



# Unlock 3x ROI with AI – with data you already own

The case for composable data services in  
financial services.

# Introduction

## How do you unlock more AI ROI?



**By Christopher Sly**  
Engagement Director - BFSI

**The idea for this whitepaper came from a flurry of questions we've been getting from FSI organisations. They're chasing the next shiny data and AI thing, whether it's Databricks, Snowflake, automating data pipelines, spinning up sandbox environments or implementing data lakehouses to support agentic.**

The race for good AI has never been higher priority. Yet the excitement about that next shiny thing keeps getting tempered by the age-old computer science concept of garbage in/garbage out (GIGO). Data volumes are increasing exponentially, yet 67% of organisations don't completely trust the data used for decision-making. And if we don't trust the output, how can we expect AI models to deliver the value outcomes we want?

There's clearly a gap between the appetite for innovation and the insight quality stakeholders are getting. This whitepaper is about bridging that gap.

**Our thesis is that a composable, microservices approach to data is key to bridging that gap.**

For every **\$1**

organisations invest in generative AI, they're realising an average of \$3.70 in return

**67%**

of organisations don't completely trust the data used for decision-making

A reason there's such a big gap in trust and insight effectiveness is because underlying data infrastructure doesn't allow governance to scale with the required consistency or flexibility.

In this whitepaper, we make the case for addressing this with **composable data services.**

Composable data services are modular, reusable data building blocks available through a governed self-service platform. The platform makes it quick and easy for teams to assemble new data flows, insights or applications.

## **As a result, you get a data infrastructure that fulfils 4 core principles of good data:**

- 1 Decentralised data ownership**  
No one team holds the keys to the data kingdom
- 2 Data as a product**  
Because data is productised rather than bespoke, it's easier to use
- 3 Self-serve platform**  
People can access data themselves as they need it
- 4 Federated computational governance**  
This is the principle that the right people get the right part of the data banquet

## **In the following chapters, data experts will explain this in more detail, including:**

- How composable data services work
- Why this microservices approach to data is so effective
- Best practice data architectures
- How to achieve the right security

# Chapter 1

## Composable data services explained



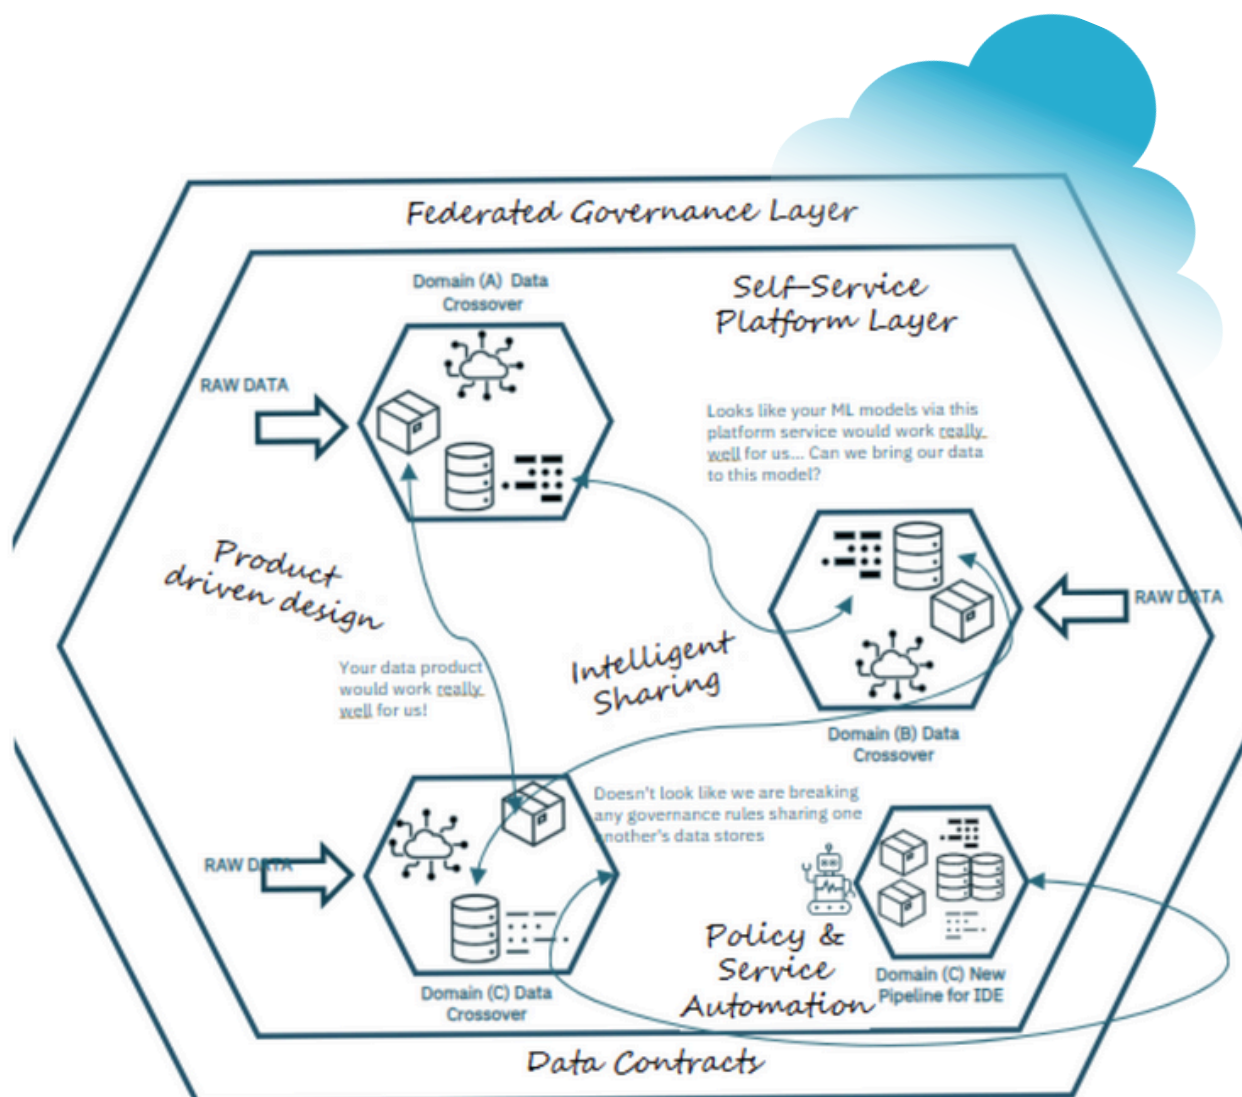
**By Michel Tardio**  
Industry CTO, Data SRE Lead  
and Infra Specialist

**Composable data services are modular, reusable data building blocks – like APIs, data products, transformation pipelines and ML models – made accessible via a governed self-service platform.**

These services enable teams to rapidly assemble new data flows, insights or applications without duplicating effort or relying on central IT. At the core is a hyperscaler-native platform, configured with federated governance across identity, access controls and data-sharing policies – a critical foundation for secure, domain-to-domain interoperability. This approach supports decentralised ownership while maintaining enterprise-wide compliance and oversight.

## Using the self-service platform, developers and scientists have the tools to:

- Discover and search existing data product components (data as raw/tables, scripts, models)
- Access and request data via governed workflows
- Compose or build their own data pipelines or apps from pre-approved components
- Monitor usage and ensure compliance



only **41%**

of live AI solutions  
are delivering value

**49%**

of organisations cite  
data availability as  
the #1 obstacle to AI  
implementation

The self-service experience is accelerated by automation in discovery, deployment and quality assurance. Users interact through a graphical UI-based portal that supports quick experimentation with sandbox environments, enabling teams to test, learn and promote data services into production with speed and confidence.

## Benefits of engineering a self-service platform to deliver composable data services

- ✓ Reduces manual ETL
- ✓ Enables domain expertise
- ✓ Removes approval and accelerates build times
- ✓ Is repeatable as standard
- ✓ Saves large amounts of money thanks to re-useable components
- ✓ Unlocks cross-collaboration that drives innovation

## Chapter 2

# Microservices worked for infrastructure, so why not data?



**By Christopher Cooke**  
CTO – DevSecOps

**When I first heard the term ‘composable data services’, I thought it might be an abstraction too far. But it’s actually a pragmatic shift in how data work can be improved – especially for organisations serious about scaling insights without scaling chaos.**

As someone working extensively in DevOps, it makes sense that the benefits realised from migrating monoliths to microservices would reap similar benefits for data products. With expanding requests to adopt DataOps and MLOps, I’ve seen first-hand how infrastructure, applications and data practices have evolved – and where they’re converging.

# 149 ZB

The total amount of data consumed in 2024

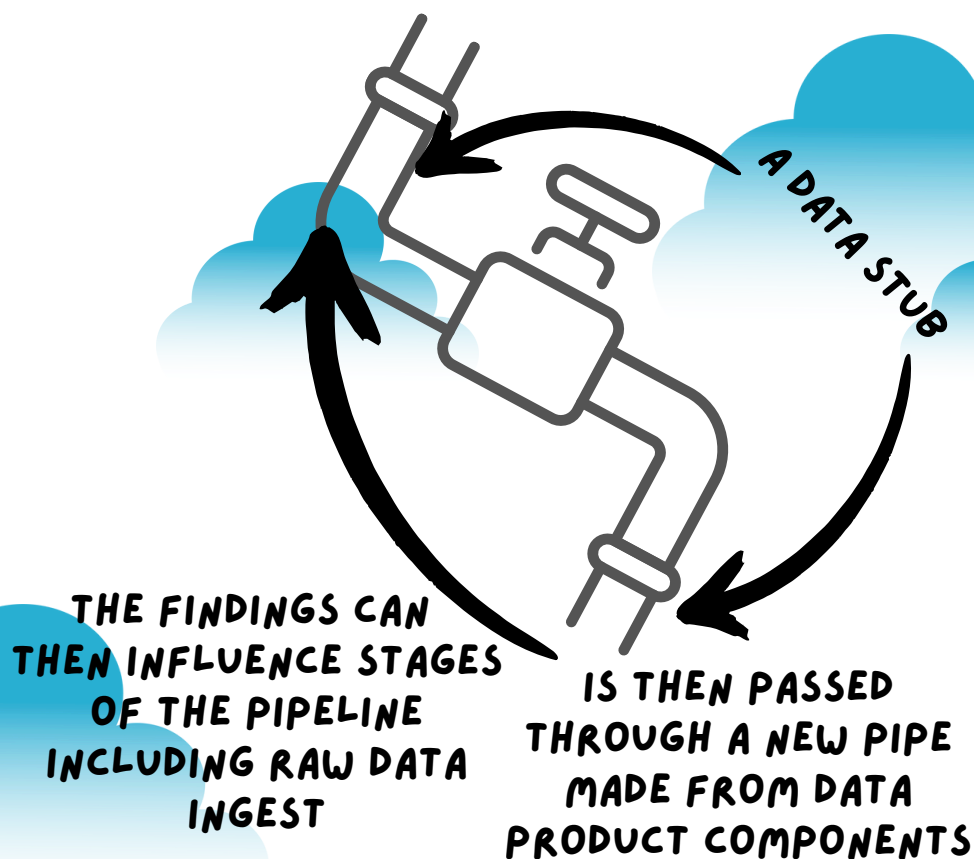
# 1 ZB

Is roughly as much information as there are grains of sand on all the world's beaches

In DevOps, we've always believed in modularity, reusability and CI/CD. Composable services bring that same discipline to the data layer. You define your data products like components – exposed through APIs, versioned, validated and observable. They're not just tables in a lake anymore, they're data services with SLAs, contracts and owners. That changes everything from the 'people versus product' perspective.



What really clicks for me is how this aligns with federated governance. In DevOps, we've always had to balance autonomy with control – giving teams the freedom to ship fast, but within guardrails. Composable data services do the same for data teams. Domains own their products, but access, sharing and compliance are centrally governed and automated. It's how you avoid spaghetti pipelines and rogue dashboards while still enabling cross-functional collaboration. It's DevOps, but for data.





**Raiffeisen  
Bank**

## **Raiffeisen Bank International: A self-service data platform story**

To meet the needs of a rapidly changing customer environment along with growing competition, RBI needed to change the way it worked.

**This included embracing new technologies to improve speed and agility for:**

- Bringing new products and services to market
- Modernising IT architecture
- Optimising new cost models

We helped them design and implement a delivery model combining a Cloud Centre of Excellence for the platform engineering backbone with over-the-top business enablers for compute and data services.

As a value outcome, agility was the priority – and with a DevOps approach to data services and infrastructure-as-code, RBI teams could deploy data platform services in a single click.

**[Read the full story](#)**

Nordcloud was an essential partner in building RBI's data platform. Without their experience, this journey would not have been possible.

**Philipp Pfneisl**  
Delivery Lead, RBI



# Chapter 3

## Re-thinking end-state data architecture



By Allan Chong  
Head of Data & AI

**Global banks process over 463 exabytes of data annually, yet only 15% of that is being used with AI models – and only 16% of UK banks classify AI as high materiality.**

While FSI organisations have been investing in centralised improvements to drive materiality with AI, the challenge is bigger than a single technology. There's no one-tool-fix-for-everything scenario. Therefore, every forward-thinking bank or financial services firm needs to be asking this question:

*'Is our underlying data architecture truly equipped to unleash AI's full, revolutionary potential?'*

For most organisations, the answer is no.

**Only  
16%**

Of executives say AI has high materiality to their core operations.

## **Redefining the architectural rule book**

Traditional architectures like data lakes, data warehouses or the growing blueprint of data lakehouses often focus on a tool-led end state – typically a single repository or platform where all data resides with integrated and/or in-built analytics. While these models offer benefits like centralisation and scalability, they can be a reason why FSI organisations struggle to address the dynamic needs of modern AI systems.

### **Why? Because these models:**

- ✓ **Often prioritise storage over agility**
- ✓ **Assume data will primarily flow in one direction**
- ✓ **Can neglect the most important aspect of data: people**

In reality, AI thrives on dynamic interactions between data producers (sources of raw and processed data) and consumers (applications, models and analysts). **Therefore, static architectures can create bottlenecks – limiting the ability to adapt to new data sources, integrate diverse formats and/or address emerging use cases.**

## **You need a dynamic architecture to unlock AI value**

Therefore, you need to look at architecture as ever-evolving, not a one-time build or static endpoint. This enables you to treat data as a product. The product isn't just a dataset – the composable opportunity comes from a combination of raw data, tabular data, people, tooling and processes.

Consider the feedback loop between machine learning (ML) models and the data they generate. Predictions, classifications and insights derived from AI models often become new data points that feed back into the system for enriching future iterations.

**Only  
15%**

Of data banks process is being used with AI models

# An evolving architecture embraces this interplay and delivers 3 value outcomes:

1

## Scalability

Domain accessibility via sharing policies, reuseable resources and pipelines with lower costs

2

## Quality

Domain-influenced improvements that are accountable to data products and outcomes

3

## Uncovering new

An integrated development environment directly accessing new services, products and information as a tailored pipeline



**£1.56  
billion**

Predicted cloud-enabled data lakehouse spend for the FSI market in 2025

When you embrace a dynamic architecture, you have a secure way to unstick people and processes – improving producer output, unlocking hidden data and ultimately increasing insight quality. What does a best-practice dynamic architecture look like? Prof Faisal Nazir answers that question in the next chapter.



# Chapter 4

## A best-practice approach to dynamic architecture



**By Prof Faisal Nazir**  
CTO – Digital Twin

**As Allan just explained, to unlock ROI from AI, you need to an architecture that treats data as an evolving service, not an optimised and final fixed asset. Now let's look at what that looks like.**

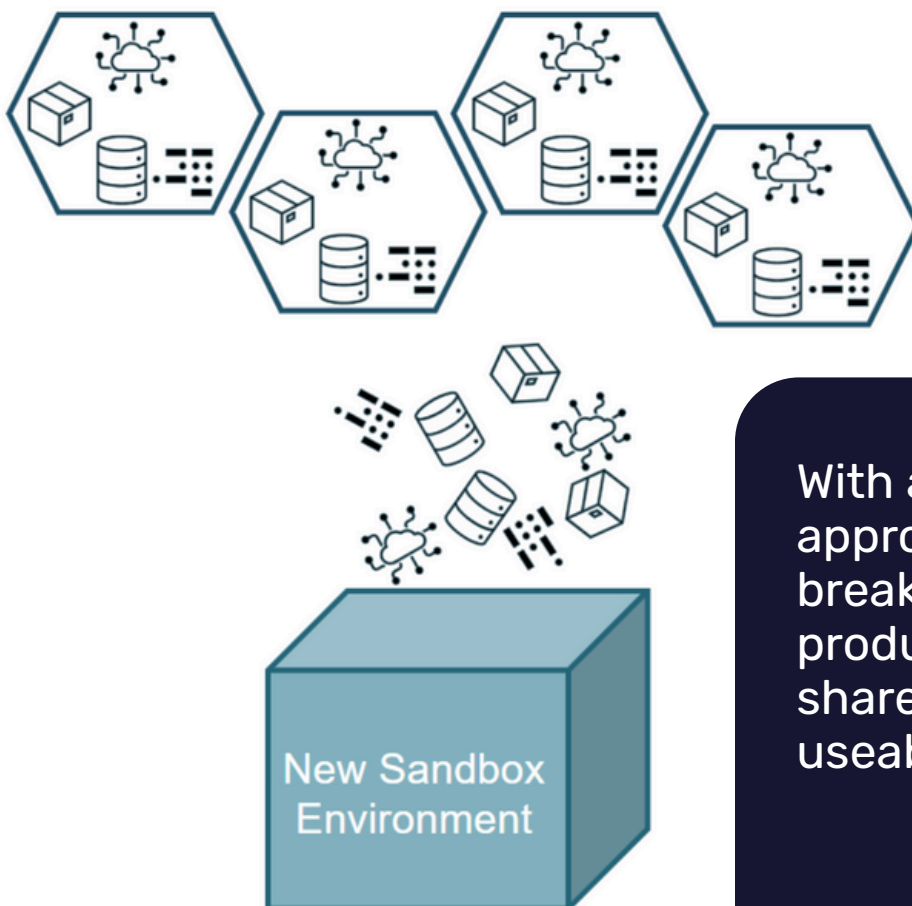
Imagine you have a factory (your data pipeline) that processes raw materials (data) into a product (insights or dashboards). In this scenario:

- Underlying governance and security enable people to connect with relevant data product components
- There's a sandbox environment to visualise new hidden insights
- Data pipelines are orchestrated as a composable data service
- You have a consolidated self-service platform where automation uses (one or many) data sources, artefacts and visualisation tools

Now imagine having to build the same complete factory over and over again in different places (environments like dev, test and production) just to test minor changes. This would be costly, slow, hard to keep consistent and wasteful.

**Introducing composable data services as modular, reusable components means you build only what you need, when you need it.**

With composable data services, you don't have to rebuild the whole factory. Instead, you plug in the machine you want to test, in a controlled environment, with a time-sensitive data stub.



**90%**

Value return from standardised, repeatable data products versus the build of individual pipelines

With a microservices approach to data, we break down data products into shareable and re-useable components.



# What good looks like

When transitioning from centralised end state to dynamic data architectures, FSI organisations should:

- ✓ **Shift mindset**  
Responsibility for data must move from a central team into domains
- ✓ **Adopt open standards**  
Use open data formats and APIs to ensure interoperability
- ✓ **Automate governance**  
Remove shadow IT with centralised non-negotiables and enable a subscription-based model where underlying patterns see variation with trusted data producers
- ✓ **Prioritise observability**  
Implement robust monitoring systems to track data flows, detect anomalies and optimise performance
- ✓ **Go green by design**  
Minimise cost and data errors by avoiding redundant pipelines and duplicated ETL processing. A core pipeline can be re-directed to environments instead of having repeated pipelines across those environments
- ✓ **Foster a culture of collaboration**  
Encourage cross-functional teams to work together, breaking down silos between data engineers, platform engineers, scientists and business stakeholders
- ✓ **Iterate incrementally**  
Start with small, high-impact use cases and scale over time, demonstrating value at each step

# Chapter 5

## The architect's perspective



**By Nachiyappan  
Lakshmanan**  
Cloud Architect – Data & AI

**As a data owner for platform architecture and engineering, I'm seeing unprecedented client urgency to better grasp collective insights.**

And this has been a more experimental period than any I've previously experienced. We used to live in a world where nearly every data request flowed through a central team –whether it was ingesting raw feeds, building dashboards or managing access policies across departments. Now, these centralised efforts seem like bottlenecks (albeit unintentional ones). And the new data approach needs to address this.

**85%**

of executives highlight a widespread issue with disconnected insights and of those customers with live AI production implementations

## The decentralisation drive

The winning organisations will be those who shift from centralised teams acting as sole providers to becoming domain-aware data consumers and enablers. This is where composable data services come in – but because they enable the necessary tooling and mindset.

**To achieve this, the following capabilities are needed:**

- ✓ **Data product catalogue**  
A searchable UI that shows available data assets, APIs, metrics and pipelines
- ✓ **Composable building blocks**  
So users can mix and match datasets, APIs, transformations and models
- ✓ **Governance and access control**  
To ensure security, PII protection and auditability
- ✓ **Pipeline builders**  
No-code/low-code tools (or SDKs) that let users create data flows
- ✓ **Packaging and versioning**  
Data products should be versioned and managed like code artefacts
- ✓ **Observability**  
Dashboards for lineage, quality, usage and performance

## Chapter 6


# A secure, compliant way to implement composable data services



By Simon Casey  
UK Industry CTO

**You don't need a rip/replace strategy to implement composable data services. In fact, hybrid allows FSI organisations to blend on-premises, private and public cloud data product components into a single, orchestrated environment.**

This unified hybrid infrastructure enables you to place sensitive data in secure, compliant environments while experimenting with public cloud platforms to optimise both security and flexibility.





# 1/3


of all composable application deployments are within banking and financial institutions

In today's rapidly evolving financial services landscape, technology isn't just disrupting – it's actively transforming core banking functions like lending, payments, treasury and capital markets. We see a clear emphasis on how financial institutions are embracing agile approaches to experiment with new technologies, fostering faster innovation and leveraging AI to revolutionise areas like trade finance and lending.

A robust and adaptable infrastructure is paramount for the governanced data sharing needed to truly harness the power of these advancements, particularly in sophisticated data and AI initiatives.

Tapping into hidden data with a component approach to data products provides a value and a story for leveraging data where it resides.





Composable data services provide a balance of control, security and flexibility – so you can manage sensitive financial data while enabling the agility required for AI-driven innovation. Ultimately, by standing up a comprehensive framework underpinned by hybrid cloud, you can move beyond mere operational efficiency to directly align cloud platform benefits with measurable business value outcomes. And that ensures every technological investment translates into tangible strategic advantage.

# Conclusion

## Ensuring ongoing AI value with composable data services



By David Coles  
UK Cluster CTO

**In this whitepaper, we've discussed how composable data services help FSI organisations unlock value from AI, with a range of perspectives and best practices. And there's substantial value from doing this right: For every \$1 organisations invest in generative AI, they're realizing an average of \$3.70 in return.**


**I'll leave you with a final thought – on ensuring your approach delivers the sustained value you need.**



In this, we can learn from the cloud hype cycle. There's no doubt that cloud has accelerated experimental innovation and enabled a larger scale of ideations from proof-of-concept to production. That said, I'm equally aware of technology trends – and the mass of cloud adopters who feel they're yet to overturn the stones of true-value ROI. Over 70% of cloud customers report they're not achieving the full value they expected from their cloud investments.

We don't want AI investments to fall into that same trap.

The shortfall in cloud value often isn't due to technology limitations; rather it's rooted in the absence of a structured, outcome-driven delivery approach. And you can fix that with your data and AI approach.



## **An outcome-driven delivery model**

In this, we can learn from a cloud value engineering model. This model embeds a core T-shaped team as a strategic service layer – one that combines deep domain expertise with broad cross-functional knowledge across data, platforms and product delivery. This core team forms the engine of repeatable and scalable value creation, enabling rapid onboarding, shared governance and consistent technical patterns.

Then, surrounding this core platform team, you have flexible specialists like ML engineers, application architects and data analysts who are integrated on demand, tailoring each delivery to product complexity, data maturity and business priority.



## Measurable, repeatable value

These cross-functional T-shaped teams can reduce rework by 25-35% and improve time to value by 30%. Organisations with federated, repeatable delivery models are 2.4x more likely to scale data-driven outcomes across business units.

As a result, you're not just adopting technologies, you're engineering an approach that delivers measurable, repeatable value.

**2.4x**

More likely to scale data-driven outcomes across the business

**30%**

Improvement in time to value

**25-35%**

Reduction in rework





**Book a free workshop to review  
the suitability of composable  
data services for your financial  
data estate**

**[Get in touch](#)**



# Sources

<https://blogs.microsoft.com/blog/2025/04/22/>  
<https-blogs-microsoft-com-blog-2024-11-12-how-real-world-businesses-are-transforming-with-ai/>  
<https://blogs.microsoft.com/blog/2024/11/12/idcs-2024-ai-opportunity-study-top-five-ai-trends-to-watch/>  
<https://www.precisely.com/blog/data-integrity/2025-planning-insights-data-quality-remains-the-top-data-integrity-challenges>  
<https://www.apexon.com/blog/the-rise-of-predictive-banking-how-financial-institutions-can-stay-ahead-in-2025/>  
[https://newsroom.ibm.com/2025-02-05-ibm-study\\_gen-ai-will-elevate-financial-performance-of-banks-in-2025](https://newsroom.ibm.com/2025-02-05-ibm-study_gen-ai-will-elevate-financial-performance-of-banks-in-2025)  
<https://developer.ibm.com/articles/data-mesh-a-peek-into-this-new-paradigm/>  
<https://www.gartner.com/en/articles/ai-ready-data>  
<https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>  
<https://www.grandviewresearch.com/horizon/outlook/ai-agents-in-financial-services-market/uk>  
<https://www.grandviewresearch.com/horizon/outlook/generative-ai-in-financial-services-market/uk>  
<https://www.cultureworkshr.com/the-disconnect-dilemma-bridging-the-gap-between-executives-and-frontline-employees/>  
<https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/transforming-it-infrastructure-organizations-using-agile>  
<https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/all-about-teams-a-new-approach-to-organizational-transformation>  
<https://blogs.psico-smart.com/blog-how-does-crossfunctional-collaboration-improve-problemsolving-and-overall-organizational-efficiency-180911>  
<https://www.snaplogic.com/blog/ai-ready-data-ai-agents-gartner-summit-insights>  
<https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-missing-data-link-five-practical-lessons-to-scale-your-data-products>