PRODUCT ORIENTED ARCHITECTURES

3rd Industry Summit on

Data Product Oriented Architectures

16 and 17 september, Antwerp, Belgium

Wednesday 16 September 2026

08.30	Registration open	
09.20	Opening and Welcome	*Bas van Gils (AMS, DAMA)*
09.30	Keynote 1	*Mathias Vercauteren (AMS, DAMA)*
10.15	Coffee break	
10.45	Implementing Data Products: Bridging Data and IT in Practice at NS	*Claudia Steghuis, Andrea Huck (NS)*
11.30	Data mesh Theory versus Practice at IMEC	*Leander Dillemans (Imec)*
12.15	Lunch	
13.15	Lessons from Building a Federated Data Organisation at ABN AMRO	*Stuti Oberoi (ABN AMRO)*
14.00	Who Governs Data in the North Sea? Interorganizational Data Sharing Investigated Using a Viable Systems Model	*Derkjan Elzinga, Marco Tonnarelli (JADS)*
14.45	Coffee break	
15.15	Modeling relationships for data product policies to enable federated automated governance	*Valeria Maria Fortina, Pierluigi Plebani (Politecnico di Milano)*
16.00	Towards Data Products for Every Purpose: Lessons Learned at KPN	*Olaf Slomp, Jaco Groeneveld (KPN)*
16.30	Announcements + relocate to AMS	*Bas van Gils (AMS, DAMA)*
17.15	Opening minds to impact the world	*Steven de Haes (AMS)*
18.00	*Closing + drinks*	*Bas van Gils (AMS, DAMA)*

In the evening, you are welcome to explore the city.
Please ask us for recommendations for drinner/drinks if you have questions.

PRODUCT ORIENTED ARCHITECTURES

3rd Industry Summit on

Data Product Oriented Architectures

16 and 17 september, Antwerp, Belgium

Thursday 17 September 2026

08.30	Registration open	
09.20	Opening and Welcome	*Bas van Gils (AMS, DAMA)*
09.30	Keynote 2	*Jan Verelst (UA)*
10.15	Coffee break	
10.45	Data Product Builder: Building Your Data Product Contract-first	*Simon Harrer (Entropy Data)*
11.30	Business Capabilities, Business Processes, Business Objects - How Does the Business Drive the Data Product Paradigm?	*Ulrich Teufel (Mercedes-Benz)*
12.15	Lunch	
13.15	Stop Building Data Products in the Data Team: Bringing Data Closer to the Decision	*Cameron Price (Data Tiles)*
14.00	Split or Merge? Scoping Data Products with Domain Driven Design	*Tom de Wolf (ACA Group)*
14.45	Coffee break	
15.15	Semantic Augmentation of Data Product Metamodels to Enable LLM-access through MCP	*Niven Ratnamaheson, Laura Schuiki, Daniel Peys, Michael Behringer (Uni Stuttgart)*
16.00	Agent-Ready Data Products: What Changes When Agents Become Consumers	*Katarina Milosevic (Data Minded)*
16.45	Closing	*Bas van Gils (AMS, DAMA)*
17.00	*End*	

In the evening, you are welcome to explore the city.
Please ask us for recommendations for drinner/drinks if you have questions.

PRODUCT ORIENTED ARCHITECTURES

3rd Industry Summit on

Data Product Oriented Architectures

16 and 17 september, Antwerp, Belgium

Abstracts day 1

Keynote 1 (Mathias Vercauteren)

The DAMA DMBOK is the international reference on data management. DAMA international has launched a project to develop the 3rd edition of this important work. Mathias is a PhD candidate at Antwerp Management School, and is also the project lead for this initiative. He will share his insights on the project and will present an update on the development of DMBOK version 3 - and will link this to data products/data mesh.

Implementing Data Products: Bridging Data and IT in Practice at NS (Claudia Steghuis & Andrea Huck)

At NS (Dutch Railways), data products have become a central element of our data ecosystem. Over the past years, we have established foundations such as structured registration (via NS Databalie), governance frameworks, and clear roles. This has enabled widespread adoption across data teams, but has also revealed practical challenges when moving from defining data products towards implementing and operating them in practice.

In this session, we share our journey from defining data products towards implementing them in a complex enterprise environment. We illustrate how we connect data product architecture to the IT stack, and how concepts such as data product distributions, metadata-driven registers, and data chain monitoring support this transition. We zoom into concrete examples of how to apply data products to operational use cases, of how we aim to productize data ingestion from source applications to our central data platform and how we implement data product management and governance, such as our NS Databalie (the front-end for all metadata about data).

We reflect on challenges we encountered in this next phase, including aligning governance with autonomy, clarifying ownership across business and IT, and ensuring data quality and usability in a product-oriented model. In addition, we explore the evolving convergence of data and application (data) products, and the opportunities this creates for better integration and end-to-end value delivery.

Data Mesh Theory vs Practice at IMEC (Leander Dillemans)

In this session, we discuss how Data Mesh is implemented at Imec. We discuss the four pillars of the Data Mesh reference artchitecture. This is the basis for discussing where we are now, what return on investment we are seeing, and what the challenges are. We will also discuss the next steps for IMEC in this area.

PRODUCT ORIENTED ARCHITECTURES

3rd Industry Summit on

Data Product Oriented Architectures

16 and 17 september, Antwerp, Belgium

Abstracts day 1

Lessons from building a federated data organisation at ABN AMRO (Stuti Oberoi)

Data products promise value, but struggle to deliver it at scale. Drawing on the journey of implementing a federated data governance model 3 years ago, this talk explores how to translate business demand into tangible data value. Using real examples from banking, I will demonstrate how data products emerge from business processes, how they must work together across domains to enable end-to-end use cases, and why governance, ownership, and organisational design are critical to success. Finally, I will share practical lessons on the realities of adoption, cross-domain collaboration, and value realisation.

Who governs data in the North Sea? Interorganizational data sharing investigated using a viable systems model (Derkjan Elzinga and Marco Tonnarelli)

The protection of critical infrastructure requires a data ecosystem wit a data-driven governance structure. Public and private companies bring together data in a data space where multiple scopes of governance interact on different levels. The Viable Systems Model (VSM) is used in a case study in the maritime security domain to understand the relationship between different decision nodes. Focusing on the data space that allows for real time maritime situational awareness. This requires bringing together and maintaining large streams of data, constrained by the fragmented landscape, strict legal requirements and need for high reliability and certainty. Based on expert interviews, a framework is created that combines different scopes of governance, like operational, technological and ecosystem governance and shows through the use of the VSM how decisions and control mechanisms and the interaction between them can be identified and improved within social-technical systems, such as data ecosystems.

PRODUCT ORIENTED ARCHITECTURES

3rd Industry Summit on

Data Product Oriented Architectures

16 and 17 september, Antwerp, Belgium

Abstracts day 1

Modeling relationships for data product policies to enable federated automated governance (Valeria Maria Fortina, Pierluigi Plebani)

In today's technological landscape, data products have become strategic assets within federated ecosystems based on data mesh architectures. However, their effective adoption is hindered by a critical open challenge: the governance of the policies that regulate their use. To address this issue, the Policy Relationship Model has been introduced to formalise data product policies and classify the types of relationships between them, including conflicts, in a data-sharing context. Building on this model, a module for self-serve data platforms is proposed that automates key steps in the policy lifecycle. Policies expressed in natural language are translated into executable Rego representations, and relationships between new and existing policies are automatically identified and resolved. Large Language Models are investigated as potential automation tools to further reduce the manual effort required for policy generation and consistency enforcement.This line of research contributes to the state of the art in data mesh ecosystems by offering an approach to policy lifecycle management and governed data sharing.

Towards Data Products for Every Purpose: Lessons Learned at KPN (Olaf Slom, Jaco Groeneveld)

Join us at the summit as we share KPN's journey in establishing Data Products as a scalable foundation for data-driven innovation. In this session, we will explore the different types of Data Products we are implementing, how each type requires tailored Data Contracts to ensure clear ownership and reliable interoperability, and the Data Quality checks needed at each level to build trust and consistency across the data ecosystem.

Through practical lessons learned and examples, we will demonstrate how aligning product thinking, governance, and quality management enables us to turn data into a sustainable strategic asset.

Opening Minds to Impact the World (Steven de Haes)

The mission of Antwerp Management School is: Opening Minds to Impact the World. We aime to make a contribution that lies at the intersection of scientific theories, heuristics/ practical theories, and real-world challenges. In this talk, prof. de Haes will share his insights of this view in relation to data product oriented architectures.

PRODUCT ORIENTED ARCHITECTURES

3rd Industry Summit on

Data Product Oriented Architectures

16 and 17 september, Antwerp, Belgium

Abstracts day 2

Keynote 2 (Jan Verelst)

Prof. Verelst is a professor at the University of Antwerp. His research aims at evolvability of sofware systems. He is one of the founding fathers of Normalized Systems Theory. In this keynote, prof. Verelst will discuss his views on Data Product Oriented Architectures.

Data Product Builder: Building Your Data Product Contract-first (Simon Harrer)

Contract-first means one thing: the contract comes before the data product, not after. You don't extract tables, explore, and document later. You agree on what consumers need first, write it down as a data contract covering ownership, schema, semantics, quality, and SLAs, and only then build the product to satisfy it.

This ordering is what makes data products trustworthy, and its exactly what coding agents need. A contract written first is a spec: clear, testable requirements an agent can build against instead of guessing from raw data. Contract-first is spec-driven development for data, and the contract is the spec.

In this talk, we'll go contract-first end to end, live. The contract is specified first, in the Open Data Contract Standard (ODCS), and only then do we implement it. Working from that spec in a live coding session, a coding agent uses skills to build a dbt project that implements the data product: discovering the upstream source data products it needs through a data marketplace, requesting access to them, and wiring them into the dbt models that satisfy the contract. The skills come from a skill repository that captures the conventions for building a data product on your stack, in our case a dbt project on Snowflake, so the agent follows your team's way of working rather than inventing its own. Because the contract is executable, it stays the source of truth: a living test suite, run via the Data Contract CLI against the actual data, that holds both the data product and the agent to the spec in your terminal and in CI.

You'll leave convinced that contract-first is the right default for building data products, especially when agents do the building, and you'll know how to start.

PRODUCT ORIENTED ARCHITECTURES

3rd Industry Summit on

Data Product Oriented Architectures

16 and 17 september, Antwerp, Belgium

Abstracts day 2

Business Capabilities, Business Processes, Business Objects - how does the business drive the Data Product paradigm? (Ulrich Teufel)

Non IT people often think in different dimensions when creating data artifacts than their IT counterparts. This begs the question: how can we approach a data mesh from a business-driven perspective, using a top down thinking rather than bottom up? At Mercedes-Benz , we now offer logical patterns to detect the gaps or pas in information flows of a company in order to derive fitting business-driven data products.

Stop Building Data Products in the Data Team. Bringing Data Closer to the Decision (Cameron Price)

For years, the data industry has operated under a simple assumption, that data engineers build data products.But this assumption breaks down in modern organizations.Most decisions are made far from the data teams who design pipelines, warehouses, and semantic models. The result is familiar, with months of engineering work, layers of transformation pipelines, and yet the final insight still arrives too late, or in a form that decision-makers cannot use.

This session challenges the underlying assumption of modern data architectures and proposes a new model, that data products designed around decisions, not pipelines.Through real-world examples, this talk demonstrates how organizations are shifting from engineering-centric data product development to decision-centric data products that bring data closer to where decisions occur.

The presentation includes: a detailed case study showing how a complex decision workflow was redesigned into a decision-centric data product, including architecture, governance, and delivery patterns, a step-by-step framework for identifying decision points and translating them into actionable data products, a practical architecture blueprint that reduces dependency on centralized data engineering pipelines by leveraging semantic layers, metadata, and governed access patterns, and before-and-after metrics demonstrating improvements in time-to-insight, adoption, and operational outcomes.

PRODUCT ORIENTED ARCHITECTURES

3rd Industry Summit on

Data Product Oriented Architectures

16 and 17 september, Antwerp, Belgium

Abstracts day 2

Split or Merge? Scoping Data Products with Domain Driven Design (Tom de Wolf)

'The Data Mesh paradigm promised a decentralized data landscape, but many organizations are currently stuck in architectural confusion. Without clear boundaries, the pendulum swings between two extremes: massive, unmanageable data monoliths, or a hyper-fragmented nightmare where every single database table pretends to be a "Data Product."

Finding the "just right" boundary is the defining challenge of data product-oriented architectures, and answers can be found in a discipline software engineers have used for decades: Domain-Driven Design (DDD). In this session, we will bridge the gap between DDD and data product design, providing you with concrete design heuristics to determine what belongs inside a single data product and when it's time to split them apart.

Semantic Augmentation of Data Product Metamodels to enable LLM-access through MCP (Niven Ratnamaheson, Laura Schuiki, Daniel Peys, Michael Behringer)

In data mesh architectures, leveraging data products through LLM interfaces requires standardized access and robust semantic context. This presentation addresses this gap by extending the ProMoTe metamodel to integrate LLM-based access via the Model Context Protocol (MCP) and semantic models. We evaluate three architectural variants, (1) a non-augmented baseline, (2) a RAG-based approach and (3) a knowledge graph (KG)-based variant with and without re-ranking. To evaluate these approaches, we leverage a subset of PrimeKG, a knowledge graph designed for medical analysis, as ground truth. In this talk, we show our results and discuss how augmenting data products with semantic models can drastically improve LLM performance.

Agent-ready data products: what changes when agents become consumers (Katarina Milosevic)

In this talk we offer a practical view on what data leaders should prioritise now to make their data product landscape ready for a world in which AI agents are data product consumers. When AI agents become consumers of data products, most of what we already know about good data management still holds, but a few things break, and a few things finally get the attention they deserve.

In this talk we share what we have learned from building agent-ready data product flows with clients. Agents will not walk over to Tom from sales to ask what a renewal really means. The tribal knowledge that has lived in scattered files and people's heads has to become an artefact of the data product itself. We look at how the producer–consumer relationship shifts when an agent becomes a consumer, and which data product principles need to be extended. We discuss the renewed role of semantic models, where to put them, why vendors interpret them so differently, and where the industry seems to be converging.