



Enterprise Data Mesh Services for the Mission at the Edge

Productize data to accelerate value

Contents

1.0 Overview

2.0 Mesh Architecture

3.0 Core Principles

4.0 Lifecycle of a Data Product

5.0 Data Mesh in Action

6.0 Management & Governance

7.0 Summary

Overview

Data is essential for organizations to survive and thrive in edge environments where competitors can rapidly organize data to overcome traditional advantage.

Every organization should be as strategic about acquiring, growing, refining, safeguarding, and deploying data as they are about traditional system capabilities.

Data is generated in all-domains. Data comes from humans, machines, Internet of Things (IoT) devices, edge systems and more—from earth and even in space. All that data adds up. In fact, a staggering 2.5 quintillion bytes of data are generated each day.

CEOs who see data as capital can grow and reinvent their businesses to be data-driven. The [2021 Chief Data Officer \(CDO\) Survey by Accenture and MIT](#), 78% of CDOs told us that their most critical responsibilities are driving business growth and value creation.

For the Department of Defense (DoD) growth means greater adoption of VAULTIS goals and value creation means more capabilities.



Stop creating Platforms

Advanced for their time, these are rapidly turning into silos of excellence



Start creating Services


Enterprises drive capabilities through services offered to data domains to enable data products



Build data Products


Enable the rapid creation of data products from complex scenarios to enable the strategic edge

Current State: Keeping up with organisational change and business needs...



Traditional Ideal State

A Systems View



Today, we approach data with a systems mindset.

A mission function has a specific problem that needs to be solved with data.

We utilize physical operations centers to manually mine data through a systems-driven approach.

Mission thread operators acquire data from the relevant sources, cleanse and prepare it, then analyze it for the specific use case.


And each time a new mission problem or use case arises, it follows a similar process to acquire, prepare, and analyze data for specific needs.

This systems approach has some significant drawbacks:

- May lead to duplication in resources
- Relatively slow to inform mission functions
- The outputs from each system typically cannot be repurposed to solve other use cases

Traditional Objective State

A Platforms View



Data platforms are designed to integrate modern capabilities for the preparation and analysis of data to rapidly address long-standing use cases

Multi-variate data, however, **increases the number of sources** from which domain data is acquired (e.g., all-domain, multi-system, sensor data, third party data), which poses notable challenges:


- Integrating **all data in one centralized platform becomes unmanageable**, with prohibitive governance, security, and cost
- Inability to meet rapid **time to mission** for new data

Through platform growth, the architecture becomes constrained by its **decomposed, mechanical functions** (ingestion, cleansing, aggregation, serving, etc.)

- **Adding new data often requires a change in all components** of the pipeline
- **Teams** at each processing stage **cannot work independently**, as the platform is not decomposed by its architectural quantum
- **The platform is decomposed orthogonally on the axis of changes**, which leads to lower scalability and velocity due to increasing volume over time

Future State

A Services View



Siloed talent means we must move beyond platforms:

- Data platform engineers are **organizationally siloed** from the source domains that generate the data and **hyper-specialized** on their technical expertise
- Platform engineers often lack business and domain knowledge

Data Mesh is a new approach based on a modern, distributed architecture for analytical data management. It enables data users to easily access and query data where it lives without first moving data to a data lake or data warehouse.

The faster access to query data directly translates into faster time to value without needing data transportation into a platform.

The decentralized strategy of Data Mesh distributes data ownership to domain-specific teams that manage, own and serve the data as a product utilizing enterprise services.

Data Mesh Operational Benefits:

- **Eliminates the challenges of data availability and accessibility at scale.** Data mesh allows data users to access, analyze, and operationalize business insights from any data source
- Makes data **discoverable, secure, and interoperable**
- **Addresses data pipelines issues**, which often become problematic over time by creating point-to-point systems
- **Addresses enterprise issues, such as business areas disagreeing on core facts of the business.**
- Brings much needed order to a system, resulting in a **more manageable and evolvable data architecture**

Core Service Principles

The core principles make Data Mesh the convergence of **distributed domain driven architecture, self serve platform design, product thinking and governed data.**




Data products provide organizations a better way to address their data needs. First, the data team determines the organization's data requirements, looking at individual business units and across the enterprise. Then, they create reusable data products that help multiple business functions solve different use cases.

Data products can be:

- **Data**—reusable datasets, streaming, feeds, and APIs (e.g., for design, manufacturing, finance and operations) that meet the needs of the whole enterprise, as well as each business function.
- **Code**—feature code and transformation snippets (small blocks of reusable code) or data models (representations of the logical structure of how different data elements fit together).
- **Analytics models**—reusable machine learning models (e.g., for predictive maintenance)
- **Dashboard reports**—reusable dashboards and other visualizations (e.g., a risk dashboard).

Domain Driven Distributed Data Mesh



Data mesh, at its core, is founded in decentralization and distribution of responsibility to people who are closest to the data in order to support continuous change and scalability.

----- CONSUMPTION ALIGNED Domains

- Fit for consumer purpose
- Changes more often
- Can be re-created

----- AGGREGATE Domains


- Integration across domains
- Aggregations, projections, models
- Derived / secondary data products

----- SOURCE ALIGNED Domains

- Facts and reality of business
- Immutable timed events
- Historical snapshots
- Change less frequently
- Permanently captured

Product Defined

Analytical data provided by the domains must be treated as a product, and the consumers of that data should be treated as customers.




Data Product Characteristics

- Discoverable [I can find it]
- Addressable [I can consume it]
- Self-Describing [I understand it]
- Trustworthy [I can trust its quality]
- Secure [It is protected from unauthorized access]
- Interoperable [I can use it with other data products]

From Mesh to Product


The Mesh is an architectural concept **defined by the services an organization creates to enable data products**. These services can derive from a platform, but they are agnostic to any technology stack. Data products are creations on the mesh. We define Data Products as artifacts created for specific domain-driven purposes:

-  → **Core Data Products** – foundational data owners for the business
-  → **Derived Data Products** – datasets created by combining multiple core data products
-  → **Use Cases** – mission-oriented



Lifecycle of a Data Product

Much like commercial products, data products have a lifecycle that consists of multiple stages. The product is conceived, then designed, developed, launched and managed in the market. Finally, the product is stored in the data product catalog when the real-time need for it diminishes.



Conceptualize—understanding the needs of the mission, to ensure that the product will meet the needs of end-users and deliver the expected ROI. At this stage, a product manager is assigned.



Design—team assesses the requirements and decides whether the product is a core foundational product (that needs to be created from scratch), or a derived data product that is composed of multiple core data products, the team also create a conceptual data model for the product.



Manufacture—a cross-functional team comes together to create the product. Depending on the type of product, the team can include everyone from analysts, data engineers and DevOps engineers to data modelers, data scientists and visualization engineers.




Rollout—typically, the team provides an initial minimum viable product (MVP) to a select group of users. This involves putting the MVP into production, setting up support teams, training users, and establishing mechanisms to track product usage. As the team iterates on the product, it rolls out new versions.



Service—once the product has been rolled out and people are using it, it needs proper support and services. The support team should monitor the data pipelines associated with the product. Centre of excellence (CoE) teams should be ready to advise people on how best to use the product. Stewards should continue to monitor product quality. And operations personnel should track product usage. When the product is no longer needed, it's retired.

Lifecycle of a Data Product

Much like commercial products, data products have a lifecycle that consists of multiple stages. The product is conceived, then designed, developed, launched and managed in the market. Finally, the product is stored in the data product catalog when the real-time need for it diminishes.



Example RACI Responsible, Accountable, Consulted, Informed


		Data Product Manager	Executive Sponsor	Central Data Team	Data Architect	Product Team: SW Developers, Data Engineers, Data Scientists	Operations
Conceptualize		R	A	C	C		
Design		R	A	C	R	C	
Manufacture		C	A	C	R	R	C
Rollout		R	A	C	C		
Service		I	A	C	I		R

Accenture Data Mesh Pattern

At Accenture, we use these same practices to manage our own data products on the Accenture Data Mesh. We’ve created ready-to-use data products to support our data users and consumers such as executives, strategists and analysts.

We’ve also built exploratory data products—e.g., curated data, AI models, APIs, and data models—for our advanced data users, such as data scientists and data analysts. We use the Data Mesh to orchestrate the commerce of these data products within our organization and our 650,000+ people. The Data Mesh provides a coordinated marketplace for product consumption and usage analytics that gives us powerful insights, such as:



- **how well each product is being used.**
- **which new products we should introduce, based on frequent searches with no results.**
- **products that should be recommended to customers based on their interactions with other products.**




Service: Mercury Project

The core principle of Mercury is **platform abstraction**: this refers to the ability to separate business logic from the underlying network, platform components, databases and open sources frameworks and tools:

- Open-source project:
<https://github.com/Accenture/mercury>
- This means freedom of choice for developer & data product teams to extract data and operate tools that are domain-specific on an event driven framework.
- Mercury powered applications are designed to evolve with the mission context



Data Mesh Architecture Example




Enabling digital transformation

Linking real-world event-driven data published to the mesh could be used in digital twins to support rapid modeling and simulation.

Outputs from the digital engineering environment are then developed using a DevSecOps framework and integrated as new features back into fielded living systems on the mesh, for effects that support recent user needs. The Data Mesh underpins the domain-driven ecosystems from data products that are requested to produce effects that support the mission.

For starters, data enables the mission to introduce new remote diagnostics, optimization, and predictive maintenance services in the digital engineering domain.

What's more, living systems at the edge unlocks unique visibility of the mission use cases using the users' perspectives. This enables the organization to provide rich insights into critical mission success factors at the time of need.



Data Mesh Enables Products that are:

- **Composable:** Composability means the individual components can be developed and evolved as products. Core data products are based on foundational data that is key to supporting the business. Other products can be derived (composed) by combining multiple core data products
- **Discoverable:** This ensures data visibility and democratization, and prevents duplication of effort
- **Interoperable:** Different data products can connect, because they interoperate (E.g., they have similar interfaces). This interoperability supports the composability point we mentioned earlier. It also prevents the creation of data silos using non-standard technologies and standards
- **Governed:** Data products need proper governance. That means clear ownership, lifecycle management, security and access policies, as well as well-defined consumption mechanisms
- **Purposeful:** Just like any other product, data products should be created to fulfill clearly defined purposes and deliver demonstrable value to identified consumers.



The right skills and technology to create and manage data products

To create data products, organizations need a broad range of data management capabilities. These include data acquisition, data integration, data quality, metadata management, data modeling, data cataloging, marketplace, data access and APIs, data rights management and data tracking. Organizations also need cross-functional skills, such as business analysis, DevOps and data engineering. Plus, the teams creating data products need knowledge of the relevant industry and domain.

The fundamental goal is to use data as a competitive differentiator. That means organizations will want to be able to update and release new data products all the time. So, it makes sense to create them in a data platform that is built in the cloud. [Cloud enables scale, agility and the opportunity to drive reinvention.](#)

It allows for data to be connected as a part of a larger continuum. And by tapping into the [Cloud Continuum](#) organizations are enabled to productize their data—wherever it resides. Indeed, many of our [clients are unlocking value by harnessing the power of data and AI in the cloud.](#)

Summary

DATA Platform



Monolithic

- Data stored in one giant repository

Centralized Ownership & Governance

- Data loaded in one central storage repository in raw format, which is managed by a centralized IT organization

Hyper-specialized Silos

- The data engineering function is typically separated from both data producers and consumers

Pipeline as first-class concern

- One data pipeline for all the services and domains
- Produce, transform and serve all data inside this pipeline

Data as a by-product

- While storing raw data in a data lake, it is modeled and transformed **only while** insights are generated
- Analytical data is a by-product of modeling operational data



DATA MESH

Distributed

- Data stored in repository of each domain

Decentralized Ownership & Federated Governance

- Both operational and analytical data is maintained by each service on its own

Domain shared knowledge

- Domain data owned and served by a cross-functional team

Domain as first-class concern

- The sources of data (e.g. business units) represent the data domain and are responsible for preparing and serving the data

Data as a product

- Data is prepared and treated as a product
- Focus on data quality, decreased lead time of data consumption and data user satisfaction

About Accenture

Accenture is a global professional services company with leading capabilities in digital, cloud and security. Combining unmatched experience and specialized skills across more than 40 industries, we offer Strategy and Consulting, Interactive, Technology and Operations services — all powered by the world’s largest network of Advanced Technology and Intelligent Operations centers. Our 674,000 people deliver on the promise of technology and human ingenuity every day, serving clients in more than 120 countries. We embrace the power of change to create value and shared success for our clients, people, shareholders, partners and communities. Visit us at [accenture.com](https://www.accenture.com)

Copyright © 2022 Accenture. All rights reserved. Accenture, its logo, and High Performance Delivered are trademarks of Accenture. This document makes descriptive reference to trademarks that may be owned by others. The use of such trademarks herein is not an assertion of ownership of such trademarks by Accenture and is not intended to represent or imply the existence of an association between Accenture and the lawful owners of such trademarks. This document is produced by consultants at Accenture as general guidance. It is not intended to provide specific advice on your circumstances. If you require advice or further details on any matters referred to, please contact your Accenture representative.