MLOps and trustworthy AI for data leaders

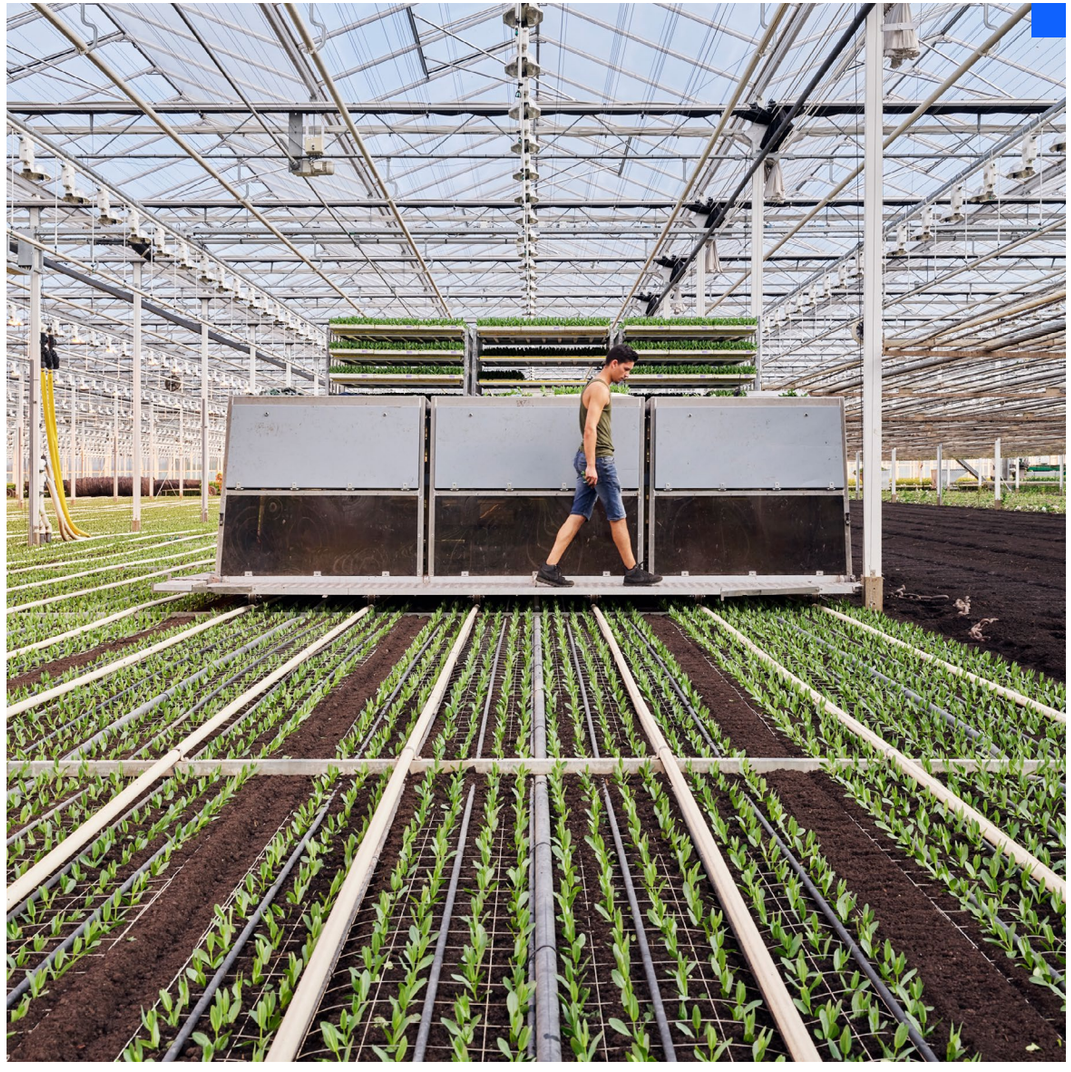

Contents

01
Introduction:
IBM Data & AI on AWS

02
Introduction—a data fabric approach to MLOps and trustworthy AI

03
Why establish MLOps and trustworthy AI?

04
The building blocks of MLOps and trustworthy AI

05
Data fabric—
a holistic approach

06
MLOps and trustworthy AI success stories

07
Consider these components

08
How it works

09
Create your ideal MLOps and trustworthy AI solution

01
Introduction—IBM Data & AI on AWS

IBM and AWS are bringing together the #1 leading AI portfolio with the largest global cloud infrastructure, to help businesses find a better way to put their data to work.

In this guide, we'll look at the most common governance and privacy challenges modern organizations face, how IBM Cloud Pak for Data on AWS can help you create an effective solution and approach to building an automated, integrated data governance and privacy layer across all the data in your enterprise.

With IBM Cloud Pak for Data on AWS businesses can:

Create a robust hybrid cloud data architecture
With access to multiple data sources to choose from, including IBM & AWS, with IBM Cloud Pak for Data your business can create a robust hybrid cloud data architecture.

Connect all data
When using IBM Data and AI on AWS, you can connect to all your data sources, including Amazon S3, Amazon Redshift, Amazon RDS, Amazon Aurora, Snowflake, MongoDB, Teradata, Apache Hive, IBM DB2, Netezza performance server and more.

Additionally, data from hybrid cloud databases & services can be connected using inbuilt connectors available in Cloud Pak for Data.

Use single interface
Create a single interface for integrating data sources coming from multiple IBM and AWS infrastructures.

Get started with ease
IBM Cloud Pak for Data on AWS is available with an AWS Quick Start deployment, which ensures that secure, comprehensive analytics and AI platform is ready within four hours.

02

A data fabric approach to MLOps and trustworthy AI

In an era where trust has gained tremendous importance to customers, every organization has a responsibility to adhere to ethical, explainable AI, respecting individual rights, privacy, and non-discrimination.

When trust in AI is established, revenues and customer satisfaction increase, time to market shrinks and competitive positioning improves. However, when this is overlooked, trust for the organization can erode and result in failed audits, and regulatory fines resulting in the loss of brand reputation and revenues.

Success in building, deploying, and managing AI/ML models is based on trusted data and automated data science tools and processes—and this requires a technology platform that can orchestrate data of many types and sources within hybrid multicloud environments. Data fabric is a technology architecture approach that helps ensure quality data can be accessed by the right people at the right time no matter where it resides.

Data fabric provides a strong foundation for MLOps and trustworthy AI. Keep reading to get the full story or try it for yourself with our free MLOps and trustworthy AI trial.

In addition, data fabric enables multicloud data integration, data governance and compliance, and 360-degree customer intelligence (these use cases are covered in other ebooks.)

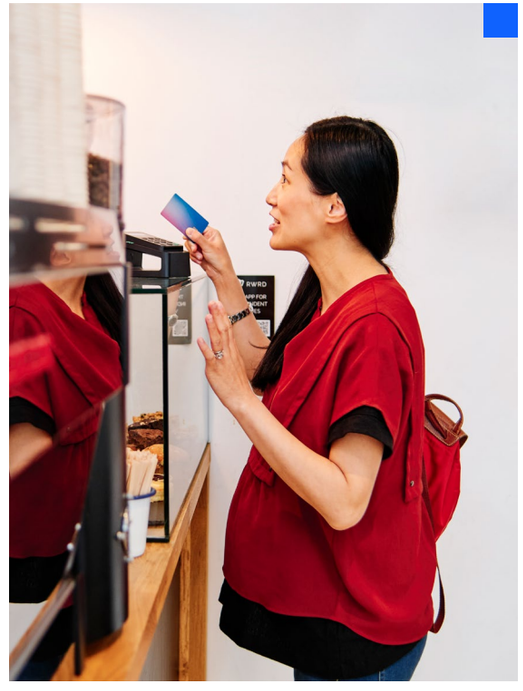

03
Why establish MLOps and trustworthy AI?

Well-planned, executed and controlled AI that is built to mitigate risks and drive desired analytic outcomes requires building trust in the data, models, and processes.

Trust in data
Strength and trust in AI outcomes require the connection to data that is accurate, of high quality, and ready for self-service consumption by the right stakeholder(s). AI model strength depends on the ability to aggregate both structured and unstructured data from disparate internal and external sources, from on-premises, public, or private clouds. Successful data collection and usage hinges on fairness in training data, tracking lineage and ensuring data privacy when offering self-service analysis by multiple personas.

Trust in models
To ensure transparency and accountability at each stage of the model lifecycle, MLOps-automated, integrated data science tools help to operationalize the building, deploying and monitoring of AI models. MLOps increases the efficiency for continuous integration, delivery and deployment of workflows to help mitigate bias, risk, and drift for more accurate data driven decisions. Some unique MLOps implementations also infuse the AI model process with fairness, explainability and robustness.

Trust in processes
Across the model lifecycle, from model building, to deploying, managing and monitoring, lack of automated processes can lead to inconsistency, inefficiency and lack of transparency. AI Governance provides automation that drives consistent, repeatable processes to decrease time to production, increase model transparency, ensure traceability and drive AI at scale.

IDC predicts that by 2025, 60% of enterprises will have operationalized their ML workflows through MLOps/ModelOps capabilities and AI-infused their IT infrastructure operations through AIOps capabilities.[1]

04
The building blocks of MLOps and trustworthy AI

#1

Data collection and access
Models and algorithms are only as good as the data that is used to create them. Data may be incomplete, or biased, which can lead to inaccurate algorithms and analytic outcomes. A method of tying these sources together is vital.

Data scientists, analysts and developers need self-service data access to the most appropriate data for their project. Setting up privacy controls for multiple personas, providing real-time access to “right” data users, and tracking data lineage can be challenging without the right tools.

#2

Model building and deploying
Data cleaning and prepping can be a very manual process, the most time consuming and least enjoyable task for the data scientists. Most would prefer to spend their time mining or modeling data. A solution that helps to automate these tasks is essential.

A common complaint among data scientists is the lack of integrated tools for building, deploying, scaling and training models. The tedious task of using multiple standalone tools is often coupled with the lack of software documentation, FAQs and use cases for those tools. The lack of integrated tools can lead to governance and compliance issues. Once a model is built, collaboration between the model builder and other teams including software engineering, AI operations and business analysts is imperative. Building an automated structure to do this in a timely manner is key.

#3

Monitoring and managing models
Once deployed, it is imperative to continuously measure and monitor model performance. Changes can be caused by model degradation, drift, bias and more. Using manual processes to detect these changes can lead to costly inaccuracies and lack of governance resulting in regulatory consequences and customer mistrust.

Automating the process of monitoring and retraining models as well as collecting model facts across the lifecycle drives consistency and transparency, automatically addressing model fairness and regulatory requirements.

TechTarget reported that 83% of organizations have increased their AI budgets, and the average number of data scientists employed has grown by 76%. Yet the time required to deploy a model is also going up, with 64% of organizations taking a month or longer.[2]

05

Data fabric—a holistic approach

In the previous sections we looked at what it takes to build trust in data, models and processes, as well as the challenges many organizations face without the right level of automation at each stage of the AI lifecycle. We outlined the specific technological building blocks you need at each stage to establish well-planned, executed and controlled AI.

These capabilities include:

- A way to integrate data of many types and sources across diverse deployments
- Self-service access with privacy controls and a way to track lineage
- Automated model building, deployment, scaling, training and monitoring
- Automated governance to help ensure data quality and regulatory compliance

To put all of these building blocks in place across the enterprise and overcome data complexity challenges, organizations need an integrated data strategy and architecture. Enter the data fabric.

A data fabric is an architectural approach that simplifies data access in an organization to facilitate self-service data consumption. It brings together capabilities like those listed previously as part of a unified architecture, avoiding the cost and complexity of integrating a plethora of point solutions. Instead of a fragmented group of products that have been stitched together, a data fabric offers a single, holistic solution that is built to work seamlessly.

A data fabric connects, governs and protects your siloed data distributed across a hybrid cloud landscape and can help you bring the promise of your data strategy to life.

Gartner predicts that by 2023, organizations using data fabrics will dynamically connect, optimize and automate data management processes and reduce time to integrated data delivery by 30%.[3]

In the next section you'll see how other organizations have successfully implemented a data fabric approach to MLOps and trustworthy AI.

06

MLOps and trustworthy AI success stories

Healthcare: Penn Medicine

Whether it's through speeding up the pace of research, helping clinicians to make better decisions, or improving patient experience through personalized medicine, AI has the potential to reshape the healthcare industry.

Penn Medicine is a world-renowned academic medical center in Philadelphia, with hospitals ranked among the nation's top hospitals. They needed a better way to predict intracranial hypertension (ICP crises) in patients with severe traumatic brain injuries, to facilitate earlier identification and more proactive treatment of the condition.

Neurosurgeons at Penn Medicine rely on a variety of physiological timeseries data to make treatment decisions but monitoring for ICP crises is a manual and labor-intensive process.

IBM's Data Science and AI Elite team developed a machine learning model to predict increases in intracranial hypertension (ICP crises) in advance with a lead time of 20-30 minutes, providing enough time for neurosurgeons to intervene and potentially avoid crises. The data-driven solution leveraged machine learning by uncovering hidden patterns within the physiological time series data. The model could provide continuous prediction in regular intervals, and if implemented in clinical workflows, could potentially improve patient outcomes after severe TBI.

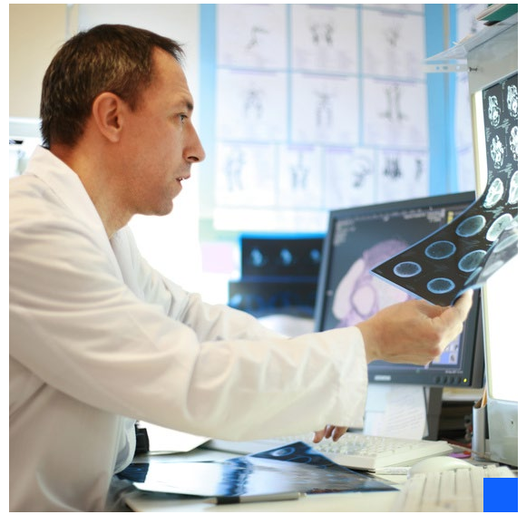

06

MLOps and trustworthy AI success stories

Manufacturing: ABB

Unpredictable customer demand, historic supply chain inefficiencies and labor shortages are driving the need for organizations to reevaluate manufacturing operations. Acceleration of digital technology using AI has the promise of driving operational efficiencies at scale, increasing productivity, and improving supply chain management. ABB is a multinational corporation headquartered in Switzerland that operates in robotics, power, electrical equipment and automation technology. They needed to leverage their pipe specification historic data to apply advanced analytics to not only preserve and maintain, but also expand and share the knowledge of their retiring sales engineers.

The company needed an end-to-end process that could:

- Unlock their data from inadequate data formatting
- Use machine learning for analytics and
- Use the output to make a real impact on their business processes

An end-to-end pipeline was created and offered insights on advanced analytical techniques such as association mining and unsupervised machine learning to uncover hidden patterns.

Implementing AI technologies helps these traditional institutions shift to digital channels that offer mobile and self-service options. By encapsulating the results in a prototype web application, ABB was able to envision the possible development of an internal application that could be used by its sales engineers.

As a result, ABB was able to:

- Reduce complexity and provided transparent data infrastructure
- Reduce additional costs and resources to train new sales engineers with a new data-driven approach
- Potentially save time for the creation of Pipe Specification from months to days
- Run advanced analytics on historic Pipe Specification data for the first time

06

MLOps and trustworthy AI success stories

Financial services: NatWest Group

Today’s customers expect exceptional, seamless omnichannel experiences. To create a lasting advantage and retain trust, retail banks, credit unions and other financial institutions need to gain insight into their customer’s current and future needs and how to meet them.

For many, owning a home is a dream come true. A bank-issued mortgage is oftentimes an essential ingredient in realizing that dream. But as regulations, products and processes evolve, issuing and obtaining a mortgage can be complicated. Banks are tasked with being able to access and apply accurate policy information to each customer’s unique needs in real time, throughout the home buying process.

Royal Bank of Scotland (RBS), now Natwest Group, is solving that problem with a plan for digital mortgage support. It created an AI-powered, cloud-based platform that empowers mortgage call center employees with real-time digital mortgage support for home buyers. The platform is called “Marge”—and is a valuable new asset to RBS’ digital transformation.

She is intentionally personified as a member of the RBS team and even has her own evolving personality. Marge was built directly on the cloud, embedded in RBS’ existing data structures, with access to new data added every minute via content updates and customer interactions. For RBS mortgage call center employees, she is one single point of access for digital mortgage support for their cognitive enterprise.

While on the phone with a customer, RBS employees get quick digital mortgage support from Marge by typing key words into a console. With cognitive enterprise technology at their fingertips, they can support new and existing home buying customers.

Since implementing the digital mortgage support tool, RBS has seen a 20% improvement in customer NPS, and a 10% decrease in call duration. As Marge evolves, Royal Bank of Scotland employees will be empowered to do great work during the digital transformation.

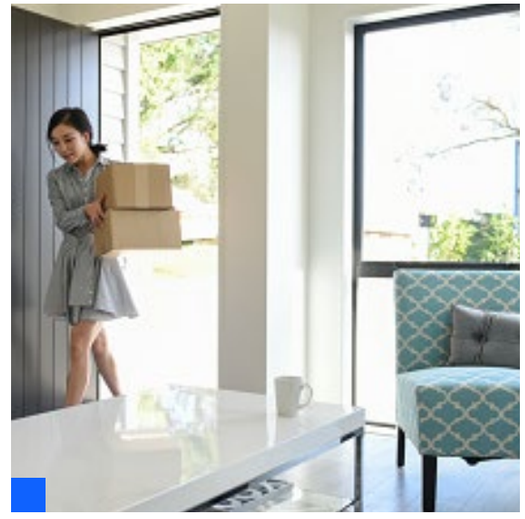

“In the mortgage industry, change is constant. Regulation changes, product changes, process changes. It’s imperative that the customer has the support and the info they need to allow them to focus on their home buying journey.”

MaryAnn Fleming
Head of Homebuying Services
Royal Bank of Scotland

07
Consider these components

IBM offers an end-to-end data science platform that provides user self-service access and integrated, supported and automated tools for model building and continuous delivery of models to production.

This set of capabilities:

- Unifies data and AI capabilities together on an integrated platform
- Increases data integration by automating engineering tasks
- Automates governance, data protection and security enabled by active metadata
- Provides collaborative tools
- Provides purpose-built AI model risk management

IBM's data fabric approach is built on a cloud-native data platform embedded with AI to help improve productivity and reduce complexity by connecting siloed data of all types, sources and workloads across a hybrid cloud environment for self-service consumption and AI models.

IBM Cloud Pak for Data
IBM Cloud Pak® for Data is a platform built specifically with a data fabric architecture in mind to predict outcomes faster and allow you to collect, organize and analyze your data, no matter where it may reside. The platform thus helps to improve productivity and reduce complexity by building a data fabric that connects siloed data distributed across a hybrid cloud landscape.

Make technology operations more agile and impactful while reducing the costs and risks of bringing AI and DevOps operations together. With its extensive ecosystem, IBM Cloud Pak for Data helps you extend your open and third-party investments and tap into IBM's latest AI innovations. As part of building your AI engineering, you can modernize your information architecture to be future-ready.

Learn more about
IBM Cloud Pak for Data →

Get started with
Cloud Pak for Data on AWS →

IBM Watson Studio
IBM Watson® Studio includes tools to explore data, build models (either visually or with code), deploy and monitor these models with end-to-end lifecycle explainability and fairness. Watson Studio uses MLOps to simplify model production from any tool, provides automatic model retraining, drives transparency and monitors models over time for accuracy and bias.

As an integral feature of Watson Studio, IBM Factsheets tracks how each model is developed and deployed. Increased transparency provides information for AI consumers to better understand how a model or service was created to help determine if is appropriate to use for a specific situation or need.

Learn more about
Factsheets →

Learn more about
IBM Watson Studio →

Get started with Watson Studio
on AWS with Cloud Pak for Data →

07
Consider these components

IBM Watson Knowledge Catalog
IBM Watson® Knowledge Catalog provides intelligent cataloging, with automated metadata collection and policy management to ensure the details of a model are automatically collected and stored for maximum transparency and repeatability. It ensures that models are impartial, address bias, are explainable and adapt to changing model parameters.

Learn more about
IBM Watson Knowledge Catalog →

Get started with
Watson Knowledge Catalogue on AWS
with Cloud Pak for Data →

IBM OpenPages with Watson
IBM OpenPages with Watson helps you identify, manage, monitor and run reports on risk and regulatory compliance across the entire model development lifecycle, with no changes for the requirements to the current AI/ML tools used today.

Learn more about
IBM OpenPages →

08
How it works

Scale trustworthy AI

With IBM Cloud Pak for Data, IBM Watson Studio enables organizations to synchronize application and model pipelines, and monitor Amazon SageMaker models for fairness, drift, and explainability.

Monitor and govern AI models to meet regulations, manage risk and enhance transparency.

IBM® Cloud Pak for Data lets you connect to your data no matter where it lives. It supports a variety of data sources such as Amazon S3, Amazon Redshift, Amazon RDS, Amazon Aurora, Snowflake, MongoDB, Teradata, Apache Hive, IBM DB2, Netezza performance server and more.

Enable your organization to:

- Maintain a data collection with self-service analytics
- Ensure quality, fairness in training, data lineage and privacy
- Automate model building, deployment and monitoring
- Drive transparency and accountability through the AI lifecycle

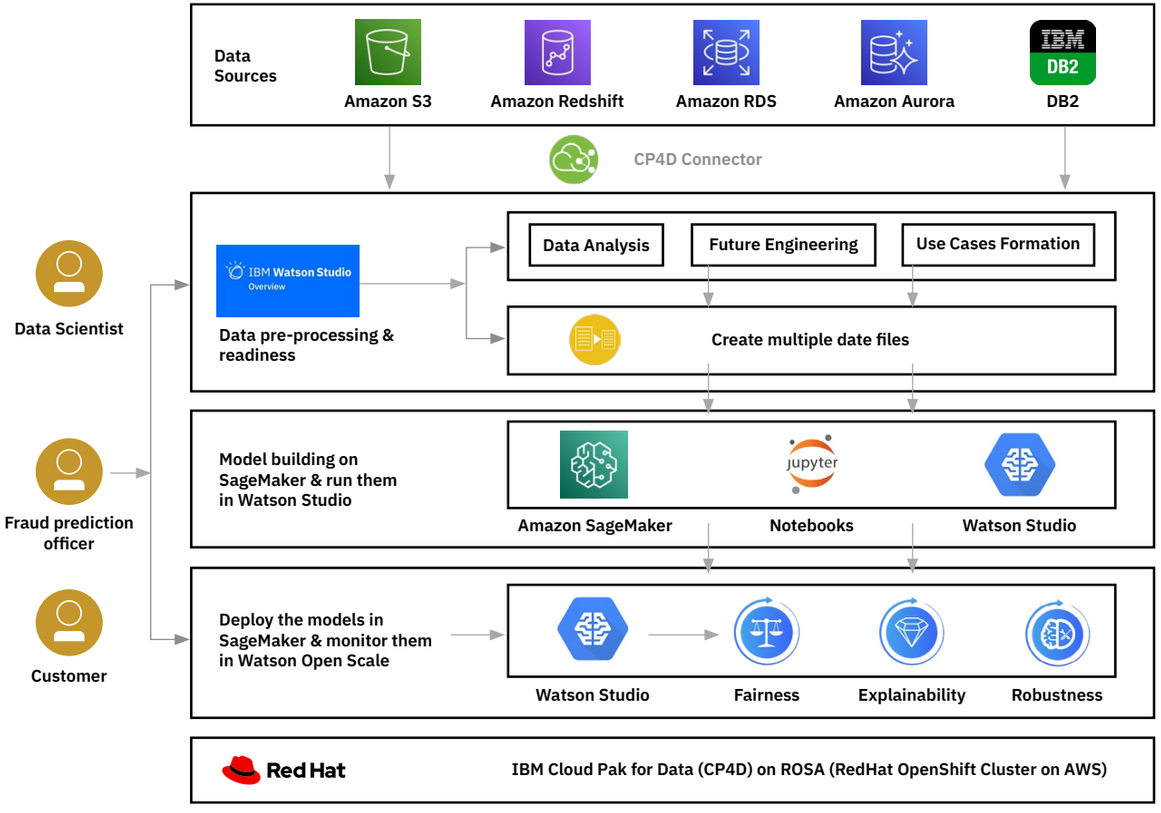

09

Create your ideal MLOps and trustworthy AI solution

If you're ready to embrace a unified strategy and architecture approach to improve the accessibility, security and compliance of your data of all types and sources, we encourage you to take advantage of a few resources.

Get started on AWS Marketplace.
Schedule a consultation with a data fabric on AWS expert.

Check out this data fabric use case ebook:

Data Governance and Privacy

IBM Corporation
Route 100
Somers, NY 10589

Produced in the United States of America
May 2022

01 IDC Future Scape: Worldwide Artificial Intelligence and Automation 2022 Predictions—European Implications IDC, Jan 14, 2022, https://www.idc.com/getdoc.jsp?containerId=EUR148675522

02 TechTarget, “Addressing 3 infrastructure issues that challenge AI adoption,” May 14, 2021

03 Magic Quadrant for Data Integration Tools, Gartner, 9 March 2021.