Snowflake and SAP: Powering Enterprise AI with Unified Data and Context



Table of Contents

Introduction 3
Why AI Is Changing the Way Enterprises Think About Data 4
Why SAP Data Is Critical and Historically Challenging to Activate 6
How SAP and Snowflake Work Together to Build a Unified Business Data Fabric 7
A Partnership Build On the Success Of Joint Customers 9
Five Ways to Transform Business Processes with SAP and Snowflake 13
Turning Business Data into AI Advantage 15
A Long-Term Foundation for Enterprise Data and AI 16

Introduction

As enterprise AI moves from experimentation to implementation, it's becoming increasingly clear that the future of software isn't programmatical, it's conversational. Instead of composing queries using arcane syntax and then generating reports to view the results, business users will ask sophisticated questions of their data using natural language and get back detailed answers in near real time. That in turn will enable faster, data-driven decision making, enabling leaders to respond faster to rapidly changing market conditions.

The emergence of agentic AI will streamline this process further by taking actions based on what that data is saying. In some cases, agents will be able to automate tasks that currently require technical or analytic expertise. Instead of expending their limited resources collecting and analyzing data, business users will be able to spend more time evaluating agent recommendations and focusing on strategy.

But conversational and agentic AI are only as good as the data and semantics they rely on, and enterprises continue to face challenges generating clean, semantically rich, well-governed data that's ready for analysis and automation. This is especially true when dealing with the most business-critical information in the enterprise: data stored in the SAP platforms used by more than 90% of the Fortune 500.

This ebook will discuss why traditional data architectures are no longer sufficient, and how building strong connections between SAP and Snowflake can help establish the enterprisewide data foundation necessary for doing business in the age of AI.

Why AI Is Changing the Way Enterprises Think About Data

With the emergence of LLMs, enterprises now have to reimagine their approach to the information they use to run their business. Data is now expected to power real-time decisions, autonomous actions, process automations and predictive intelligence — which means a data architecture built around batch movement and periodic reporting is no longer fit for purpose. Here are three reasons why:

1. Traditional enterprise architectures are breaking down

Legacy enterprise data architectures were designed for a world where data was replicated and kept up to date in predictable, scheduled batches among a manageable number of systems. Today's AI-driven environment is exposing the fundamental limitations of those designs.

Replication-heavy pipelines increase cost, unexpected latency and risk. When data has to move across multiple systems, every hop adds infrastructure cost, introduces delays and increases the risk of errors and vulnerabilities. What's more, business meaning erodes as data moves across platforms. As data travels through systems, the business rules, hierarchies and relationships that exist from the source are frequently lost or distorted.

Further difficulty comes when there are changes to SAP systems that create issues downstream. Because SAP systems sit at the center of so many critical business processes, even small changes — like modifying a field or business rule — can cascade through the pipeline and system that depends on that data, triggering expensive and time-consuming remediation work and adding avoidable business risks.



2. Isolated platforms have hit a natural limit

When integration needs were limited and stable, connecting enterprise systems one at a time made sense. But as data estates grow more complex, this approach creates a tangled web of brittle, one-off connections that are expensive to maintain and nearly impossible to scale.

Traditional architectures simply can't keep up with real-time AI needs. Data pipelines were designed to move data in scheduled batches. That may be acceptable when the final product is a quarterly report, but it's fundamentally incompatible with AI applications that require fresh, continuous data for driving real-time decisions.

Governance and lineage fracture across systems. When data passes through multiple platforms and transformation layers, the ability to track where it came from or how it changed degrades at every step — making it extremely difficult to meet compliance requirements, audit usage, or trust that the data is accurate and complete.

Context doesn't survive movement, causing AI pilots to fail. When data is extracted from SAP systems and moved into an AI environment, the relationships, hierarchies and business rules that make that data meaningful rarely survive the journey intact. AI models are left to work with technically correct but semantically hollow data that may produce unreliable outputs and stall AI adoption. The massive effort and expense required to rebuild these semantics causes unacceptable delays. Before an AI model can do useful work, someone has to reconstruct the business meaning that was left behind during data movement — a labor-intensive effort that can consume months of work before a single model is trained.

3. AI success depends on context, semantics and trust

The primary barrier to enterprise AI success is not the technical capability to build and run AI models, it's ensuring that those models are operating on meaningful, consistent and trustworthy data. Here's what a lot of enterprises are still getting wrong:

AI models need business meaning, not just raw data. A raw database field containing a random number requires context that tells the AI what that number represents — days of inventory, a customer tier, a contract value or something else. The richer and more accurate the business meaning attached to the data going in, the more reliable and useful the intelligence coming out.

Example: Abstract field names like MATNR, BUKRS or VBELN, or a numeric code like "01" or "A3," are common in SAP ERP. Combined with years of bespoke naming conventions and custom terminology, these fields unfortunately provide no inherent meaning unless the AI understands whether it refers to a product, company code, document number, status or another business concept.

Without shared semantics, AI results can be inaccurate and untrusted. When different systems define the same business concepts differently, AI models operating across those systems will produce inconsistent results. Trustworthy AI is built on shared semantics, where every system agrees on what a term means and how a metric is calculated.

Example: An SAP system may define "revenue" as invoiced or recognized revenue, while a CRM system like Salesforce may define revenue as booked deal value. An AI model combining both without understanding the difference may produce forecasts or recommendations that appear precise but are fundamentally incorrect.

Enterprise AI requires security, governance and explainability by design. In an enterprise context, an AI model that occasionally produces unexpected or inconsistent results can have serious legal, financial or reputational consequences. Security, governance and the ability to explain why a model reached a particular conclusion need to be architectural commitments built in from the start.

Example: If an AI model automatically approves vendor payments or invoices without proper controls or auditability, an error or manipulation could lead to fraudulent payments, financial loss or regulatory violations — with no clear explanation of how the decision was made.

Why SAP Data Is Critical and Historically Challenging to Activate

SAP systems sit at the core of most enterprises, capturing the financial, supply chain, procurement, workforce and manufacturing data that allows organizations to function. This makes it uniquely valuable as a foundation for AI.

However, because SAP data accumulates across numerous systems and is often organized in proprietary structures, it can be challenging to extract the data along with the business context — context that AI models need to be accurate, secure, well-governed and explainable.

It is costly and slow to rebuild semantics, and traditional architectures are inadequate for the demands of AI. The gap between what these architectures were designed to deliver and what modern AI applications actually need can lead to several problems, including increased cost and risk, the erosion of business meaning, broken pipelines, and the failure of AI pilots.

While data replication has made SAP data more accessible, it has also introduced its own set of problems. When SAP data is replicated across warehouse, data lake, analytics tools and AI environments, governance becomes a multidimensional problem that grows harder with every additional copy. Each duplicated dataset requires its own access controls, audit trail and policy enforcement — creating an ongoing operational burden that many enterprises significantly underestimate as they build out their AI programs.





How SAP and Snowflake Work Together to Build a Unified Business Data Fabric

The urgency to scale AI across core business functions has never been greater. Leaders are under pressure to deliver faster decisions, agentic automation and measurable results — all while maintaining governance, compliance and security controls. Because businesses run on SAP, their most mission-critical processes already contain rich business meaning.

Together, SAP and Snowflake unite this essential context with the scale and intelligence of the Snowflake AI Data Cloud, creating a foundation designed to provide trusted, governed support for intelligent, enterprisewide execution. With the introduction of SAP Snowflake, an SAP-certified solution extension for SAP Business Data Cloud (BDC), and SAP BDC Connect for Snowflake, mission-critical business data can be made available for analytics and AI while reducing the expense and risk involved with moving data or the months required to build dedicated data pipelines.

SAP and Snowflake enable seamless, bidirectional integration of data and metadata, keeping both the data and its underlying business meaning synchronized between each platform. This means organizations can combine their SAP data with data from other systems, without the expense and effort of rebuilding all that business context from scratch. This unlocks the true potential of SAP data to empower advanced analytics, AI/ML models and intelligent business applications, while also simplifying the data landscape and shifting costs to value-focused data usage instead of data movement.

The result is a unified business data fabric that offers many benefits, including:

- **A unified layer of semantically rich data products.** Data is available in a consistent, meaningful and governed form, with business meaning, lineage and access controls intact.
- **Data that is usable across multiple systems**. The same governed data products can serve a BI dashboard, a customer-facing application and an AI model without requiring separate pipelines or redundant transformation work or fragmentation.
- **A single source of truth for business users, engineers and AI systems.** Today, finance leaders, data engineers and AI systems pulling the same revenue figures will often get slightly different numbers — because each is drawing from a different pipeline with its own version of the data. A business data fabric fixes this by establishing a single authoritative source from which everyone and everything can draw.

A fully managed platform for data and enterprise AI

By combining SAP Business Data Cloud (BDC) with Snowflake, organizations can put their most critical business data to work — bringing together SAP's deep reservoir of operational data with Snowflake's powerful analytics, AI and data collaboration capabilities to build solutions that are ready for the demands of the modern enterprise.

Here are four benefits to enterprises leveraging this powerful combination:

Access the right data foundation to harmonize your data. A unified, semantically rich foundation enables analytics, AI and intelligent applications to deliver on their promise at enterprise scale. You can build a business data fabric by bringing your SAP and non-SAP data together in a single, AI-ready foundation that simplifies your data landscape, reduces data silos and preserves business meaning, with zero-copy access.

Build AI-powered intelligent applications. Organizations can move from experimenting with AI to actually deploying it — building intelligent applications to automate decisions, generate forecasts and surface new insights. This enables you to deliver adaptive experiences and build AI-powered applications that draw on near real-time SAP data.

Deploy applications faster and continuously learn from your data. Combine business-ready SAP data with third-party sources, using Snowpark and Snowflake Cortex AI to quickly build predictions, forecasts and intelligent insights without having to rebuild the data foundation every time. Continuously learn from your data by optimizing your intelligent applications over time, using a single connected service that learns from the metadata flowing across your business data fabric.

Ground AI in your enterprise context and leverage AI agents. By connecting your data and metadata between SAP BDC and Snowflake, business meaning becomes available in both directions — giving AI the ability to understand your data not just numerically but also its relation to everything else. Using SAP's Joule intelligent assistant and Snowflake Cortex allows you to deploy the right AI agent for any task — whether you're trying to automate routine decisions or power fully intelligent applications.

A Partnership Built On the Success Of Joint Customers

"AstraZeneca is constantly pushing the boundaries of science and is pioneering in life-changing medicines. Data and AI are central to achieving this aim, and our close collaboration with SAP and Snowflake turbocharges our ability to access, process and analyze real-time data."

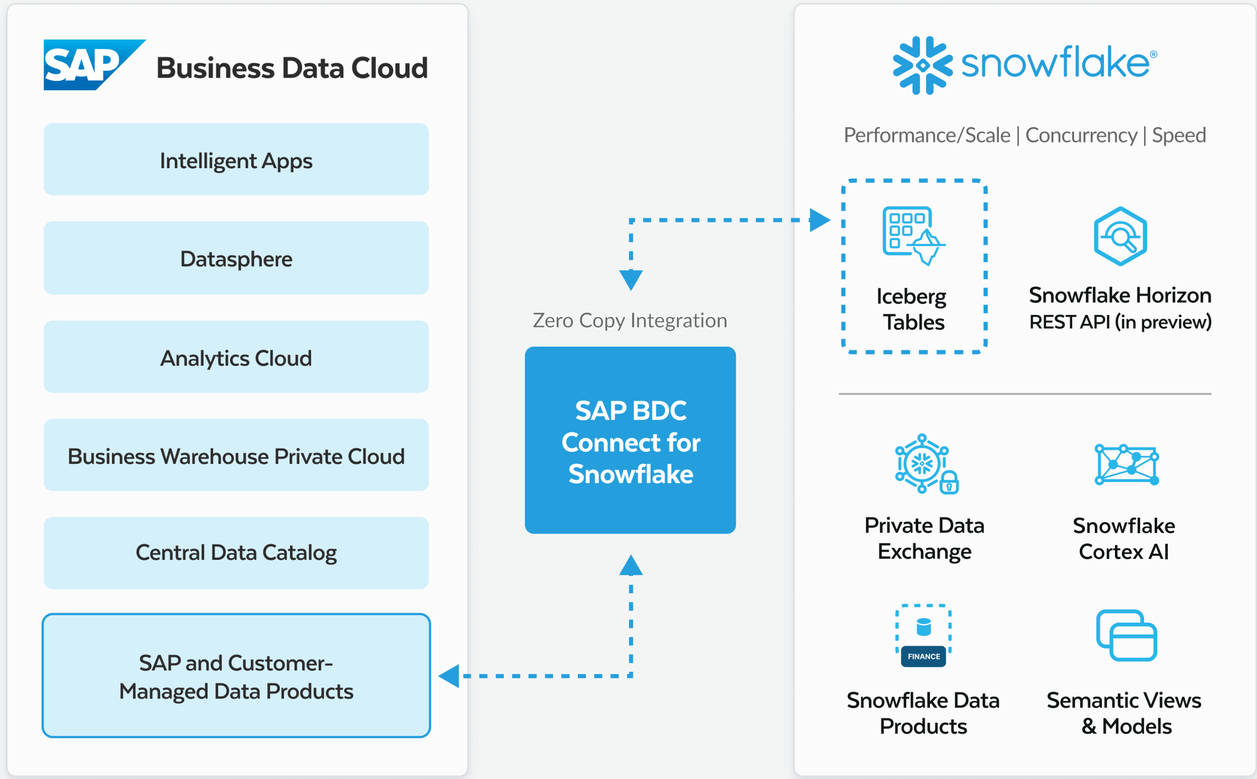
—RUSSELL SMITH
Vice President, ERP Transformation Technology, AstraZeneca

AstraZeneca

The breakthrough: zero-copy, bidirectional data sharing

In the past, extracting data from SAP platforms and getting it ready to use with AI was a costly and complex process. With SAP and Snowflake connected, there's no need to move or duplicate data. The original data remains intact, along with its context and business rules, with zero-copy, bidirectional data sharing. This preserves context so that semantically rich assets retain their business meaning, definitions and relationships and can be used and reused without needing to rebuild the underlying context.

With bidirectional data sharing, metadata, business semantics and lineage information are shared in both directions between SAP Business Data Cloud and Snowflake, and the two platforms remain synchronized. Enriched AI outputs generated in Snowflake can instantly be activated back into SAP BDC as new data products, with their meaning and lineage fully intact.



The SAP BDC Connector for Snowflake simplifies the bidirectional flow of semantically rich data.

Easy governance and reduced risk

This offering also makes governance easier and reduces risk by maintaining a consistent governance layer across both SAP and Snowflake. The SAP BDC admin governs which data products can be shared or un-shared with a Snowflake account. Once the data products are shared, customers can use RBAC in Snowflake to govern which users/roles have access to the shared data in Snowflake. When data exists in only one place, there is only one set of access controls to maintain, one audit trail to manage and one version of the truth to govern — helping reduce the compliance burden.

Greater customer choice and flexibility

The ability to keep data in sync while running workloads in either SAP or Snowflake offers organizations greater control and flexibility over their data estate. Organizations can choose which platform is best for their analytics, AI and applications, instead of having to standardize on a single platform. For example, SAP may be best for operational processes and Snowflake for advanced analytics and AI. They have choice and flexibility — without sacrificing data consistency or governance.

Zero-copy, bidirectional data sharing also supports hybrid strategies. Organizations that have built their data strategy around SAP can extend into Snowflake's AI and analytics capabilities without abandoning their existing investment, while Snowflake-centric organizations can take advantage of SAP's semantically rich operational data. The architecture is designed to meet organizations where they live, rather than force them to start over.

And because the data foundation is consistent, governed and reusable, the same underlying integration framework that enables a supply chain AI model today can also support a financial forecasting application or a customer intelligence platform, so you can deploy a single architecture for multiple outcomes.

How zero-copy changes the rules of the game

Embracing a zero-copy bidirectional model fundamentally changes how SAP data is used within an organization. Teams move faster, with greater efficiency, lower costs and reduced risk. From data teams to business users, productivity gains compound across the organization, including fewer pipelines, minimal duplication, less rework, re-use at scale and faster hand-offs.

When data products are shared in place rather than extracted on demand, there are fewer custom extracts or brittle integrations. A governed data product built once for a finance use case can be immediately reused by a supply chain team, an AI model and a customer analytics initiative — with no additional preparation work required.

Because data products carry their business meaning with them, teams no longer have to spend weeks reverse-engineering what fields mean or negotiating definitions with the team that produced them. Faster hand-offs are possible when business-ready data products are available on demand. The back-and-forth between data engineering and analytics teams compresses from weeks to days or hours.

Faster time to value

By eliminating the need to reinvent the wheel for each new project, teams are able to put their data to work more quickly.

- **Business-ready data from the start.** Rather than beginning every new initiative by wrangling raw SAP data into something usable, teams start from a governed and semantically enriched foundation that is already fit for analysis and AI.
- **Easier onboarding of AI projects.** When the data foundation is already in place and governed, standing up a new analytics or AI initiative becomes a matter of connecting to existing data products rather than rebuilding the underlying plumbing from scratch.
- **Repeatable use cases.** The same architecture that powers the first SAP data use case can also power the next 10. Each new initiative builds on the last rather than starting over, so the return on the original investment compounds with every addition.
- **Shorter iteration cycles.** When the data foundation is stable and reusable, teams can focus their iteration cycles on refining models and improving outputs rather than fixing the pipelines those models depend on.

Reduce risk and focus on what matters

Eliminating data movement and duplication greatly reduces an organization's exposure to data breaches or compliance violations. That means a lower risk of data loss, fewer points of failure and more consistent controls. And when data products have clear lineage and exist in a governed state, audits become smoother.

Zero-copy also reduces the overhead and risk of moving and managing data, so teams can focus more on the analytics and AI outcomes that matter. The ultimate promise of zero-copy integration is not just technical efficiency — it's the organizational freedom that comes from redirecting engineering talent and budget away from data plumbing and toward the analytics and AI work that actually moves the business forward.

How to achieve faster time to value, at lower cost and risk

Data teams working with SAP and Snowflake will also be able to shift their mindset from moving data to *using* data, from asking "How do we get this data out of SAP?" to "How do we show value from this data right away since we can use it right away?" This change in mindset will empower organizations to realize even more benefits:

- **Shift from investing in data movement to prioritizing business impact.** The less time your engineering teams spend on building and maintaining data pipelines, the more time they can spend building the analytics, AI and applications that actually deliver business value.
- **Evaluate total cost of ownership across the entire life cycle.** Traditionally getting data out of SAP was only the beginning of the cost — the real expense came in everything that followed: cleaning it, rebuilding semantics, governing it, maintaining it as systems evolved and doing it all over again when something broke. Adopting the new approach makes most of those tasks unnecessary.
- **Optimize for reuse and governed access, not one-off pipelines.** A data product you can build once and make available to every team is exponentially more valuable than a custom pipeline built for a single use case.

Five Ways to Transform Business Processes with SAP and Snowflake

Together, SAP and Snowflake give organizations a powerful foundation — combining the most mission-critical business data with powerful analytics and AI to drive transformation across every core business process. Leading retailers, manufacturers, healthcare, financial services firms and others already rely on SAP and Snowflake to drive value. Here are five concrete scenarios where the combination is delivering real business impact:

1. Transform supply chain operations from reactive to predictive

Snowflake can combine SAP's deep operational data — purchase orders, vendor lead times, inventory levels, procurement history and so on — with external data sources like weather feeds, geopolitical risk scores and shipping carrier performance data. The result is an AI model that can predict, as much as weeks in advance, that a critical component is at risk of shortage because a key supplier is in a region experiencing port delays.





2. Drive an intelligent and continuous financial planning workflow

Finance teams using SAP for their general ledger and financial reporting may find manual month-end close processes challenging. By activating SAP financial data in Snowflake in near real time and combining it with data from banking systems and rate feeds, companies can move from a backward-looking monthly close process to continuous accounting. Rather than waiting for a quarterly planning cycle, a CFO can see accurate, up-to-date views of financial position at any point, while forecasting models can automatically adjust revenue projections as new order data flows in from SAP.

3. Deliver proactive, personalized customer experiences

As transaction histories and service records are shared from SAP into Snowflake, they can be enriched with call center transcripts and third-party identity data to build a 360-degree view of each customer. AI models running in Snowflake can analyze that data to predict which customers are at risk, then share those predictions back into SAP in near real time — giving customer-facing teams the intelligence they need to step in at exactly the right moment.

4. Streamline workforce planning and boost retention

By combining SAP's HR data (headcount, roles, compensation, performance ratings, tenure) with external labor market data and internal tools usage, companies can build models that flag employees at risk of resigning. For example, by identifying when a specific engineering team's compensation has fallen below market rate, HR can intervene before losing critical talent.

5. Improve manufacturing quality and increase yields

SAP tracks production orders and quality inspection results, while IoT sensors continuously generate machine performance data — temperatures, pressures and cycle times. By bringing both together in Snowflake, manufacturers can determine how conditions on the factory floor correlate to product quality. For example, by identifying that machines running above a certain temperature consistently generate products that fail inspection, a food manufacturer can correct the problem before it reaches the customer.



Turning Business Data into AI Advantage

AI success depends on trusted, contextual business data. No matter how sophisticated the model or how powerful the platform, AI is only as valuable as the data that fuels it. For most large enterprises that means making SAP's rich operational data available in a form AI models can trust and understand.

SAP and Snowflake together deliver what enterprises need now. The integration between the two platforms delivers semantically rich data products and enterprise-ready AI capabilities using enterprise infrastructure that's already in place.

One architecture, flexible adoption, real AI outcomes. Whether an organization enters through SAP or through Snowflake, they are building on the same proven technical foundation — which means the path they choose today does not limit the outcomes they can achieve tomorrow.

Enterprise-ready AI by design. The hardest problems in enterprise AI — connecting disparate data sources, preserving business context and establishing governance — are solved at the platform level, rather than left to individual teams to figure out on their own.

Built-in security, governance and observability. Governance policies and full auditability of data and AI outputs are structural properties of the architecture — built in from the start, not layered on later.

Proven at scale. The SAP and Snowflake integration isn't a promise of future potential, it's an architecture joint customers across industries have already deployed and proven, putting SAP's operational data to work within Snowflake's analytics and AI platform to deliver measurable business outcomes.

A Long-Term Foundation for Enterprise Data and AI

SAP and Snowflake together solve the most fundamental challenge in enterprise AI — making the richest and most business-critical data in the organization available for analytics and AI, without stripping out the meaning, governance and trust that makes it valuable in the first place. The result is a foundation that reduces the cost and complexity of traditional data movement, and accelerates the path from AI pilot to production and gives organizations the flexibility to build on what they already have rather than starting over.

Organizations can choose the commercial and operational entry point that fits their current strategy — either through Snowflake or SAP — without making technical trade-offs that would limit their ability to expand or evolve the architecture as their needs change.

By preserving business meaning at the source, eliminating unnecessary data movement and providing a governed and reusable data product layer, the joint SAP and Snowflake integration gives organizations a foundation upon which they can build continuously — compressing the time from AI pilot to production and expanding their return on every data investment they have already made.

Ready to explore how this partnership fits into your data, AI, and SAP transformation plans? To take the next step, contact Snowflake sales or connect with your Snowflake team.



Snowflake is the platform for the AI era, making it easy for enterprises to innovate faster and get more value from data. More than 13,300 customers around the globe, including hundreds of the world's largest companies, use Snowflake's AI Data Cloud to build, use and share data, applications and AI. With Snowflake, data and AI are transformative for everyone.

Learn more at **snowflake.com**

(NYSE: SNOW)

© 2026 Snowflake Inc. All rights reserved. Snowflake, the Snowflake logo, and all other Snowflake product, feature and service names mentioned herein are registered trademarks or trademarks of Snowflake Inc. in the United States and other countries. All other brand names or logos mentioned or used herein are for identification purposes only and may be the trademarks of their respective holder(s). Snowflake may not be associated with, or be sponsored or endorsed by, any such holder(s).

ebook-snowflake-and-sap-powering-enterprise-ai-with-unified-data-and-context-r04