

# How Microsoft Is Stitching Together SAP Business Data Cloud, Fabric, and Databricks to Unify Data



**Robert Hernandez**  
Principal Product Manager  
Microsoft

Robert Hernandez spent nearly 20 years at SAP, rising to Vice President of Innovation Solutions, before joining Microsoft as a Principal Product Manager in 2021. His current role sits on both sides of the SAP-Microsoft relationship: he oversees a team of product managers in Walldorf working with SAP on infrastructure design for RISE with SAP, Business Technology Platform (BTP), and SuccessFactors on Azure, while coordinating co-development of offerings like Analysis for Office, Joule Copilot, and cross-platform single sign-on.

In a conversation with ASUG, Hernandez laid out how SAP Business Data Cloud (BDC), [Microsoft Fabric](#), and the Databricks partnership fit together for customers whose data extends well beyond the SAP estate, and made a case that identity management, not security architecture, is the foundation most organizations still haven't laid.

*This interview has been edited and condensed for length and clarity.*

## **Q: From your perspective, what is SAP BDC, and what problem is it solving?**

BDC is SAP's go-forward solution for analytics across the SAP stack. SAP will use BDC to unify information across its Suite applications to enable end-to-end analytics. Within BDC, SAP will deliver prebuilt "data products" for analyzing business objects from its cloud suite. This enables unified views of entities that span applications—such as a customer across Commerce Cloud and S/4HANA, or a vendor across Ariba and S/4HANA—by bringing those definitions into a consistent data foundation.

Many of the components within BDC are capabilities customers may already have in place today, such as SAP

## **Key takeaways:**

SAP Business Data Cloud (BDC) aims to unify analytics across SAP's suite via prebuilt "data products." When data and processes span non-SAP systems, Microsoft positions Fabric and OneLake to help unify and analyze data across sources and clouds. Identity and single sign-on across systems are often the main prerequisites that customers underestimate—especially for AI scenarios. Databricks fits in via tight SAP partnerships (e.g., Delta Sharing) and Microsoft's first-party-like Azure Databricks experience, with deployment choice depending on required features such as ETL for non-SAP data.

Datasphere, SAP Analytics Cloud (SAC), or SAP Business Warehouse (BW). With most versions of BW reaching end of maintenance in 2027, customers can instead deploy BW into BDC to extend its supported lifespan while planning a transition to newer capabilities.

## **Q: BDC unifies analytics across the SAP suite. But what about customers whose data extends well beyond SAP?**

Data unification across different applications is a challenge for many customers—especially as we talk about AI.

That can only work if you have reliable data across the enterprise. SAP solves that for SAP business processes with BDC, but what if a process spans multiple systems? Siloed, disconnected information creates a challenge for humans to interpret, let alone for AI.

I had a conversation with a customer earlier today, and he said, "Hey, Robert, 85% of my data is SAP data. I do not have a lot of data outside SAP in my organization." Well, I said, "Hey, you're going to be one of those customers who wants to take that 15% and put it into BDC, because you've got all your SAP data in there already."

If, however, you are like a lot of customers I talk to, who say, “Hey, my data is 50/50 or 60/40 SAP,” now we have to talk about how to bring together SAP with the rest of the data. Because frankly, my data gravity is not in SAP, but I need to bring it together with the other pieces. And I really want to be able to consume it with some of these tools. How do I do some of that?

**Siloed, disconnected information creates a challenge for humans to interpret, let alone for AI.**



**Q: For those customers who need to go beyond BDC, what is Microsoft bringing to the table?**

We brought Power BI into Fabric because people found that with Power BI, it's easy to analyze, gather, and pull data together. In addition to that, [we created a capability called OneLake](#), which makes it a lot easier to unify data from multiple sources and multiple locations. I can bring together information from different locations and various sources and consume it within Fabric—within Power BI, but then also within the Microsoft data stack for ML, forecasting, all those capabilities.

With OneLake, not only do I have access to Microsoft data—data that exists in Azure—but I can also consume data that may exist in AWS or Google Cloud Platform (GCP). We can connect to that data in GCP as if it were in Microsoft, without having to pull that data across. It really becomes this nexus point where I can bring together information across my enterprise.

**Q: What is the piece of this puzzle that customers tend to overlook?**

People underestimate this component: the first piece you must think about is identity. Even before I start talking about security, I've got to talk about identity, because while those two seem similar, they're actually two different concepts that you need to bring together.

I can go out there and create security, roles, authorizations, and permissions. But if I do not know who you are, it does not do me any good to create that.

This has tripped up so many of our customers who have tried to deploy AI as well. They say, “Hey, Robert's logged into Copilot or Teams, and he's trying to connect over to Joule, or he's trying to connect over to the back-end source system,” and when it passes my ID from Copilot to whatever, if that thing can't associate that user ID across both, you're going to stop right there, and the process is going to break down. And it will not be a situation where I get access to things I should not; it is where I do not get any access.

We want to be in a situation where, as much as possible, we drive the authentication back into that source application. That is not always possible. I had a customer who was joining together SAP data and some IoT data. The problem is, to bring those things together, they had to transform them, because the definitional elements that came out of SAP were not consistent with their wellhead information that they were getting from their offshore platforms. In that case, we needed a different security model in Fabric because definitionally, I could not bring those things together at the level and element that SAP secured and authorized it.

The goal would be to push that authorization back down to the source application if I can. But even if I do, I need to think about: is my user ID the same across those two? Do I have single sign-on authentication? I am with Microsoft, so


take it with a grain of salt, but [do I have Entra ID](#)? Do I have that same ID, that same user across all those applications?

If I do, my life will be pretty easy. [If I do not have that same ID, that same user across all those applications, then it is a nightmare just to try to get the systems to talk to each other.](#)

That is one thing that you aren't thinking about right now, if you want to try and do one of these projects, and you should think about laying some of that groundwork before you do.

**People underestimate this component: the first piece you must think about is identity. Even before I start talking about security, I've got to talk about identity[...]**

Connect with Robert Hernandez on LinkedIn.



### **Q: Where does Databricks fit into this, and what does the Microsoft-Databricks partnership mean in practice?**

[When I deploy Databricks within Microsoft, Microsoft treats Databricks as a first-party product.](#) When you get the bill, it comes in as a Microsoft product. You would be hard-pressed, if you covered the name Databricks, to recognize that what you are deploying is not actually the Microsoft offering.

Databricks is important as we talk about integration with BDC because SAP launched a tight partnership with Databricks, including the concept of Delta Sharing. There are a couple of different ways to bring Databricks together with BDC. One option is SAP Databricks, which runs within your BDC subscription. If you do not have Databricks yet and want more advanced AI capabilities alongside SAP data in BDC, that is one route.

At Microsoft, since we have such a big Databricks footprint, we have been having more conversations with customers who already have Databricks up and running. They have the

standalone Databricks, or what we call native Databricks, and that has more functionality and capability than the SAP Databricks.

SAP has this functionality called Delta Sharing, where I can take data that is loaded into BDC. I do not copy the data over to Databricks. It remains within BDC. But via Delta Sharing, I can view that information that is stored in BDC, and I can use it as if it were resident within Databricks.

If you do not have Databricks today, you have to figure out which route to go. A lot of the extract/transform/load (ETL) to bring the data in is turned off in SAP Databricks. If you need those capabilities for the non-SAP data, you need to go native. SAP Databricks may be easier if all you need are the features and capabilities that come with it.

### **Q: Once you connect Databricks to BDC, can they understand the SAP business semantics, or are you just moving raw data?**

What they will do is read the data product definition out of the BDC system. When you start talking about semantic models, this is where it gets more nuanced. If you think about SAP, you want to do things like currency conversion, or unit-of-measure conversion, or you want to think about fiscal calendars versus calendar calendars, all those sorts of business semantics. If I want to consume that via Power BI, the question is not just whether those data products exist, but how can I consume them in a way that Power BI can understand those elements?

When you start thinking about how you are going to report on this data, it comes down to use cases. If I am reporting SAP data and talking about SAP financials, I may want to utilize BDC and SAC and the intelligent apps that come with SAP. Those are specialty-built tools, capabilities, and visualizations for the SAP community. If I am a customer who runs Oracle Financials and SAP Financials, and I need to see the unified business, then I am going to have to figure out a different approach.

**Q: What does Microsoft offer to the customer who is running SAP alongside Oracle, Salesforce, or Dynamics and needs a unified view?**

This is one of those places where Microsoft, with our experience across a lot of these applications, is looking at bringing capabilities together for customers running multiple applications. In Fabric, we have brought together five business functions. We have brought together the various enterprise applications so that you can unify those together: pre-built data models, pre-built data flows, pre-built data mapping, all consumable with Microsoft Copilot.

I can consume that data via Copilot in Teams, say things like, “What are my top five customers by revenue?” This is not to replace a dashboard or report, but if I have that data in Fabric, I can consume that via Teams, via Word, via PowerPoint, and bring that data together.

**Q: What are the options right now for getting SAP data into Fabric?**

You have a couple of options right now. One piece we have is called Mirroring for SAP. That’s in private preview, where we mirror data from your SAP ERP system into Fabric. We also have a list of certified third parties that can load that data into Fabric. And then we have Datasphere replication flows, which can be written into Microsoft as well. I hope to be able to add Delta Sharing to bring the data together. These approaches all copy data.

**Q: Getting the data into Fabric is one thing. How do you ensure Power BI can make sense of it?**

SAP has its way of defining data, and some of it is in the data model, and some of it is in business logic. If I want to have the outcome, but I need it in Power BI, I need to translate that SAP-ese into Power BI-ese, if you will. That is what that component is doing—taking that data from SAP but transforming it in a way that Power BI can consume, keeping the business semantics correct.

The same is true for Oracle and for Salesforce. That’s one of the pieces that we’ve invested in, because I have so many customers who run SAP and Oracle, or SAP and Salesforce,

or SAP and Dynamics, and what I don’t want is a situation where every single one of those customers has to reinvent the wheel.

The idea would be that if I am running Oracle Financials and SAP Financials, they come together into a single place. A singular model across. My goal is not to have an Oracle Finance, an SAP Finance, a Dynamics Finance; that does not help anybody. At the end of the day, it does not matter if I am 50% SAP or 50% Oracle or 60/40; I need to bring those together to a unified financial view across the enterprise.

**Q: When you are pulling financials from SAP and Oracle into a single model, how do you reconcile back to the source?**

Well, here is the thing: it all must go into the same model anyway, because who demands it? Wall Street. That is who I must deliver the unified financials to. It does not matter what system it started from. At the end of the day, if I cannot provide reconciled financials, my CFO goes to jail. It is not just about individual source systems. It is about the overall outcome of bringing it all together.



**Visit the Microsoft website.**





Microsoft (Nasdaq “MSFT” @microsoft) creates platforms and tools powered by AI to deliver innovative solutions that meet the evolving needs of our customers. The technology company is committed to making AI available broadly and doing so responsibly, with a mission to empower every person and every organization on the planet to achieve more.



ASUG is the world’s largest SAP user group. Originally founded by a group of visionary SAP customers in 1991, its mission is to help people and organizations get the most value from their investment in SAP technology. ASUG currently serves thousands of businesses via companywide memberships, connecting more than 130,000 professionals with networking and educational resources to help them master new challenges. Through in-person and virtual events, on-demand digital resources, and ongoing advocacy for its membership, ASUG helps SAP customers make more possible.