




Beyond Code Translation

A Semantic Approach to SAP BW Cloud Migrations

How a top-down methodology built on semantic consolidation transforms the most complex SAP BW migration challenge into a structured, accelerated, and trust-centered design exercise.



Overview

Table of Content

Executive Summary	3
1. The Urgent Need for Migration	4
2. The Limitations of a Bottom-Up Migration	5
3. The Advantage of a Top-Down, Semantic Approach	6
4. From Technical Objects to Valuable Business Assets	7
5. How a Semantic Layer Operationalizes the Migration	9
6. Beyond Migration: Building a Foundation for Trusted Data	12


Executive Summary

Most SAP BW cloud migration projects overrun budgets and timelines for a reason that has nothing to do with technology: organizations do not understand what they have before they try to move it. Decades of organic growth have left business logic scattered across thousands of objects and buried in custom ABAP code, poorly documented and difficult to understand.

Traditional bottom-up migrations compound the problem. They translate code object by object, carrying forward redundancy, technical debt, and inconsistent definitions into the cloud at full cost.

This paper argues for a fundamentally different approach. **A semantic, top-down methodology** builds a unified knowledge layer from metadata. It maps the entire legacy landscape, extracts business logic into platform-independent terms, and packages the results into governed data products before migration begins. This is the critical difference: organizations know exactly what to migrate, what to consolidate, and what to retire before they pay to move a single byte. The target system receives only clean, trusted, and business-aligned assets.

The result: dramatically compressed timelines, reduced cost, and a data foundation the business can rely on – for KPIs, reports, and AI-driven automation.



1. The Urgent Need for Migration


For organizations running on SAP Business Warehouse (SAP BW), time is running out. With SAP shifting its strategic focus to cloud-native platforms, SAP Business Data Cloud (SAP BDC), SAP Datasphere, and BW/4HANA, the message is clear: the legacy BW landscape has a limited future. Maintenance timelines are shortening, the talent pool that understands classic BW is shrinking, the cost of keeping aging systems alive continues to climb, and there is no room for innovation.

Yet for all the urgency, most SAP BW migration projects move slowly, overrun budgets, and often deliver less than promised. Most of the time, the reason is not a lack of tools or technology. SAP itself offers a comprehensive toolkit for migration. But the reason migrations stall is more fundamental: organizations do not fully understand what they have before they try to move it.

A typical enterprise SAP BW environment is the product of 10 to 20 years of organic growth. Thousands of InfoCubes, Data Store Objects, MultiProviders, and transformations have accumulated – many built by people who have long since left the organization, documented poorly or not at all. Critically, a large share of the business logic lives in custom ABAP code: start routines, end routines, expert routines, and field-level transformation rules that apply calculations, filters, and exceptions that only make sense to those who originally wrote them.

This ABAP code is at the center of the migration challenge. SAP's modern platforms perform best with SQL. Migrating therefore requires not just moving data, but re-engineering the logic itself: understanding what each ABAP routine does, what business question it serves, and then reimplementing it in SQL, or other languages. This is not a syntax translation, but a paradigm shift, and it is where projects burn through time and budget.

The picture becomes clear: the migration challenge is not primarily technical. It is a knowledge problem. What does the system do? Why does it do it? And does it still need to do it? The answers are scattered, implicit, and often simply lost.



2. The Limitations of a Bottom-Up Migration


The dominant approach to SAP BW migrations treats the challenge as a code translation problem. Tools and consulting methodologies work bottom-up. They inventory the technical objects in the legacy system, map them to equivalents in the target platform, and convert or rebuild them one by one.

SAP offers capable tools like Readiness Check [↗](#) or Custom Code Analyzer [↗](#) to cover each individual step of a migration. But they share a crucial limitation: none of them consolidates the real meaning of what is being migrated. They operate at the level of technical objects without a unifying layer that connects those objects to the business questions they collectively answer.

The result is predictable. Redundant objects get migrated alongside essential ones. Teams spend months reverse-engineering logic only to discover that it serves a report no one has looked at in years. The technical debt that accumulated over two decades is carried forward into the cloud, unchanged.

Perhaps most importantly, a bottom-up approach offers no mechanism for ensuring that the data arriving in the target system is trustworthy. If the definitions were inconsistent in the source, those inconsistencies migrate too. The organization moves to the cloud but inherits the same credibility problems.

The missing piece is a layer that sits above the individual objects. One that asks: what does this data really mean and who uses it? Is it accurate and relevant? And what should the target architecture look like if we design it around business needs rather than legacy structure? This is the gap that a semantic, top-down approach is designed to fill.



3. The Advantage of a Top-Down, Semantic Approach

A semantic approach to migration inverts the traditional sequence. The goal is not to move every object from A to B. It is to understand the full landscape, extract its business value, and then design the target architecture around that exact understanding.

The foundation is a unified semantic model that connects everything the legacy system contains into a knowledge layer. It captures relationships. It traces lineage from source to report. And most importantly, it extracts the business logic buried in custom code and expresses it in platform-independent terms.

This shift also introduces something that bottom-up approaches structurally lack: trust in data. By making metrics traceable and definitions explicit, the semantic model forces transparency. Inconsistencies that were invisible in the legacy system surface during modeling, not after go-live.

The ultimate goal of any migration is to produce data the business can rely on – for KPIs, reports, and AI-driven automation. Data that carries forward undocumented assumptions undermines that goal, regardless of how modern the platform is.

4. From Technical Objects to Valuable Business Assets

Once the semantic model makes the legacy landscape visible and understandable, the real question becomes: what is the most effective, cost-efficient, and fastest way to migrate?

The answer starts with a counterintuitive insight. The fastest path to the cloud is not to move everything. It is to understand everything first and then move only what matters.

This is where the semantic approach delivers its most tangible value. Before a single object is transferred to the target environment, the full landscape has already been mapped, documented, and assessed. And the results are packaged into governed data products. Data products are not just the ideal outcome of a migration, enabling scalable results.


They are also the ideal vehicle for getting there. By defining data products before anything moves to the cloud, organizations create a clear, business-aligned blueprint of what needs to migrate and what does not. Each data product traces data to its business purpose, captures the logic that produces it, and makes quality, lineage, and ownership transparent. Because all of this is built from metadata alone, the organization gains transparency and trusted, governed assets ready for deployment, before any raw data even leaves the source system.

This sequence saves real time and money. Legacy SAP BW environments are full of redundant objects, overlapping logic, and data nobody uses anymore. A bottom-up migration moves all of it into the cloud, at full cost and with the same complexity. A semantically informed migration, grounded in data products, retires what is obsolete, consolidates, and migrates only what genuinely serves the business. The difference in scope, and therefore in cost, can be dramatic.

SAP itself has recognized data products as the key elements of modern data architecture. With the newly introduced SAP Data Product Studio, planned as part of the SAP BDC, SAP is introducing tooling for creating and governing data products natively. However, this approach requires organizations to first move their data into the SAP BDC – including all the untrusted, undocumented, and redundant data from the legacy landscape. Only then can data products be defined. This means the cost and complexity of migrating everything comes before the clarity of knowing what is really needed. A semantic approach reverses this sequence entirely. Data products are built from metadata before any raw data moves to the cloud. The organization arrives at the target system with a clean, governed, and validated set of assets and not a replica of the legacy mess.

A data product for “Customer Churn Analysis,” for instance, carries not just the underlying data, but also the rules that produce it, the definitions that give it meaning, and the quality contracts that ensure it stays reliable. Each data product also embeds governance from the start. Quality standards, freshness thresholds, structural constraints, and service-level commitments are defined before deployment. Every metric is traceable. Every source is assessed. Every definition is agreed upon, before it reaches the target system.


And because the logic and rules are captured in a platform-independent way, the result is future-proof. The same data product can be deployed to SAP Business Data Cloud today and extended to another environment tomorrow. This gives full flexibility over data assets: free to adopt new technologies, switch platforms, or expand without rebuilding from scratch.



5. How a Semantic Layer Operationalizes the Migration

The concepts of semantic modeling and data products are compelling in theory. But how does this work in practice? Building a semantic layer over a large SAP BW landscape and turning it into governed data products requires a platform purpose-built for the task.

One that delivers full transparency, governed data products, and a validated target design before migration even begins. The ideal workflow proceeds through several stages:



1. Ingestion and Knowledge Building

The platform connects to the legacy SAP BW environment and automatically collects all relevant metadata: system objects, ABAP source code, transformation logic, documentation, and structural information.

This material is semantically connected into a rich knowledge layer that combines graph-based relationships, semantic indexing, and contextual retrieval. The result is an interconnected model of the entire legacy landscape: what objects exist, how they relate, what logic they contain, and what business purpose they serve.


3. AI-Driven Intent-to-Code Automation

Once the data products are designed, the platform translates them into executable code for the target environment. The ABAP-to-SQL translation for example, that would otherwise take months, is handled with AI-automation and happens in a much shorter time frame. Because the platform understands the business's intent behind the legacy code, it generates clean, optimized SQL or target-platform-native code. Architects maintain full control while supervising and adjusting the output.

2. Use Case Design and Data Product Definition

With the knowledge layer in place, architects and business stakeholders define the target data products. Rather than starting from a blank sheet, the platform provides recommendations based on its understanding of the legacy landscape.

Where is consolidation possible? What logic needs to be preserved, and what can be retired? Business teams define priorities. Data teams receive precise input. The result: clearly scoped data products with defined ownership, quality standards, and business context.




4. Pipeline Execution and Validation

The generated pipelines are executed and validated in the target system. Governance contracts – quality rules, schema requirements, and service-level commitments – ensure that each data product meets its defined standards before going live. Because the data products were fully defined before migration, the target system receives only what was planned, scoped, and governed.

5. Visibility and Ongoing Governance

The platform provides continuous visibility into the target landscape. What data products exist, who owns them, how do they connect? Do they meet their quality contracts? This layer does not end with migration but governs the organization's data on an ongoing basis.

The value: compressed timelines, reduced cost, and a target environment that is cleaner and more trustworthy than the legacy system it replaces. All because data products were fully defined before the migration began.



6. Beyond Migration:

Building a Foundation for Trusted Data

It is tempting to view an SAP BW cloud migration as a one-time infrastructure project. Move the data, retire the old system, declare it done. But this view misses the larger opportunity.

A migration executed through a semantic, top-down approach produces data products that outlast the migration itself. The semantic model becomes the organization's living documentation of its data landscape. The data products, with their embedded quality rules, lineage, and business definitions, become reusable building blocks. Any team can discover them, understand them, and deploy them to new contexts. Without recreating logic from scratch.

This advantage comes down to the sequence. By building data products before migration rather than after, organizations avoid paying to move data they do not need, arrive at the target system with only clean and governed assets, and gain a reusable governance layer that serves them well beyond the migration itself.

This matters more than ever as organizations rely on AI and machine learning. AI agents and models are only as reliable as the data they consume. If that data carries undocumented assumptions or invisible quality issues, the AI's outputs will be untrustworthy. Governed, transparent data products provide the foundation that trustworthy AI requires.

Trust runs through every aspect of this approach. It starts with transparency: knowing where data comes from, how it was transformed, and what it means. It deepens through alignment: business teams and data teams sharing a common understanding of definitions, priorities, and quality standards. And it scales through automation: packaging logic once and executing it consistently across any target system.

When embracing this semantic, top-down philosophy during an SAP BW migration, the migration becomes the catalyst for building a data estate a business can rely on. For correct KPIs. For reliable reports. For confident decisions. The real migration is not from one system to another. It is from fragmented, undocumented data to data that is understood, governed, and trusted. Move the meaning first. The data will follow.



One Data

Whitepaper: Beyond Code Translation
A Semantic Approach to SAP BW Cloud Migrations
© One Data | ET/CF