native capabilities remain, with pre-built analytical applications based on SAP data models and the additional option to extend these via the SAP Databricks component. For example, by mixing SAP with non-SAP data in a 'Derived' Data Product, you can store this Data Product on Datasphere's Object Store and make it available to SAP Databricks through Delta Share, bypassing premium outbound functionalities. Combine this knowledge with the integration with SAP Analytics Cloud, which includes dashboarding and (financial) planning, and we can safely say that BDC will indeed be the next step in the integration between SAP Datasphere and SAP Analytics Cloud as a **unified solution**.

Practical Scenarios

SAP's Business Data Cloud combines familiar functionalities we know from Datasphere and SAC with exciting new capabilities from Databricks and integration options with various non-SAP solutions through Delta Share. Now the question remains: how do you get there? Let's say that you are a potential Business Data Cloud (or Datasphere) user who needs to make the step from SAP BW. You see the value in the combination of semantically rich Data Products combined with Delta Share and strong SAP integration. What are your options and what do they look like? In the overview below we have shared a high-level visual on what paths are possible from BW to the next of generation (SAP) data warehousing.

Customers who still have BW 7.5 will have to make a choice. The simplest one is to go into (paid) extended maintenance until 2030, but that is not an advice we would be eager to give. The gains are minimal, and one is basically postponing migration activities that will *have* to be undertaken unless a move away from an SAP data warehouse is made altogether. Therefore the viable choices are either to move to SAP BW/4HANA to remain on-premise, gradually move to the cloud (BDC with Datasphere) through BW PCE or go with a big bang towards BDC with Datasphere, the latter of which is also something that is quite challenging to do depending on one's legacy.




Figure 12 – The migration paths from SAP BW onwards

Then there are **customers that already have SAP Datasphere**. This group had the option to extend their Datasphere-only contract for a single year in 2025; any extensions afterwards would automatically move to a Business Data Cloud model (and all the credit shifts that come with such a contract). The consideration on why and when to move from SAP Datasphere to Business Data Cloud can have several factors. For example, the data center/hyperscaler that a linked SAP Analytics Cloud or S/4HANA tenant is running on can be important, as can the existing landscape (legacy) already built in Datasphere if you are for example using BW Bridge. The idea is that the move to Business Data Cloud ideally brings added functional benefits for customers, such as productive use of Data Products or Delta Share, not just a new contract.

In short: the best way to move forward is to either commit to long-term on-premise through SAP BW/4HANA, or to (gradually) make the move to the Cloud by taking either the BW PCE – BDC, or the incremental ‘Shift’ from BW-route. Our take is that SAP BW 7.5 on HANA PCE is not a smart/sustainable solution, but a temporary situation that buys customers time to ponder their future data landscape going forward. In SAP terms this would be either SAP Datasphere (as part of BDC), SAP BW/4HANA or SAP BW/4HANA PCE (in BDC). For customers that already have Datasphere and have built quite a mature data warehouse already, the step is smaller, but should still be considered carefully because of the throughput time on the infrastructure side of things (for either the PCE migration or the tenant rewiring process).

If you want tailored advice and quality in-depth support to continue your organization’s data and analytics journey, do not hesitate to [contact Expertum](#). Alternatively, feel free to read our [other blogposts](#) on SAP Business Data Cloud on our website.

Appendix: End of life dates SAP Netweaver solutions

Solution	Sunset mainstream support	Sunset extended maintenance
SAP NetWeaver BW 7.4 or older (on-premise)	End of 2020	<i>Customer-specific agreements</i>
SAP NetWeaver BW 7.5 (on HANA) (on-premise)	End of 2027	End of 2030 Additional Fee
SAP NetWeaver BW 7.5 (PCE)	End of 2027	End of 2030 No Additional Fee
SAP BW/4HANA 2021	End of 2027	<i>Upgrade to release '2023'</i>
SAP BW/4HANA 2023	End of 2030	<i>Upgrade to release '2026'</i>
SAP BW/4HANA (PCE)	End of 2040	<i>To be determined</i>
SAP BPC on NetWeaver (10.1)	End of 2027	End of 2030 Additional Fee
SAP BPC 2021 (for BW/4HANA)	End of 2030	<i>To be determined</i>