



eBook

The Definitive Guide to Success in Modern Data Engineering

By Sumeet Kumar Agrawal, VP, Product Management,
Informatica; and Valentin Moskovich, VP, Product Development,
Advanced Data Engineering, Informatica

**Where data
& AI come to** 



Contents

| | | | |
|---|-----------|---|-----------|
| Introduction | 4 | Artificial Intelligence, Machine Learning and Data Science | 32 |
| Evolution of the Data Landscape | 5 | Data Science Use Cases Supported by Data Engineering | 38 |
| The 4 Fundamentals of Data Engineering | 8 | Operationalization in Data Engineering | 43 |
| - Data Discovery and Lineage | 9 | - Business Understanding | 45 |
| - Data Replication and Ingestion | 11 | - Data Acquisition | 46 |
| - Data Processing | 12 | - Model Development | 47 |
| - Data Quality | 14 | - Model Deployment | 48 |
| | | - Model Monitoring | 49 |
| How to Become a Successful Data Engineer | 16 | | |
| Top 5 Data Engineering Trends | 20 | AI, ML and Data Science in Action | |
| - Data Mesh | 20 | - Banco ABC Brasil: Accelerating the Credit Approval Process by 70% With Enhanced Analytics | 54 |
| - Data Fabric | 23 | - Maersk: Using Data to Improve Operations and Maintain a Competitive Edge | 55 |
| - Modern Data Stack | 25 | | |
| - Data Management Platforms | 27 | How Informatica Intelligent Data Management Cloud™ (IDMC) Supports Modern Data Engineering | 56 |
| - Data Observability | 29 | About Informatica | 58 |
| Data Engineering Trends in Action | | | |
| - MC: Transforming Complex Technology into Extraordinary Business Performance | 30 | | |
| - Takeda: Delivering Breakthrough Therapies Faster With One Data Management Cloud | 31 | | |

Acknowledgements

Special thanks to the following Informatica contributors:

Special thanks to the following Informatica contributors:

- Peter Manta, Chief Architect, North America
- Abhilash Mula, Senior Manager, Product Management and R&D
- Rudra Ray, Director, Technical Marketing
- Siddharth Rajagopal, Chief Architect, EMEA
- Michelle Schallhorn, Senior Principal Content Producer

Disclaimers

The information in this eBook is made as of May 2023 and is subject to change without notice. The information provided is for informational purposes only and is based on the authors' knowledge, experience and opinions. The methods described in this eBook are not intended to be a definitive set of instructions. You may discover other methods and materials to accomplish the same result. Your results may differ.

Informatica LLC provides the information in this eBook "as is" without warranty of any kind, express or implied, including without any warranties of merchantability, fitness for a particular purpose and any warranty or condition of non-infringement.

This eBook includes information regarding the products and services of third parties. We do not assume responsibility for any third-party materials or opinions. Use of mentioned third-party materials does not guarantee your results will mirror those mentioned in the eBook.

Introduction

It's no secret — data engineering is always evolving and demand for the most advanced skills are growing rapidly. And it's on you to keep up with the latest innovations, trends and patterns. There's no doubt that a career in data engineering is rewarding and amazing. And at times it can also be overwhelming, demanding and stressful, which can inadvertently introduce a very common challenge — technical debt.

The bottom line? Your organization depends on you to extract insights that will ultimately inform critical business decisions. Decisions that can make the difference between success and mediocrity.

In today's competitive environment, successful data engineers must have the right tools, skills and knowledge in place to do their job well. In your role, you help preserve the credibility of business intelligence and advanced analytics and help support artificial intelligence (AI) and machine learning (ML) models. Data engineers who take advantage of modern technologies and architectures will not only provide more value; they will help elevate their organization in a crowded market.

This must-read eBook is designed for both early-stage professionals and veterans who have seen (and done) it all. It provides the industry knowledge and practical advice you need to stay relevant in a competitive landscape. And it will help you sharpen your skills for long-term success — personally and professionally.

You'll learn about:

- Key fundamentals and skills that can set you up for long-term success
- Data engineering best practices to enable you to build scalable, cost-performant, end-to-end data pipelines
- The latest data engineering trends, like data mesh and data fabric, to master and incorporate into your IT strategy
- How data engineering supports AI, ML and data science to help your company deliver better customer experiences and make more-informed decisions

If you're a data engineer, there's a lot to learn so let's get started.

Evolution of the Data Landscape

The evolution of the data landscape plays a critical role in understanding modern data engineering. Setting the DeLorean to travel back to the 1970s and '80s (for you “Back to the Future” movie fans), we would see mainframes and midrange machines storing most enterprise data. As time progressed into the 1990s, much of this shifted into distributed applications like ERP, SCM, CRM and other systems. Each of these applications was designed to store its own data and presented different access challenges.

Moving into the 2000s, as illustrated in Figure 1, there were on-premises **data marts** and **data warehouses**. There was furious debate about the best way to structure the warehouses — Kimball or Inmon. Regardless of the right answer, there were some common truths. Compute and storage were expensive!

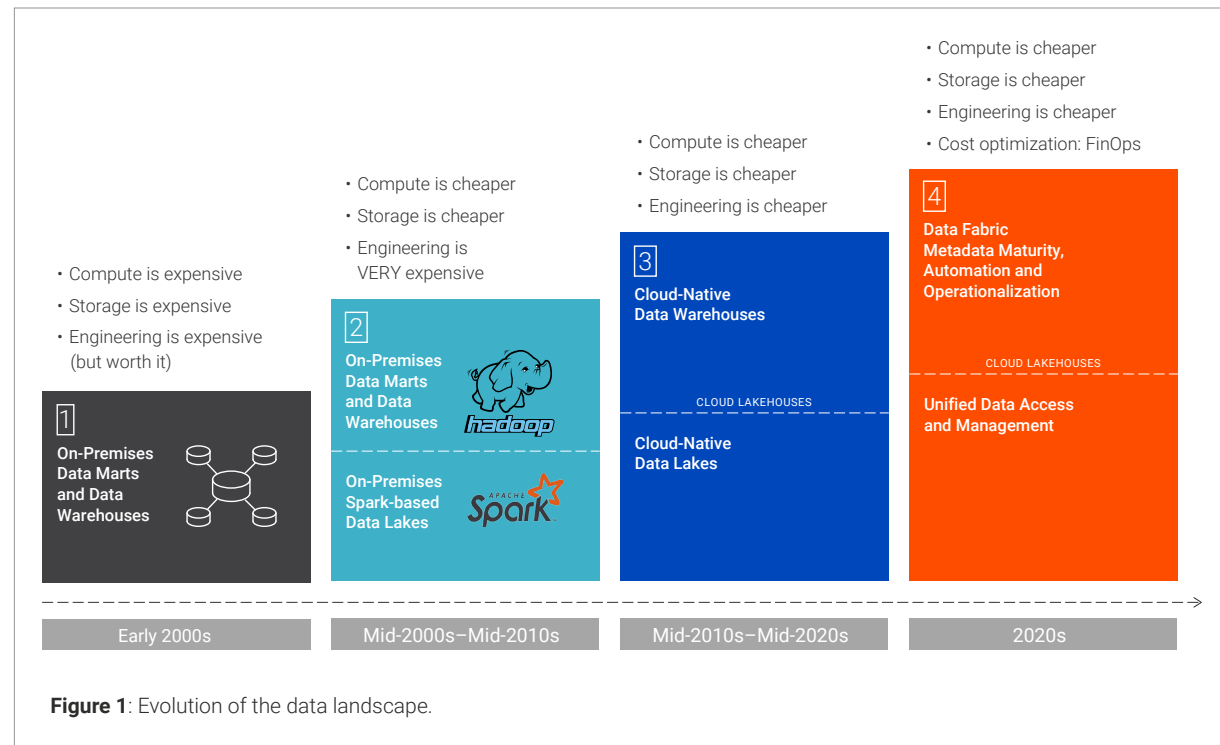


Figure 1: Evolution of the data landscape.

Evolution of the Data Landscape (continued)

But the value of the data warehouse was worth the expense. In fact, data warehouses delivered so much value the world moved towards purpose-built data warehouse appliances, and those were very expensive. So expensive that data modelers and data engineers were tasked with optimizing the systems and reducing the operational costs. Those data modelers and engineers were also expensive, but again, they were worth it.

Around 2006, Hadoop came along, and it looked like **Big Data** was going to take over the world. As we know, that didn't quite pan out, but even so, Hadoop had a massive impact on **data management**: The notion that compute and storage are expensive got flipped on its head. Storage and compute, relatively speaking, now became cheap.

More importantly, Hadoop made it OK to say, "Throw more horsepower at it." Prior to Hadoop, uttering those words was very likely a career limiting move. If you had Teradata, Netezza, Exadata, etc. and you said to your manager: "Performance is a little slow. Let's buy more processing..." Well, that was very likely a seven- or even eight-figure suggestion and data engineers would get shown the door.

Hadoop Map-Reduce wasn't well suited for data management, so Spark soon made its way into the architecture. With strong in-memory processing, it opened the door to near real-time analytics and efficient big data management.

Although compute and storage were inexpensive, Hadoop was painfully complex, and engineers who really knew how to make it sing were extremely hard to find. And if you were able to find skilled Hadoop engineers, they were extremely pricey.

Technology again evolved and here we are today rushing to the cloud. And why not? With cloud, storage is cheap and compute is cheap because it is consumption based. And the final nail in Hadoop's coffin? Engineering is handled by the cloud ecosystems.

This all leaves us in a state of mixed architectures. Most organizations today are in some stage of architecture modernization. Whether it's cloud **data lake** and warehouse, **data fabric**, **data mesh**, implementing a data science practice or something else, there are significant challenges to overcome. Mainframes, distributed applications, relational databases, cloud ecosystems, batch, change and real-time latencies, evolving technology and expense profiles and self-service data management — all while still delivering business requirements and meeting service level agreements (SLAs)? Who can successfully pull all these things together? It's the data engineer, of course.

The Role of the Data Engineer

It is estimated that there will be around 200 zettabytes of data by 2025, with 100 zettabytes of them stored in the cloud.¹ Storing zettabytes of data is challenging on its own, but it can be even more difficult to gain value from such a huge amount of information. The data that's collected will have security and governance requirements that are mandatory to protect. Poor data quality causes misinformed business decisions, which can lead to pricy mistakes. The data that is collected not only needs to be secure, but it must also be clean and consistent. This is where data engineering comes into play.

The role of the data engineering team is to take data in a raw and unusable format and transform it into a clean state so it can be used by business leaders and data science teams for forecasting and predicting. Data engineers work in the background to help answer a specific question. The more data the company processes, the more time is spent on analyzing it.

Data engineers design and implement the architectures necessary for data scientists to be successful. After all, without clean, trusted data, what's the point of running analytics? You can see why data engineers are viewed as the backbone of any data team.

¹ How Much Data Is Created Every Day in 2022? [NEW Stats] (earthweb.com)

² <https://cloudacademy.com/course/intro-data-engineer-role-1123/data-engineer-vs-database-scientist/>

Data Engineering Versus Data Science

The amount of time spent on **data preparation** versus data analytics is disproportionate. Based on our experience, less than 20% of time is spent analyzing data, while over 80% is spent collectively on searching for, preparing and governing the appropriate data.

And the responsibilities are clear: Data engineers build and maintain the systems that store and organize data. Data scientists analyze data to predict trends and answer questions that help make meaningful business decisions. The data engineer engineers the data for the scientist to work on — a little bit like a lab technician and a scientist.²

The below pyramid illustrates how data engineering assists in data science operations.

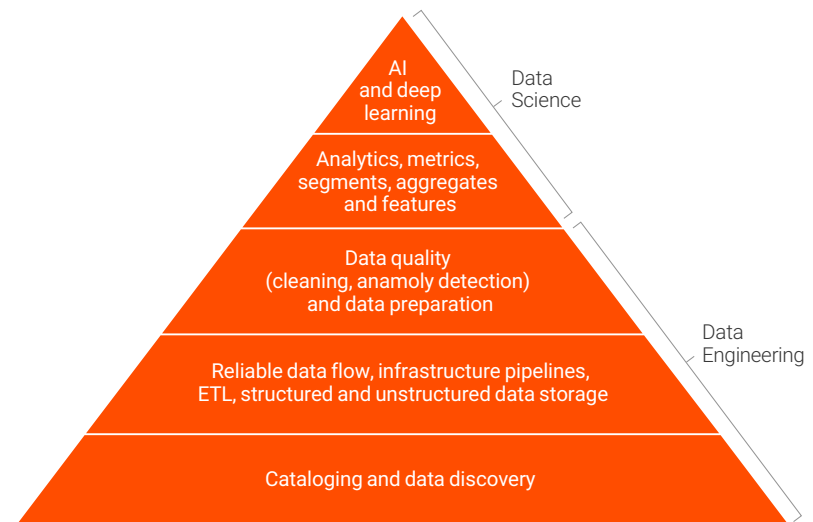


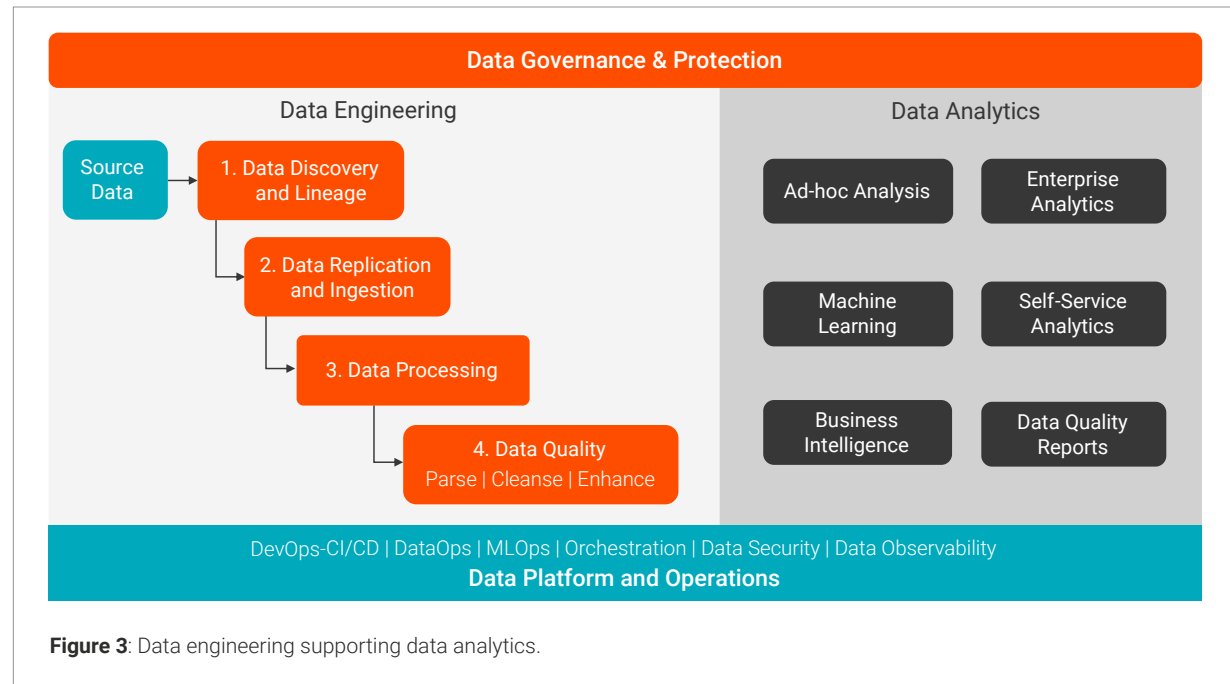
Figure 2: How data engineering supports data science projects.

The 4 Fundamentals of Data Engineering

While the data landscape and its impact on data engineers are constantly changing, the core fundamentals of the data engineering processes listed below have not changed much:

1. Data Discovery and Lineage
2. Data Replication and Ingestion
3. Data Processing
4. Data Quality

Over time, they have been enriched with the advancement of cloud storage, computing resource optimization and emerging data architectures. Figure 3 shows how these processes support data analytics and machine learning.



The 4 Fundamentals of Data Engineering (continued)

1. Data Discovery and Lineage

Before a data acquisition engineer starts building pipelines to clean the data for business needs, one of the fundamental steps in a data lake architecture is **data discovery**. This includes solving modern day data challenges, such as enterprise-wide data democratization, privacy and trust assessments for compliance, and greater data insights to ensure digital transformation success. Data discovery enables organizations to identify, catalog and classify business critical and sensitive data, so you can govern it for meaningful purposes with increased transparency.

The main steps in the data discovery process include:

- **Connect to data assets:** It is very important to have the right connectors to scan and profile data assets as part of data discovery. Manually gathering data assets and information takes longer, so out-of-the-box connectors to these data assets are critical in understanding existing data and anomalies.
- **Curate data preparation:** Once you establish connections to the required data assets, the next step is creating, organizing and maintaining data sets so they can be accessed and used by people looking for information. It involves collecting, structuring, indexing and cataloging data for users in an organization or group or for the general public.
- **Data discovery and lineage:** Data discovery begins with scanning for data across your organization's landscape. This can include on-premises or cloud-based sources from data warehouses to **extract, transform and load (ETL)** data and BI tools, SaaS applications and more.

Once you have located your data, you can discover further information about it, such as its structure, content and relationships. This information can be used to catalog data, enrich its metadata and provide context to help you understand what data exists, its source, its **lineage** and how it's related to other data.

The ability to map and verify how data has been accessed and changed is key to generating a detailed record of where specific data originated, how it changed and how it gets used. This is valuable for finding and fixing gaps in necessary data prior to data pipeline building. It is also important for responding to reporting requirements and audit requests for regulatory compliance. But how do you keep up with data lineage given the volume and scale of data today?

You need an AI-powered data lineage solution that includes a data catalog with advanced scanning and discovery capabilities to ensure you capture all the relevant metadata from your data sources.

The 4 Fundamentals of Data Engineering (continued)

The solution also needs to provide detailed, end-to-end data lineage across the cloud and on-premises.

• **Measure and optimize data value:** Unless we derive meaningful information for business decisions from an underlying dataset, the data itself is worthless. Hence it is very important to constantly measure and optimize data value through intuitive, real-time data insights. A recent poll we conducted with webinar participants indicates that the top four data asset-related information organizations are most interested in analyzing are:

- Data inventory
- Catalog adoption
- Data usage
- Data value

And it makes sense. Measuring data assets helps answer critical questions, such as:

- Who is using what asset, what feature and when?
- What is the catalog adoption rate?
- What are the most accessed assets?
- Who are the top contributors and collaborators?

Measuring and optimizing data value help you to define data as an asset. The definition of an asset is something owned or controlled, exchangeable for cash, and that generates a benefit.

Data catalog solutions such as an enterprise data catalog and cloud data governance help you measure and optimize data value by defining it as an asset. To learn more, [check out this data sheet on data asset analytics](#) and hear from industry experts in this [on-demand webinar](#).

The 4 Fundamentals of Data Engineering (continued)

2. Data Replication and Ingestion

Once data discovery is complete, the next step in data processing is to ingest the data into a data lake. Data engineers use **data replication** and **data ingestion** pipelines to handle the scale and complexity of business demands for data. Data replication and ingestion have also become key components of self-service platforms for analysts and data scientists to access data for real-time analytics, ML and AI workloads. Data ingestion is a process that extracts data from the source where it was created or originally stored and loads it into a destination or staging area. A simple data ingestion pipeline might apply one or more light transformations, enriching or filtering data before writing it to a set of destinations, a data store or a message queue.

There are many ways to replicate and ingest data. The organization's data strategy and business requirements and when the data is needed determine the appropriate data replication and ingestion method. Below are the three most common data replication and ingestion techniques:

- **Batch processing:** Collects data from sources incrementally and sends batches to the application or system where the data is to be used or stored.
- **Real-time processing:** The data is loaded as soon as it is recognized by the ingestion layer and is processed as an individual object.
- **Micro batching:** The data is divided into groups and ingested in smaller increments. This makes it more suitable for applications that require data in real time.

The 4 Fundamentals of Data Engineering (continued)

3. Data Processing

Data processing engines take data processing pipelines, abstract the business logic (either simple or complex) and process the data on frameworks such as Apache Spark, in a streaming or batch mode, on-premises or in the cloud.

More complex transformations such as joins, aggregates and sorts for specific analytics, applications and reporting systems can be done on the data inside the data lake using additional pipelines. From there, the final data is loaded into a cloud data warehouse for analytics. This process is typically done using **ETL** or **extract, load, transform (ELT)** techniques, which is often used for BI, analytics, reporting and AI/ML.

Constructing data pipelines with complex transformation and business logic is the core responsibility of data engineering. It requires advanced skills to design a program for continuous and automated data exchange. These data pipelines are commonly used for:

- Preparing data into a single location to streamline ML projects
- Integrating data from various Internet of Things (IoT) systems and connected devices
- Moving the prepared data into a cloud data warehouse
- Bringing data into the data warehouse to make informed business decisions

There are also different types of data processing techniques based on the source data:

- **Transactional processing** is deployed in mission-critical situations. These are situations which, if disrupted, will adversely affect business operations. An example is bank transactions.
- **Real-time processing** is like transaction processing, where output is expected in real time. However, this differs from transactional processing in terms of how data loss is handled. Real-time processing computes incoming data as quickly as possible. If it encounters an error in incoming data, it ignores the error and moves to the next chunk of incoming data. GPS-tracking applications are the most common example of real-time data processing.

The 4 Fundamentals of Data Engineering (continued)

- **Batch processing** is when chunks of data, stored over a period of time, are analyzed together or in batches. Batch processing is required when a large volume of data needs to be analyzed for detailed insights. For example, a company's sales figures over a period are typically processed in batches because there is a large volume of data involved, and the system will take time to process it.
- **Distributed processing** breaks down large datasets and stores them across multiple machines or servers. Distributed processing can be immensely cost-effective since businesses no longer need to build expensive computers and invest in their maintenance.

You need a single platform that provides the following data processing techniques and can:

- Support just about any integration patterns, replication, ingestion and ETL/ELT with thousands of metadata-aware connectors
- Support just about any data processing patterns, such as batch, real-time and near real-time through various engines
- Develop pipelines on day one with zero infrastructure footprint with serverless processing
- Achieve unlimited scale with auto scaling and auto tuning
- Ensure business continuity with high availability and tenant isolation
- Build once and run mapping in ETL or ELT mode
- Process data incrementally and efficiently as it arrives from files or streaming sources or databases or applications
- Infer schema and detecting column changes for any data source and file formats
- Cleanse data including data quality rules, data masking and data domain rules with out-of-order (OOO) transformations

"With Informatica, we achieved our key business objective of democratizing self-service access to the cloud data platform, which is responsible for delivering standardized, accurate, on-time data to various core business units."

Shashank Kulkarni

Data Engineer

Nutanix

The 4 Fundamentals of Data Engineering (continued)

4. Data Quality

Data quality is a critical component of any architecture. It is also unique because **data quality** is not the sole responsibility of data engineers. Done properly, data quality bridges the gap between business and IT. Business users understand the rules and context of data but do not usually have the technical ability to implement data quality rules in a production process. On the other hand, data engineers can implement rules but do not always have the business understanding necessary to create the rules.

For example, we may have a customer identifier formatted as XX999 — two characters followed by three integers. Creating a data quality rule to check for this format is straightforward for the data engineer, but rules are often more nuanced than simple format checks. What if those first two characters mean something specific and only the letters A, R and Q are allowed but Q is only used when the second integer is 7, in which case it indicates that customer has opted out of being contacted?

In the example above, the business user understands the rule to implement but does not know how to implement it enterprise-wide. This is where the business and IT must collaborate to design and implement the data quality rules. The keys to doing this properly are role-specific interfaces, shared metadata and loosely coupled architecture:

- **Role-specific interfaces:** The business user should have a thin client interface designed around ease of creating and testing rules. Rule creation should be graphically driven rather than requiring scripting knowledge. Rule testing must then be integrated into the data quality workflow. The data engineer's interface is more technical in nature with a full palette of both data quality and **data integration** functions.
- **Shared metadata:** Typical collaboration usually involves multiple rounds of back-and forth discussion where business and IT are not exactly on the same page. What typically happens is the business user tries to describe the requirements and the data engineer tries to explain the results. By sharing metadata, the business user and data engineer are always aligned because they both see the exact same thing, but through their role-specific lens.
- **Loosely coupled architecture:** Because data quality rules are represented as logical metadata, they should be loosely coupled from their sources and deployed via different mechanisms. Data quality is not specific to a database, SaaS application, mainframe or other sources. Data quality rules must also be easily applied to batch data, real-time data or exposed via an API to enable upstream, proactive data quality. Only a metadata-driven solution supports the same logic regardless of data source or deployment mechanism.

The 4 Fundamentals of Data Engineering (continued)

Data quality ensures your data is fit for purpose. To help you achieve this, you need a quality solution for multi-cloud and on-premises data with key features like:

- Discovery, search and profiling
- Data enrichment
- Role-based capabilities
- A rich set of transformations
- Reusable rules and accelerators
- Exception management
- AI-driven insights to automate data quality rules

Understanding these four data engineering processes — data discovery and lineage, data replication and ingestion, data processing and data quality — will help you accelerate your data engineering journey. To maximize these key fundamentals, data engineers should leverage out-of-the-box capabilities whenever possible and automate data pipelines using AI/ML- based engines to identify data changes faster.

“Thanks to Informatica [cloud solutions] and Google Cloud, we now have much more capacity for advanced analytics, giving us the insights we need to compete in the fast-changing solar power industry.”

Harish Ramachandraiah

Director, Engineering & Analytics
Sunrun

How To Become a Successful Data Engineer

We have covered the fundamentals of data engineering. Now let's look at what it takes to become a successful data engineer. From aspiring to early-stage data engineers to data and analytics leaders who plan to build data engineering teams, this chapter is for you!

Data engineering is a rapidly growing profession. From large public cloud companies to innovators, data engineers are in high demand. There are over 220,000 job listings for a data engineer in the U.S. on LinkedIn. In fact, data engineering is the fastest growing tech job, beating data science hands down, and the demand has only increased since 2020.³

The bottom line: Your skillset is in demand. According to The New York Times, U.S. unemployment rates for high-tech jobs range from slim to nonexistent. On average, each tech worker looking for a job is considering more than two employment offers.⁴

What Exactly Does a Data Engineer Do?

Data engineers enable data-driven decision making by acquiring/ingesting, transforming and publishing data. A data engineer discovers, designs, builds, operationalizes, secures and monitors data processing systems. They do this by focusing on security and compliance; scalability and efficiency; reliability and governance; and flexibility and portability. A data engineer is also responsible for leveraging, deploying and training pre-existing machine learning models.

³ <https://medium.com/codex/4-reasons-why-data-engineering-is-a-great-career-move-in-2022-3ef07b1e14f3>

⁴ The New York Times, OnTech with Shira Ovide, June 14, 2022

⁵ <https://cloudacademy.com/course/intro-data-engineer-role-1123/key-traits-of-a-data-engineer/>

"The expert in anything was once a beginner."

– Anonymous

Data engineers find trends or inconsistencies in data sets and develop algorithms to help make raw data more useful to the enterprise. Along with technical skills, data engineers also convey data trends, quality issues and patterns to help the business make meaningful use of the data it collects.

The role is very outcome orientated. A data engineer is a superhero of sorts because you can bring all this data to life.⁵

Some examples of data projects you may be involved in include:

- Analytical and visualization projects, which require knowledge of how to share data with data visualization and BI tools such as:
- Data aggregation
- Website monitoring
- Real-time data analytics
- Event data analysis

How To Become a Successful Data Engineer (continued)

- Data science and ML-focused projects, which require knowledge of how to train ML models continuously with the right set of cleansed data, etc. Some examples include:
 - Smart IoT infrastructure
 - Shipping and distribution demand forecasting
 - Virtual chatbots
 - Loan prediction

Data Engineer Qualifications: What You Need

Data engineers wear many hats throughout the various phases of the data lifecycle, so must have a diverse background that goes beyond education. While a degree in computer science, engineering, applied mathematics, statistics or related IT area is critical, here are key technical skills that every data engineer should have:

- Deep understanding of data management concepts focusing on data cataloging, data replication and ingestion, data integration (e.g. ETL, ELT) and data quality
- Experience in database management (relational/non-relational database management system concepts), data warehouse and data lake concepts
- Proficiency in scripting/coding languages such as SQL, R, Python, Java, etc.
- Cloud computing skills in one or more cloud service providers (e.g., Amazon Web Services, Microsoft Azure, Google Cloud, etc.)
- Basic understanding of machine learning algorithms, statistical models and some mathematical functions

*"A fundamental reality of the data engineer job market is that demand far exceeds supply. There are too many companies looking for too few data engineers."*⁶

⁶ <https://www.filtered.ai/blog/how-to-hire-a-data-engineer-fd>

How To Become a Successful Data Engineer (continued)

- Knowledge of data discovery and profiling through data cataloging and data quality tools

Intermediate and advanced data engineers should have the following skillsets, which can be learned by shadowing or creating your own test project by downloading test datasets:

- Emerging modern data architecture frameworks like data fabric, data mesh and modern data stack
- API and real-time streaming
- Data governance concepts such as data sharing, data access and data asset analytics
- Data compliance and security knowledge
- Operationalizing data processing, data observability and DataOps

Traits of a Successful Data Engineer

Being a great engineer goes beyond technical skills and advanced degrees. Having the right personality is just as important. A career in data engineering can be rewarding and amazing. It can also be overwhelming, demanding and stressful. Here are five key traits of a data engineer who is poised for success:

1. **Curiosity.** Data engineers must keep up with the latest trends surrounding technology, tools, datasets and its usage. Things change fast and you need to be able to quickly understand, evaluate and learn new tools. You should be eager to learn, grow and always ask, "Why?"
2. **Flexibility.** There is constant change in the data industry. Data engineers should be able to go with the flow and be comfortable with pivoting strategies, changing priorities and adjusting timelines.

"Engineering is easy. It's the people problems that are hard."

Bill Coughran

Former Google Senior Vice President
of Engineering

How To Become a Successful Data Engineer (continued)

3. **Problem-solver.** Data engineers are responsible for testing and maintaining the data architecture that they design and looking for ways to improve data processes. This requires a mind for creative problem solving and thinking outside of the box.
4. **Multi-tasker.** Not all data engineers come from a computer science or data science background. Other fields include IT, statistics, math and computer engineering. It helps to be well-versed in all facets of data and proficient in tools and technologies focused on automation.
5. **Strong communicator.** Data engineers are an integral part of the data team and must present and explain concepts to non-technical and technical stakeholders ranging from peers to executive leadership. The ability to confidently state your case will help break down silos across the organization and lead to better business decisions.

In addition to having the above traits, being proficient in several (up to 30) technologies and knowing what tool to use when is critical. You must have a strong sense of ownership. This is a job where you are literally given a framework to work with and expected to come up with the program to run.⁷

Now that we have covered what it takes to be a competitive, well-rounded data engineer, let's shift gears to the latest trends and how you can incorporate them into your IT strategy.

"Utilizing Informatica [cloud solutions] to integrate the warehouse information from an on-premises operation data store (ODS) into Google BigQuery, analytic execution time was driven down to minutes resulting in a highly, scalable easy-to-use analytics solution."

Data Engineer

Large Grocery Retailer

⁷ <https://cloudacademy.com/course/intro-data-engineer-role-1123/key-skills-of-a-data-engineer/>

Top 5 Data Engineering Trends

So you've learned the fundamentals of being a successful data engineer. Now it's time to look at the trends you need to master.

As the complexity and volume of data grows, data engineers need to invest a significant amount of time maintaining legacy data pipelines, which can take them away from focusing on new requirements and supporting AI/ML initiatives. In addition to that, 60% of infrastructure and operations (I&O) leaders will encounter public cloud cost overruns that negatively impact their on-premises budgets.⁸

To help deal with these challenges, data engineers must be aware of newer technology trends and patterns that can help simplify **data pipelines**, reduce costs and provide meaningful data for analytical use. Let's explore the top five **data engineering** trends we believe will impact the way a typical data engineer works today.

Let's explore the five most critical trends every modern data engineer should know and how they can influence the way you work:

1. **Data Mesh**
2. **Data Fabric**
3. **Modern Data Stack**
4. **Data Management Platforms**
5. **Data Observability**

1. Data Mesh

A domain-driven analytical data architecture treats data as a product and is owned by teams that most intimately know and consume the data.⁹ In a **data mesh** architecture, data warehouses and data lakes are treated as nodes by data domain on the mesh rather than the central point of the overall architecture.

As the need for real-time data-driven decisions becomes more critical, many organizations require data to have business domain-specific ownership. This can enable data product owners to better manage and decide how their data is used. It can also encourage teams to share data, as opposed to simply copying it, and provide better visibility into where specific data is being used across the enterprise.

⁸ <https://www.informationweek.com/cloud/10-cloud-strategies-to-avoid-cost-overruns>

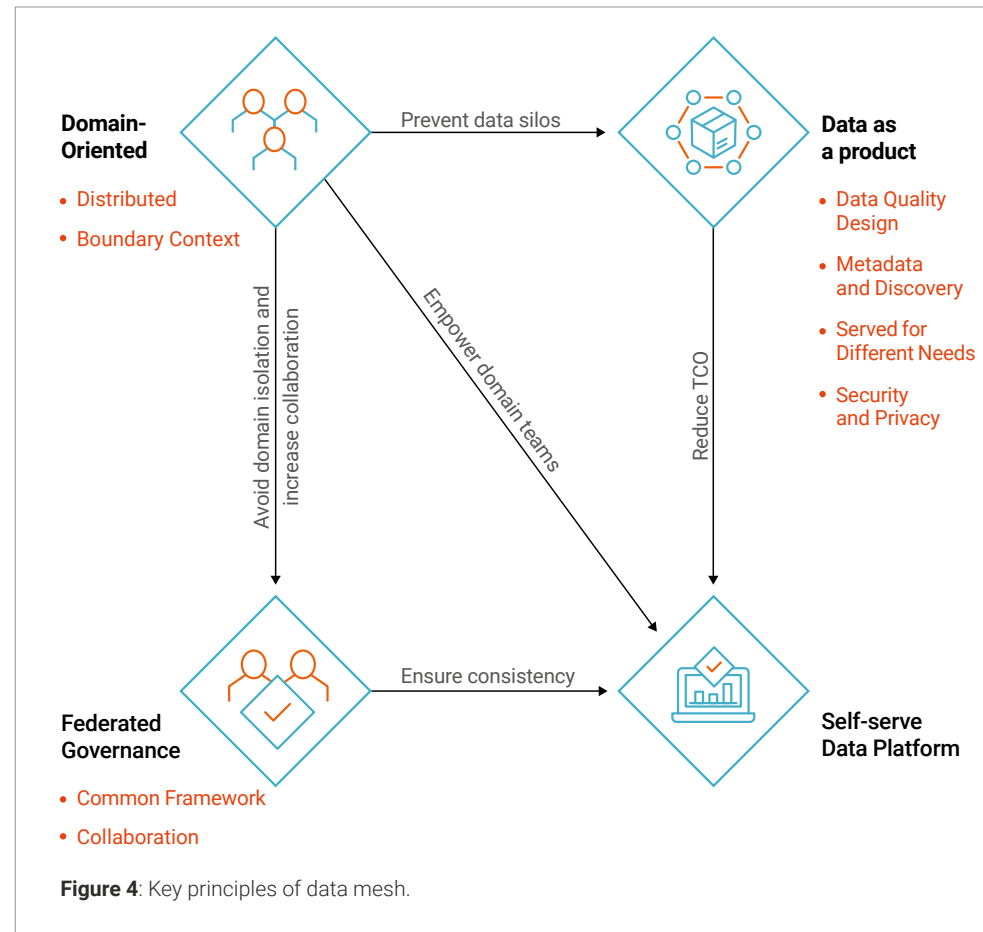
⁹ <https://www.thoughtworks.com/en-us/what-we-do/data-and-ai/data-mesh>

Top 5 Data Engineering Trends (continued)

As highlighted in Figure 4, the 4 key pillars of a data mesh include:

1. **Domain-oriented:** This is the ability to decentralize the ownership of sharing analytical data to business domains closest to the data, usually represented by either the source of the data or its main consumers.
2. **Data as a product:** There are differences between a “data product” and “data as a product.”¹⁰ A data product is what facilitates an end goal with data, whereas “data as a product” is the result of applying product thinking into datasets, and making sure they have a series of capabilities including discoverability, security, explorability, understandability, trustworthiness, etc.

“Data as a product” contains the code, its data and metadata, and the necessary infrastructure to run it. “Data as a product” is a subset of all possible data products and it belongs to the raw or derived data type of “data product.”



¹⁰ <https://towardsdatascience.com/data-as-a-product-vs-data-products-what-are-the-differences-b43ddb0f123>

Top 5 Data Engineering Trends (continued)

3. **Self-serve data infrastructure as a platform:** This allows the ability to enable a new generation of self-serve data infrastructure of choice to empower domain-oriented teams to manage the end-to-end data life cycle (from data acquisition to data democratization) of their data products
4. **Federated computational governance:** Federated data governance ensures data is discoverable, accessible, secure, trusted and reusable. Each of the data domain teams can manage the implementation of their own, local data products. Concurrently, there is a need for central data discovery, data marketplace, analytics and auditing to facilitate finding relevant data by users

With a data mesh, a data engineer needs to be more domain-focused on creating data products. Underlying data management needs for a data mesh include:

1. **Standardizing data products:** A product that facilitates an end goal with standardized and clean data. Examples of data products are:
 - **Google Analytics:** Analyze data using insights and machine learning capabilities to make the most of your data.
 - **Movie genre recommendations:** Predict the genre of the movie (action, drama, comedy, thriller, etc.) based on review text.

Data products require a standardized dataset through:

- a. Integrated data discovery and lineage
- b. Standardized templates and widgets
- c. Reusable data quality libraries

2. **Self-service infrastructure:** Product teams and application owners can rapidly provision and maintain application infrastructure, without depending on the IT operations team. For example: Infrastructure as Code (IaC) using Terraform open source (OSS) with the **Terraform Open Registry** for cloud provisioning. Technical capabilities required to build self-service infrastructure includes:
 - a. Metadata-driven design principles: Identify business metadata elements/ attributes that help define and qualify the data in terms of its source, ownership, load frequency, attributes, access, data sensitivity, etc. for infrastructure requirements.
 - b. Abstraction to data infrastructure: The goal is to abstract away the physical underlay and provide platforms — as opposed to servers — for developers to work on. For example, you can abstract infrastructure using container orchestrations like Kubernetes.
3. **Multiple integration patterns:** Integration patterns supply a standardized method for integrating data, such as:
 - **Migration:** Moving data from one system to another
 - **Broadcast:** Moving data from one system to another
 - **Aggregation:** Process of receiving data from different systems and merging it into one system for a centralized view
 - **Bi-directional synchronization:** Process of combining datasets from two distinctive systems to act as one

Top 5 Data Engineering Trends (continued)

- **Correlation:** Synchronization of datasets only happens if the records exist in both systems

Key requirements to support various integration patterns are:

- Data latencies
- Data volumes
- Operational and analytical use cases

Data engineers should focus on building flexible data pipelines by leveraging different integration patterns and selecting

data infrastructures to support cloud computing such as **elastic** and serverless.

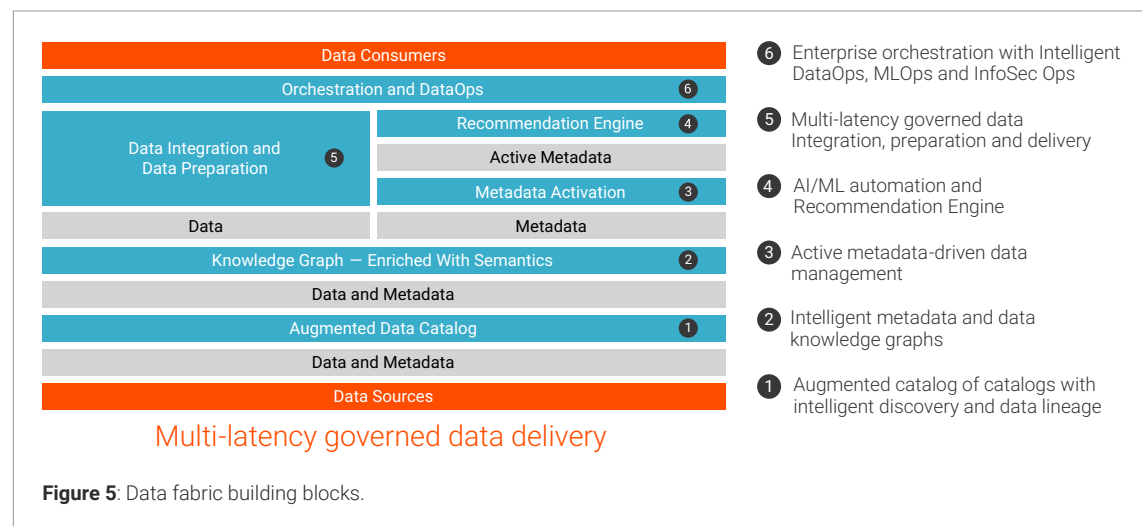
2. Data Fabric

A **data fabric** integrates and connects all your organization's data intelligently and efficiently by abstracting underlying complexity. It minimizes disruption by enabling a highly adaptable data management strategy. Data fabric is agnostic to deployment platforms, data processing methods, data delivery methods, locations and architectural

approaches. As such, it can help enable faster data-driven decisions through automated data management and broader data sharing, as well as optimize **data integration** and data preparation to improve productivity in a cost-effective manner.

As highlighted in Figure 5, the key pillars of data fabric are:

1. **Augmented data catalog:** An AI-driven **data catalog** enables you to find, understand and prepare all your data with automated metadata discovery and data cataloging.
2. **Intelligent metadata and knowledge graphs:** An enterprise knowledge graph that leverages AI and metadata foundation puts data in context by linking and enriching semantic metadata to deliver intelligence to data management functions. This includes data cataloging, **data governance**, data integration, **data quality** and **master data management**.



Top 5 Data Engineering Trends (continued)

3. **Active metadata-driven data management:** The metadata which actively moves to the places where the people need to access it becomes part of and adds context to the data management tasks. It has many use cases such as helping in automating data governance and protection for data quality improvements, data curation, data classification, policy enforcement and more.
4. **AI/ML automation and recommendation engine:** An AI engine learns your data landscape to automate thousands of manual tasks and augment human activity with recommendations and insights, allowing you to scale your data management to meet your organization's needs.
5. **Multi-latency governed data integration, preparation and data delivery:** Governed data integration and enterprise **data preparation** enable you to simplify and speed up data preparation with advanced ML-based automation and data cataloging.
6. **Enterprise orchestration:** Enterprise orchestration and XOps enables automatic orchestration of all data delivery flows by employing DataOps, MLOps and InfosecOps in support of continuous analysis and monitoring.

The four key data fabric principles for data engineering are:

- **Abstraction**, which is the process of removing characteristics from something to reduce it to a set of essential elements. An abstraction layer provides a common business understanding of the data. It requires templating and modularity without specifying underlying data storage or infrastructure locations at design times. Data abstraction enables you to seamlessly integrate data preparation capabilities, to help produce operational data pipelines. Such data pipelines can be easily scalable and maintain high performance while handling increasing volumes of data.
- **Strong DataOps capabilities**, such as advanced orchestration and DevOps with continuous integration and continuous delivery (CI/CD) capabilities.
- **Data discovery and metadata awareness**, which require catalog services to be integrated. Catalog services also provide automated lineage and transformation for active metadata and catalog solutions.
- **Data marketplace**, which provides selfservice data access and delivery for data engineers and data scientists.

Top 5 Data Engineering Trends (continued)

3. Modern Data Stack

The modern data stack (MDS) refers to the technology and tools that are used to collect, process and store data in modern data and analytics. MDS brings many benefits to data engineers, such as greater efficiency, faster and cost-effective ways to validate or experiment hypothesis and reduced technical debt and frustration. Let's look at what makes the data stack "modern" compared to traditional.

IDC predicts that the amount of data generated is growing nine times annually and expected to reach 221 zettabytes by 2026.¹¹ The growing volume, newer data types and format results in cost overruns, resource constraints and technology/implementation complexity. Traditional data stacks such as on-premises Hadoop (ecosystem) and SQL warehouses result in multiple challenges, including slow response to new information, lack of flexibility and difficulty to adapt to changes.

This brings us the concept of MDS, which helps in resolving existing challenges with traditional data stack. Key characteristics of MDS are:

- **Supports scalability:** As data volume is increasing and organizations are moving to the cloud for elasticity and scalability of cloud service providers such as AWS, Azure, Google, Snowflake, Databricks, etc., the performance should not be impacted, even if dealing with billions or trillions of records.
- **Easy trial and deployment:** Most of the tools to support data stack are software as a service (SaaS)-based. This makes it easy to deploy without thinking about server infrastructure.

- **Easy integration with other tools:** MDS tools should provide simple and easy integration with other tools and work cohesively as a common data platform.

As illustrated in Figure 6 on the following page, key components of MDS include:

1. **Flexible data replication, ingestion and integration:** Replicate, ingest and integrate data from multi-cloud, hybrid or SaaS applications in massive scale with features like version control and reusability. This should help handle the unpredictable data workload automatically without breaking and scaling in and out on demand. This also helps control cost and time with its data driven recommendations, which is referred to as elastic data integration.
2. **Single or multi-cloud data platforms:** MDS consists of single or multicloud application and storage platforms. It includes data warehouses and data lakes such as AWS Redshift, Snowflake Data Cloud, Azure Synapse, Google BigQuery, Databricks Lakehouse, etc.
3. **Data delivery:** Once data is standardized, cleansed and transformed reverse extract, the **extract, transform and load** (ETL) process helps deliver data to analytics platforms.
4. **Data analytics:** Modern data analytics tools should support self-service and line of business (LOB)-specific analytical capabilities.

¹¹ IDC, Global Datasphere 2022-2026 Forecast, Doc #US49018922, May 2022

Top 5 Data Engineering Trends (continued)

5. **DataOps:** DataOps is a set of practices, processes and technologies that combine an integrated and process-oriented perspective on data with automation and methods from agile software engineering to improve quality, speed and collaboration. It also promotes a culture of continuous improvement in data analytics.¹² DataOps provides a way to operationalize your data platform by extending the concepts of DevOps to the world of data.
6. **Data governance and catalog:** In a traditional data stack, data processing concerns were more important than data governance. Data teams should focus on making data discoverable and secure through the simplification of managed cloud services and flexible resource management. A solid governance strategy and framework protects data lakes and warehouses from becoming a data swamp.

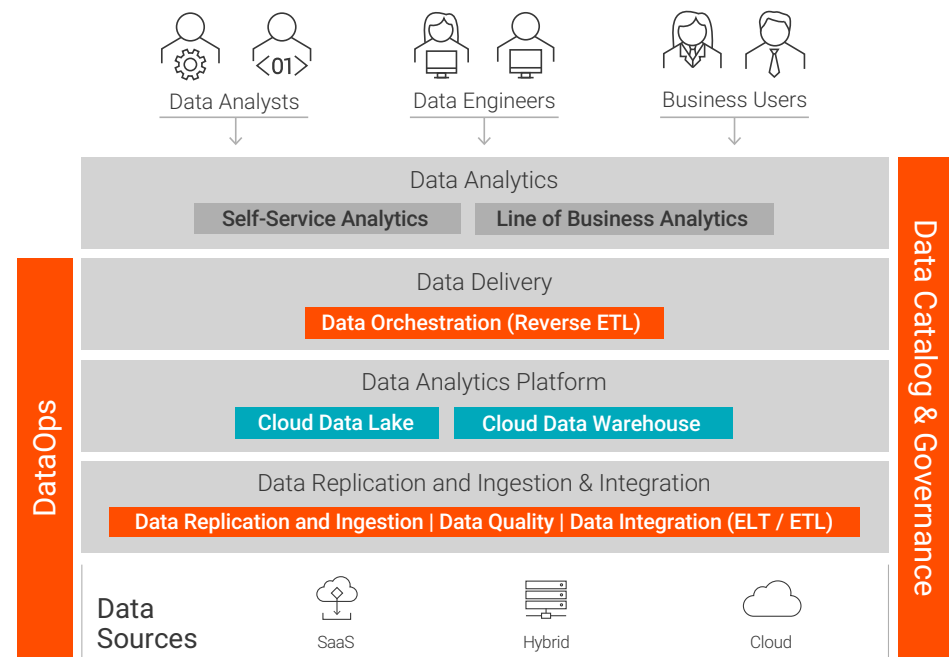


Figure 6: Modern data stack.

¹² <http://ceur-ws.org/Vol-2191/paper13.pdf>

Top 5 Data Engineering Trends (continued)

4. Data Management Platforms

Cloud modernization initiatives often times fail due to technical debt. Multiple point solutions make it even more complex to stitch information together. For example, multi-cloud and hybrid environments using different vendors for different solutions (such as data governance, integration, master data management) may seem like the easy choice at the initial stage. But once point solutions come together, it creates multiple problems such as compatibility issues, continuous rework due to version upgrades and specific unsupported features.

To realize successful business outcomes, enterprises require a data management strategy that encompasses the entire end-to-end data lifecycle — from the edge to AI — to enable a connected strategy. This can be achieved through one data management platform that supports all data management functionalities across multiple cloud vendors, any source/target across different SaaS application and on-premises systems.

The platform should offer comprehensive, best-of-breed data management applications, such as:

- Data catalog to discover, organize and curate data assets
- Data integration with virtually any pattern, streaming data, batch data, ETL, ELT, data engineering, and hybrid and multi-cloud integration
- Application integration to connect business applications, automate and streamline business processes and workflows and provide built-in API management functionality

- Data preparation to enrich and prepare data
- Data quality to profile, cleanse and improve the quality of data
- Master data management and tailored solutions that provide a 360-degree view of your business
- Data marketplace and data services to democratize data and facilitate data sharing
- Data privacy to discover, classify, protect and tag sensitive data
- Data governance to build business glossaries for data standardization to manage and enforce policies, maximize the value of data and enable a foundation of trusted data
- A metadata knowledge graph that delivers consistent data intelligence across all data management applications to provide a unified metadata foundation

Benefits of One Data Management Platform

There are many benefits of establishing one data management platform for data engineers, including the ability to:

- Work with connected, composable data strategies:** Build composable architectures and enable connected data strategies with a fully integrated, interoperable data management cloud that enables the end-to-end data lifecycle.

Top 5 Data Engineering Trends (continued)

- b. Scale data management with AI-powered automation:** Simplify and scale data management by automating thousands of data management tasks using an AI engine.
- c. Modernize to the cloud to enable next-gen analytics:** Accelerate modernization to cloud data warehouses, lakes and applications with seamless integration across hybrid, multi-cloud, and pre-built automation tools. Capitalize on proven methodologies and reference architectures.
- d. Effectively govern data lakes, data warehouses and lakehouses:** Discover, catalog, curate, classify data assets and manage policies with cloud data catalog and governance capabilities to prevent data swamps and deliver accurate insights.
- e. Gain a 360-degree view of critical data domains:** Manage all master data domains at scale in a single SaaS solution. Increase productivity and efficiency with low-code/no-code experiences and AI powered automation.
- f. Empower data collaboration and sharing:** Connect data and organizational silos and democratize and share data assets and products via data marketplaces and data services.
- g. Capitalize on data-in-motion for real-time insights:** Quickly ingest, integrate, prepare and deliver insights from streaming and device data with mass ingestion and data streaming capabilities.
- h. Connect ecosystems and automate end-to-end business processes:** Connect business applications and build, secure and automate intelligent business processes with cloud native application integration capabilities.
- i. Accelerate time-to-insight with data fabric and data mesh architectures:** Enable data fabric and data mesh architectures with capabilities like continuous integration, knowledge graph, governance automation, data marketplace and self-service infrastructure.

Top 5 Data Engineering Trends (continued)

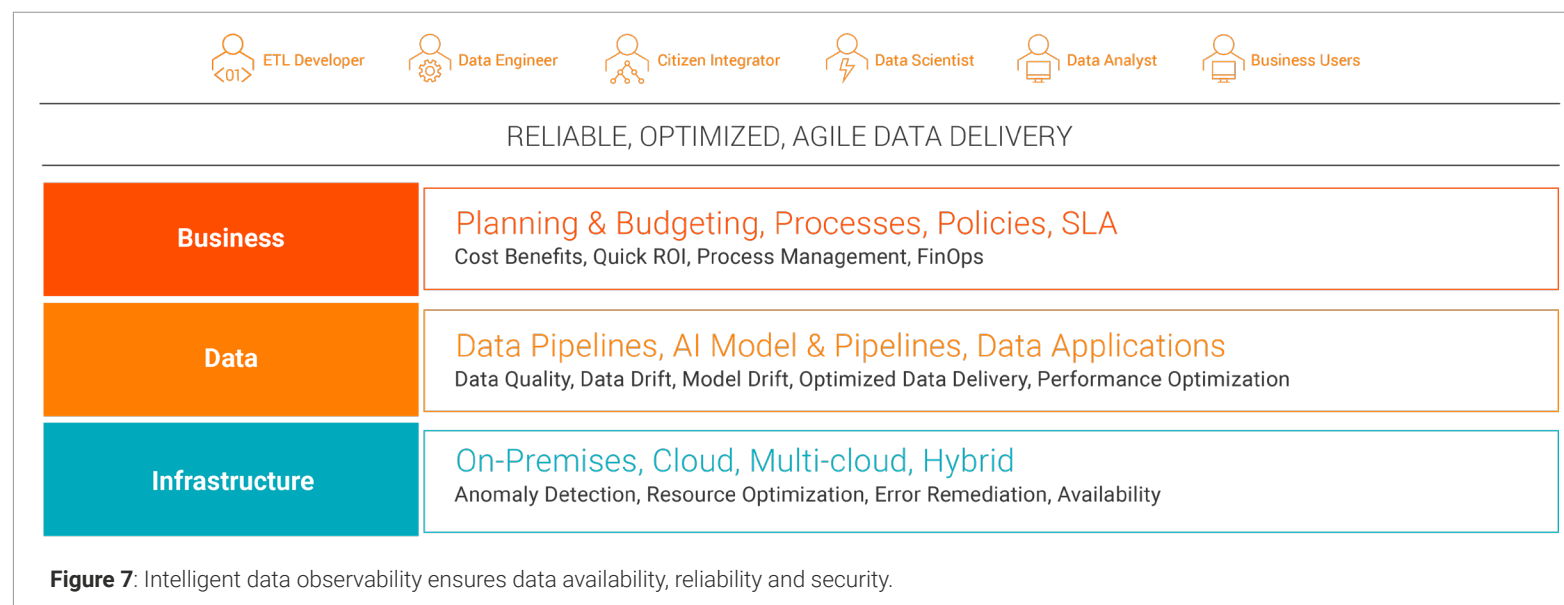
5. Data Observability

Data Observability is the ability to understand the health of an organization's data landscape, data pipelines and data infrastructure. This is done by continuously monitoring, tracking, alerting, analyzing and troubleshooting incidents to reduce and prevent data errors or downtime. Data observability is useful to engineers because it does not just stop at describing the problem. It instead provides context and suggestions to help solve the issue through the following:

- Discoverability of the data assets

- Understanding the data context (such as where the data originated from and its related business processes)
- Accessing and moving the data where needed
- Ensuring trust and fit for usage
- Protecting and ensuring privacy and data security
- Governing data consumption and optimizing data usage

Let's take a look at how some of these data engineering trends can be used in real-world use cases.





Case Studies | Data Engineering Trends in Action

BMC Transforms Complex Technology Into Extraordinary Business Performance With a Data Fabric

BMC Software, Inc. is an American multinational IT services and consulting and enterprise software company based in Houston, Texas. BMC works with 86% of the Forbes Global 50 and customers and partners around the world. BMC's goal is to help them run and reinvent their businesses with open, scalable and modular solutions to complex IT problems.

Challenge: The company's accounts payable and generic ledger operations were handled by decentralized regional services centers using manual processes. This, in turn, caused a lack of standardization across countries. It impacted the BMC treasury team's ability to view current account balances. This resulted in the need to maintain excessive cash reserves to cover any unpredicted cash needs.

Solution: With Informatica data integration and mass ingestion capabilities, BMC built a functional system in a very short period. Then, it layered on more sophisticated capabilities such as data fabric to increase cloud data availability for anytime access. The company dramatically improved visibility into actual and projected cash flows. This enabled BMC to better manage cash positions and optimize the use of its working capital.

Results: BMC saved hundreds of thousands of dollars and now has much better reporting and control across hundreds of bank accounts. With accurate and timely visibility into its cash holdings, it has also elevated the rigor behind its risk management and mitigation strategies. BMC plans to dive deeper into its data fabric to make data more available to more users with data catalog and data governance powered by **Informatica Intelligent Data Management Cloud™ (IDMC)**.

"Over the years, we have amassed terabytes of data, which is at the core of everything we do. And the crown jewel is our data warehouse, which helps us support our customers' growing data demands and drives insights that will make them future-ready."

Jeff Gheen

Senior Director of Data Analytics and Data Warehousing, BMC



Delivering Breakthrough Therapies Faster With One Data Management Cloud

The Takeda Pharmaceutical Company Limited is a Japanese multinational pharmaceutical company, with partial American and British roots. It is the largest pharmaceutical company in Asia and one of the top 20 largest pharmaceutical companies in the world by revenue.

Challenge: To speed research and development, Takeda needed to build an enterprise cloud data platform that enables quick, easy access to analytics tools and also reduces cost.

Solution: Takeda invested in Informatica's data management cloud to move to the cloud through a cloud data lake for advanced analytics. This solution helped them transition to a serverless architecture and optimize compute resources using Informatica's data engineering capabilities and **Databricks**. This set a solid foundation for a digital and multi-cloud journey.

Results: By creating an integrated, central repository for enterprise data with a cloud data lake, cloud data warehouse and cloud data integration using Databricks and Informatica, Takeda achieved an estimated 40% to 50% cost reduction by avoiding clusters. They were also able to accelerate decision making with easy access to critical analytics in a cloud environment. This enabled them to optimize cloud compute resources based on changes in demand using auto scaling.

We've examined how the latest trends can help organizations thrive. Now let's explore how data engineering supports AI, ML and data science for better customer experiences and more-informed decisions.

By using Informatica and Databricks, Takeda optimized cloud compute resources and achieved an estimated 40% to 50% cost reduction by avoiding clusters.

Artificial Intelligence, Machine Learning and Data Science

Now that you're up to speed on the latest patterns, let's shift gears and jump into AI. As evidenced from innovations like robots, facial recognition, smartphones and driverless cars, AI is poised to transform organizations across the globe and change the way people work. According to a PwC study, AI has a \$15.7 trillion potential contribution to the global economy by 2030.¹³ The amount of data generated by both humans (real data) and machines (synthetic data) is estimated to reach 463 exabytes globally daily by 2025.¹⁴ But a large part of that data is meaningless until it's converted into valuable information.¹⁵

The bottom line: AI is a game changer for society and the economy, for sure. But how will it impact data engineers? Let's dive in to better understand how AI, ML and data science help bring modern data engineering to life.

AI and ML

AI is a broad area of computer science applications that perform complex tasks that once required human input. It focuses on acquiring data and creating rules for how to turn the data into actionable information and it "learns" on its own, with no human intervention.

ML is a subset of AI that allows machines to learn from data without being programmed. For example, self-driving cars use deep learning, a subset of AI, to recognize the space around a vehicle to avoid accidents.

Historically, data engineering responsibilities involved a lot of complexity. While there is repetition in data operations tasks, no two projects are the same. There are many tasks, such as running **ETL (extract, transform, load)** jobs, **data preparation** and integrating with third-party APIs before any AI/ML can happen.

The data demands for ML and AI are so high that companies are financing digital transformation initiatives and building a data-first stack¹⁶ to store and analyze data. They are also investing billions of dollars in building AI functionality.

There are three key challenges that make data operations more important than ever:

1. The data explosion we're seeing will continue with new data formats and types emerging from IoT, connected devices and more. This includes sensors being integrated into everything from electronic devices to retail stores. Deriving insights for decision

¹³ <https://saabrds.com/how-much-will-ai-contribute-to-the-global-economy-and-the-industrial-market>

¹⁴ <https://techjury.net/blog/how-much-data-is-created-every-day/#gref>

¹⁵ <https://www.thehansindia.com/hans/young-hans/the-value-of-artificial-intelligence-data-science-in-todays-world-712209>

¹⁶ <https://moderndata101.substack.com/p/evolution-of-the-data-stack-the-story>

Artificial Intelligence, Machine Learning and Data Science

making becomes difficult when dealing with large data volumes and a variety of formats.

2. The evolution of **business intelligence** from descriptive analysis to AI/ML will accelerate the need for more **data integration** to feed AI/ML models. Garbage in is garbage out, so an AI/ML model is only as good as the data that its being fed into.
3. AI/ML is not niche anymore. It is now an integral part of critical, enterprise-wide decision making. As AI/ML advances, more data will need to be moved. This puts untenable demands on the data operations function if they are working with antiquated data management systems.

Rather than having more resources to perform repetitive tasks, you need a framework to automate tasks and simplify data pipelines. A modern approach to data management includes auto scaling the infrastructure up and down, smartly shutting down the infrastructure when not in use and automating performance tuning through an auto tuner with elastic capabilities.

Data Science

To understand the role data engineers play in data science, look at auto racing. A car races around the track and pulls into the pitstop for fuel, new tires and windshield cleaning. After a few seconds, the car zooms off again and the process repeats.

The finish is exciting — the winner jumps on the car and sprays Champagne. But what about the pit crew who changed the tires? Do they get to celebrate with Champagne? They may not get credit for winning the race but they can't be blamed for losing the race.

Data science is sort of like this. Data scientists get accolades when their recommendations to leadership move the business forward. But behind the scenes, it's the data engineer who "changes the tires and fills the tank." Data engineers take data from various stores and profile, standardize, transform and cleanse it so that the data is usable by the scientist. Data engineers design and implement the architectures necessary for data scientists to be successful. Data engineers may not win the race, but without them, the data scientist never gets to do a victory lap.

Artificial Intelligence, Machine Learning and Data Science (continued)

Let's walk through the reference architecture (Figure 8), which outlines the data engineer's responsibilities in data science.

Exploratory zone. Data scientists want to spend as much time as possible operating in the exploratory zone – the center teal box. This is where data scientists perform exploratory data analysis, create and evaluate models and so much more. But to do these exciting things, they need complete, trustworthy data. How does this happen? You guessed it – the data engineers make it possible by working diligently behind the scenes.

Data sources. Starting on the left, our data sources are usually a grab bag of various technologies. The data engineer must figure out how to access and extract data from all these sources. With systems like mainframes, IoT devices, Avro, JSON, Hadoop, relational databases and more, even this first step can be daunting.

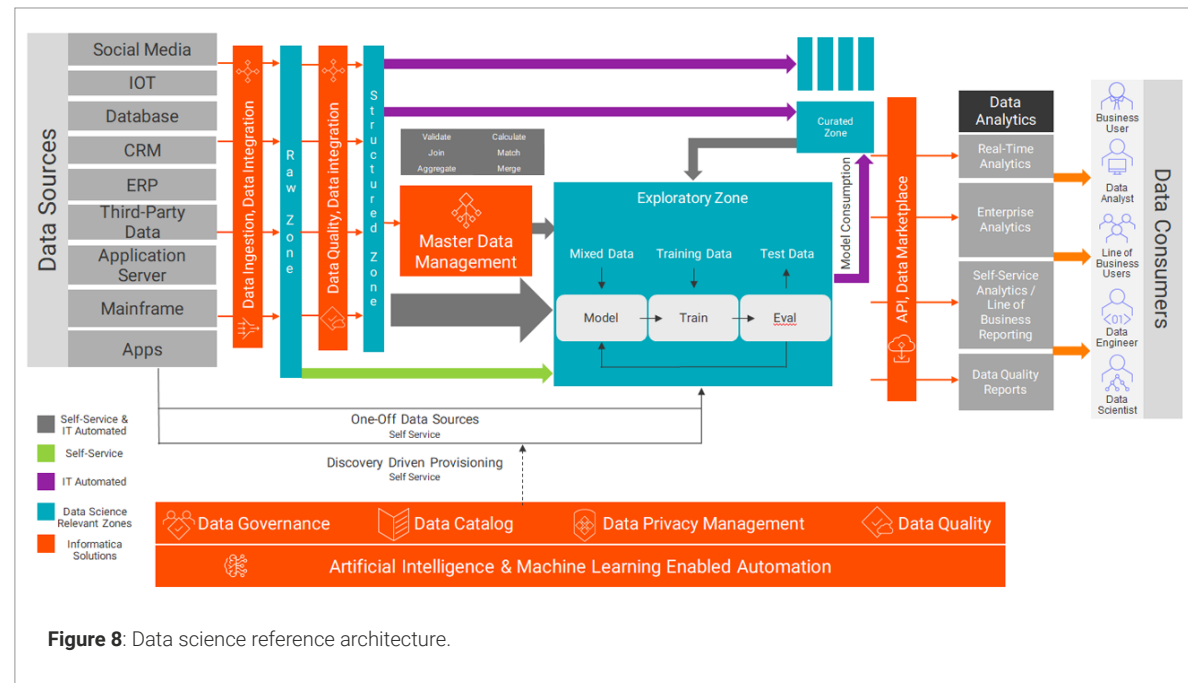


Figure 8: Data science reference architecture.

Artificial Intelligence, Machine Learning and Data Science (continued)

In a **data lake architecture**, the approach is to grab all data from various sources and move it unchanged into the raw zone. However, not everything always makes it over. For example, for an organization that streams vibrational measurements from a container ship's engines, it's just a huge file full of numbers. This is not something that would ever be governed and curated or used in traditional analytics, but it ended up being extremely useful for the data scientist.

In this case, the data files were saved to a USB drive, then downloaded onto a file system on the ship and eventually moved onto a corporate file server. The data never made it into the data lake. Luckily, the file server was automatically cataloged, and only by accident did a data scientist discover these files.

Through the catalog, they self-provisioned the data into the exploratory zone for analysis and ended up creating a predictive maintenance model that dramatically decreased operating costs of these massive container ships. One-off provisioning is not a sustainable process, so the data engineer created an automated, ongoing process for getting data from vibrational data loggers into the raw zone and ultimately, the exploratory zone.

The data engineers are responsible for the data we know about but also sometimes for the data we don't know about. Mass ingestion, **data integration**, data cataloging and **data governance** are all important capabilities for data science with respect to getting data out of its native sources.

Raw zone. The general practice is to replicate and ingest data as-is from the native data sources into the raw zone. As a result, it contains more data than any other zone.

Data in the raw zone is generally not suitable for direct consumption. For example, you may have a column of strings like '11/02/2021'. We know that's a date, but in the raw zone it's just another string. Speaking of which, is that date November 2, 2021, or February 11, 2021? It's impossible to know because the data is cataloged, but not curated, in the raw zone.

"Speed is everything in our sector. Through Informatica, we're accelerating product development and reducing time to market for our therapies, which helps to enrich the lives of people around the world."

Shyam Dadala

Enterprise Analytics
Architecture Engineer
Shire Pharmaceuticals

Artificial Intelligence, Machine Learning and Data Science (continued)

The data engineer's primary responsibility here is creating a productionized process to move data into the raw zone. A mass ingestion approach is most appropriate because being comprehensive is more important than format or cleanliness.

The data engineer works with a wide variety of data sources — which can be very challenging. Files of various formats, RDBMS, mainframes, streaming data, incremental data and more all present different issues. For example, ingestion must be a low-maintenance process resilient to interruption and schema drift. If a new file appears in a directory or a new table is created in a schema somewhere, the data engineer does not have to explicitly find that file or table and create a new process to move it into the raw zone. Similarly, if a new column is added to a data source, the ingestion process should move the data autonomously without requiring intervention.

Again, not all data makes it past the raw zone. Data scientists will sometimes access the raw zone directly when they cannot find relevant data elsewhere. Governance then becomes interesting because, as noted earlier, raw data typically isn't governed or curated. However, data in the raw zone should be cataloged to help data scientists and engineers discover attributes that may not exist anywhere else in the data lake.

Structured zone. This is the first stage of transformed data. In the structured zone, data is typically stored, typed and cleansed in a tabular structure of rows and columns.

Data engineers are responsible for the integration process to transform and move data from the raw zone to the structured zone. Unlike when we moved data into the raw zone, this process includes applying quality functions to standardize and cleanse data. For example, all dates are typed as a date and may be formatted dd-mm-yyyy.

The structured zone is also where we start applying curated governance. For example, we might know that an element formatted as 999-999-9999 is an identifier, rather than a phone number.

This process should be as automated and dynamic as possible so as sources change or are added, the data can be quickly profiled, cleansed and structured for data science and analytics.

Although advanced analytics are not performed here, most of the data ultimately used by data scientists will come from the structured zone. Often, data scientists want to further manipulate data, so we see data preparation used to extract and transform data from the structured zone into the exploratory zone. In this context, data preparation delivers value as an agile, self-service activity. The data scientist needs to quickly discover, manipulate and provision data into the exploratory zone to try things out. At this point, the data may or may not be useful so we can't insert a delay into the process.

Artificial Intelligence, Machine Learning and Data Science (continued)

However, once the scientist determines a data set is critical to analysis and modeling, the data engineer is responsible for taking whatever happened in data preparation and turning it into a production process. Ideally, data preparation creates a recipe behind the scenes that the engineer can directly convert into production.

In leveraging structured zone data, the engineer must understand all the steps taken in moving the data from source to raw, raw to structured and then whatever additional manipulation the data scientist may have applied. They also need to make sure the process is repeatable, well designed, resilient, dynamic and as automated as possible.

Curated zone. Think of data in the curated zone as **data marts** and **warehouses**. This data is usually aligned with a model and as a result, comes from multiple sources. This data is directly used by consumers, so it comes from trusted zones such as structured, master and in some cases, the exploratory zone. Data in the curated zone also leverages ML models which are consumed to provide insights.

Here, data has been aggregated and augmented with calculations and experienced data quality checks. Because of this, engineers and data scientists will extract data from the curated zone and move it to the exploratory zone, so they don't need to separately perform those calculations again.

Data engineers need to design integration routines to populate the warehouse and marts. In the context of data science, they also need to understand and manage data between the curated and exploratory zones.

Master data zone. Master data isn't often considered in the world of data science, which can be shortsighted. Many times, data scientists will analyze trends in the context of a master data domain. We've all seen analysis by state or country, which are typical reference data subjects. Within a specific enterprise, we may want to analyze how specific products are selling in a specific customer demographic. In this case, products and customers should be as relevant as possible to the enterprise. And these specific products and customers come from the master data zone.

In this context, data engineers are responsible for interacting with the **master data management** (MDM) solution. Mature MDM solutions include a service framework that makes working with the master data very easy. Data scientists figure out how to relate master to analytic data, but the data engineers are then responsible for integrating the MDM service framework into a production process.

MDM, data integration, data catalog and data governance are all important capabilities with respect to leveraging master data in the data science architecture.

Data Science Use Cases Supported by Data Engineering

Data science and AI/ML are used in various industries, particularly:

- **Healthcare:** Providers, payers and other healthcare organizations have large amounts of data such as medical records, diagnostic information and medical claims. When used properly, this information can help reduce patient readmittance, detect fraudulent payments and find propensity of illness.
- **Banking and financial services:** Fraud detection, credit and loan approvals and blockchain are key applications of data science.
- **Retail:** From special offers to campaign effectiveness and churn analysis, data science can help companies provide the best customer experience while increasing their top line.

Similarly, other data science use cases exist in pharma, automotive and transportation, insurance, energy, government, sales and supply chain management.

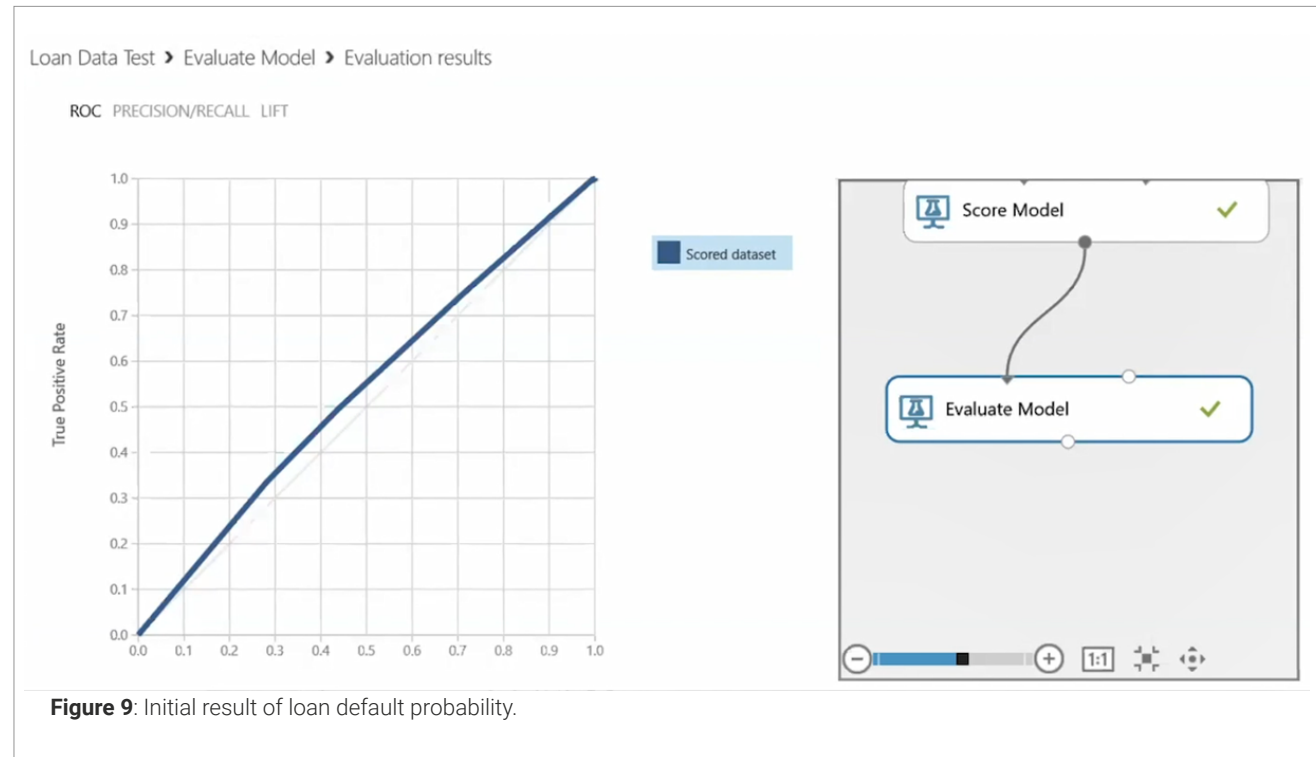
The Anatomy of a Successful Data Science Use Case

Here is an example of how using data quality dramatically improved the data science outcome.

A financial institution wanted to predict the probability of loan default for their customers. This was a very straightforward model based on various attributes including a field called 'years_in_current_job', which indicates years of experience at their current job. The data was first sourced and split, and a decision tree was applied. The model was then trained, scored and evaluated.

Data Science Use Cases Supported by Data Engineering (continued)

The results were not initially encouraging. In Figure 9, the chart on the left shows a dark blue receiver operating characteristic (ROC) curve. If you look closely, you will see a faint gray line just below that. That gray line represents a random classifier. Because the ROC curve is only slightly above it, the model is performing only slightly better than a coin flip.



Data Science Use Cases Supported by Data Engineering (continued)

It's tempting to implicate the model and try something different. However, in this case, we profiled the data and noticed the 'years_in_current_job' field was poorly distributed as seen in Figure 10.

The histogram on the right shows the data is skewed with the majority falling in the "10+ years" classification. Further, on the left, we see most data is non-distinct which seems odd since this field measures years of service.

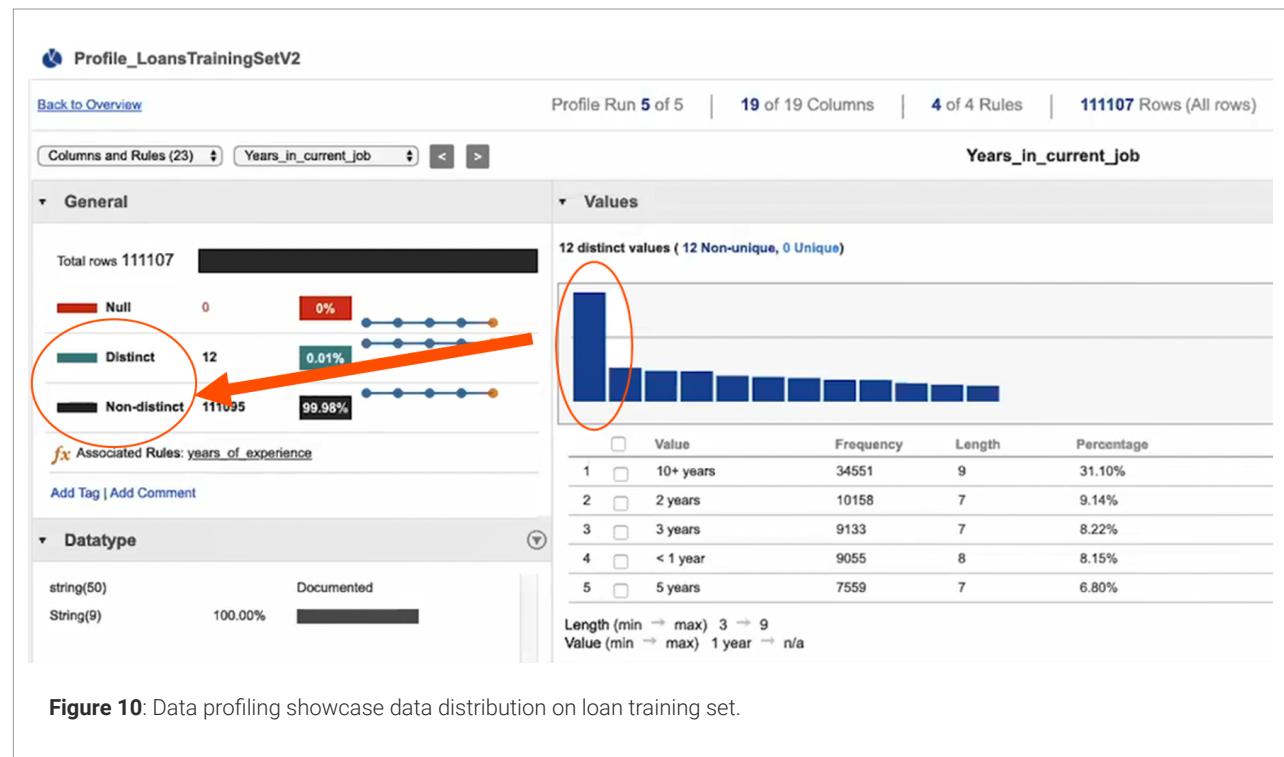
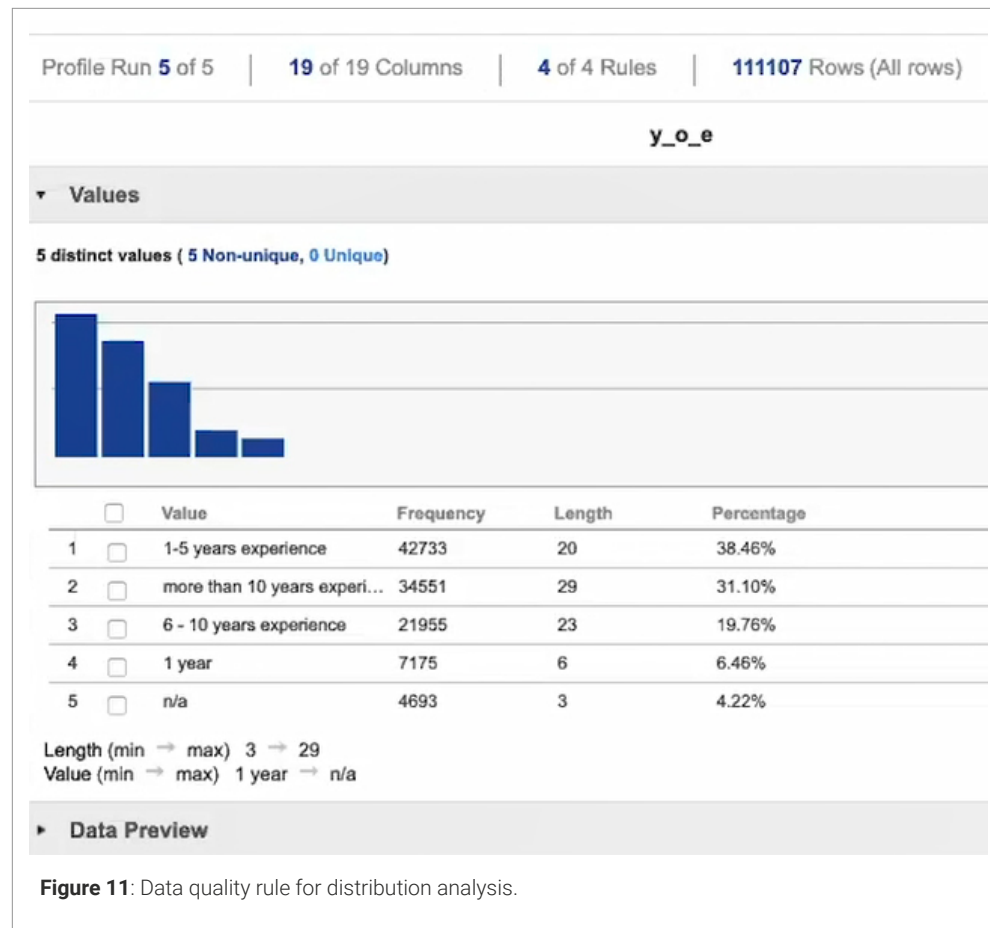


Figure 10: Data profiling showcase data distribution on loan training set.

Data Science Use Cases Supported by Data Engineering (continued)

To better distribute the data, we created a simple **data quality** rule. After application, the histogram now looks like Figure 11.



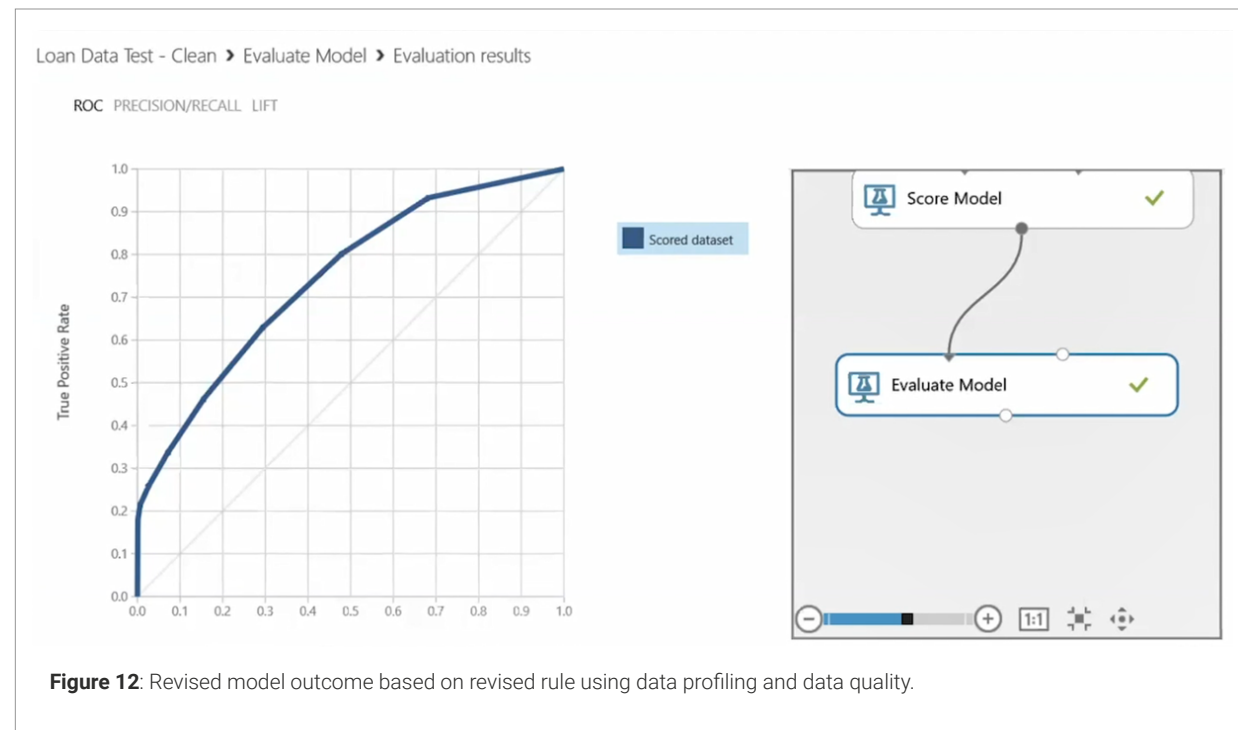
Data Science Use Cases Supported by Data Engineering (continued)

It's important to note that we did not touch the actual data. The raw data remains unchanged. We simply created a rule to distribute it more evenly.

Re-evaluating the model produced the following results, as shown in Figure 12.

The ROC curve is now significantly better than the random classifier "coin flip" line. Better yet? Loan default predictions will be much easier.

So, what does this mean? The model wasn't incorrect, but the data was! It would be very tempting to condemn the model in this case. But by profiling the data, we were able to identify an attribute that was poorly distributed. Plus, adding one simple data quality rule resulted in a much more predictive model.



Operationalization in Data Engineering

The true potential of data engineering processes can be realized once they meet the needs of the business and are accepted by key stakeholders. Now let's look at the operationalization in the context of data engineering.

Operationalization is the process of bringing together the right data, at the right time, for the right users — all in a repeatable and collaborative fashion that can be trusted for business insights and actions.

MLOps: Operationalizing ML Models

In 2020, a McKinsey Global Survey found that only about 15% of respondents have successfully scaled automation across multiple parts of the business. And only 36% of respondents said that ML algorithms had been deployed beyond the pilot stage.¹⁷ And today? The same challenges still exist, and we do not see much, if any, improvement.

Another reason AI/ML projects fail is because they lack a framework and architecture to support model building, deployment and monitoring. We call this very necessary framework machine learning operations (**MLOps**), a core function of ML engineering. MLOps focuses on taking ML models to production and then maintaining and monitoring them. It is a collaborative function, often comprised of data scientists and other IT professionals.

Most organizations engaged in data science have defined a process to build, train and test ML models. The challenge has been what to do once the model is built. Integration, deployment and monitoring are essential aspects for providing continuous feedback once the model is in production. This is where the entire process of building ML models aligns more closely with the software development lifecycle than with an analytics project. Many organizations think data science projects are limited to creating models. After being developed and deployed, many other aspects are needed to operationalize them. For example, the model must be:

- Managed and monitored to ensure it performs optimally within the thresholds defined by the business
- Monitored by model drift or degradation with a feedback loop
- Tweaked based on the above steps and periodically retrained

¹⁷ The imperatives for success with automation technologies, McKinsey and Company

Operationalization in Data Engineering (continued)

As shown below in Figure 13, MLOps operationalizes the ML model development process to establish a continuous delivery cycle of models that form the basis for AI-based systems.

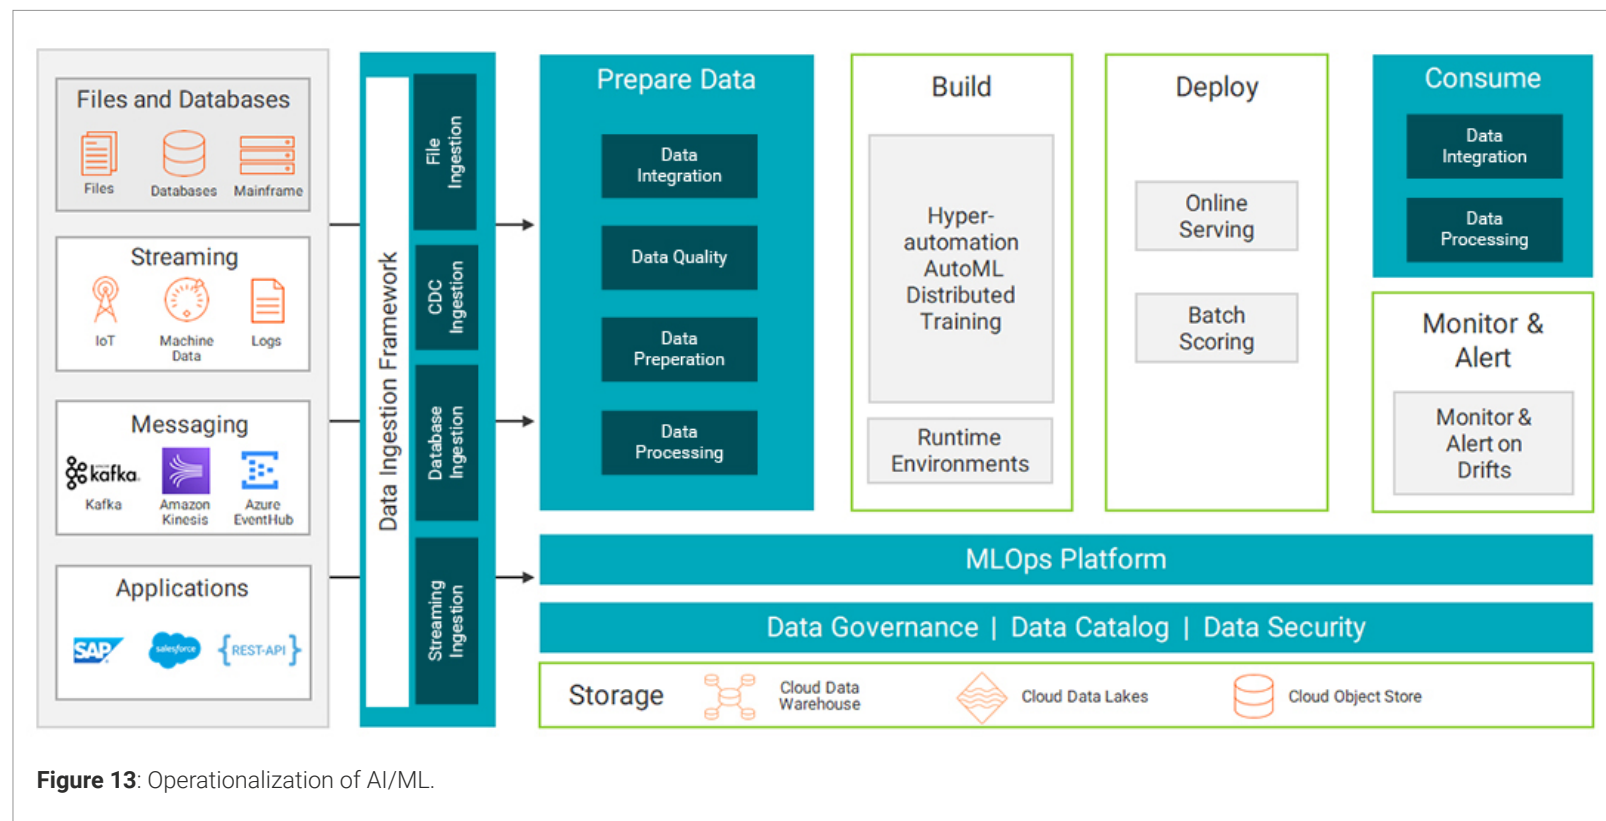


Figure 13: Operationalization of AI/ML.

Operationalization in Data Engineering (continued)

There are five steps of the MLOps process that are necessary for a successful data science project:

1. Business Understanding
2. Data Acquisition
3. Model Development
4. Model Deployment
5. Model Monitoring

1. Business Understanding

The right data governance solution facilitates collaboration between data governance and data stewardship practitioners and the subject matter experts (SMEs) with line of sight into relevant systems, data processes and owners. It provides an entry point for exploring context in which the business problem is first identified. It also helps explain the problem space as well as the approach to addressing project KPIs.

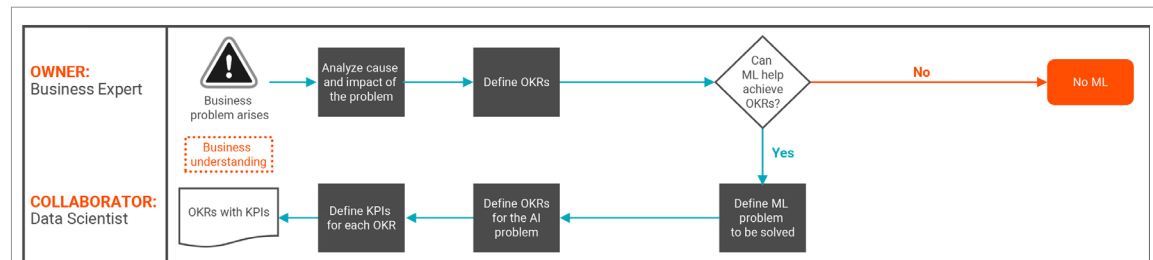


Figure 14: Step 1 of the MLOps flow: Business understanding.

Operationalization in Data Engineering (continued)

2. Data Acquisition

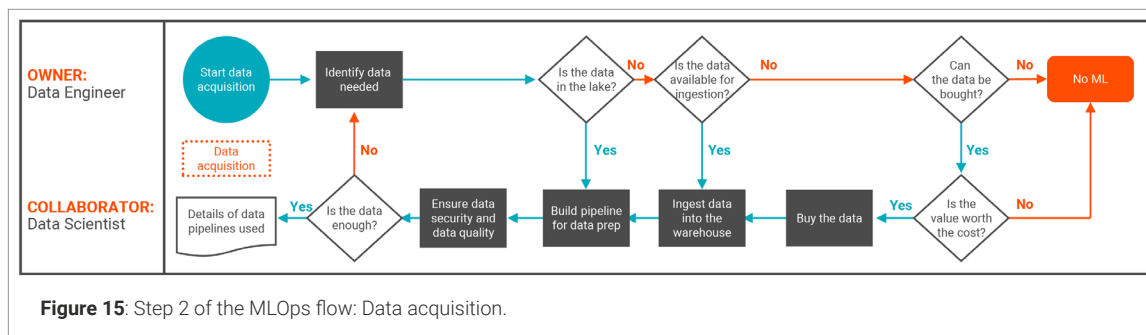
During the data acquisition phase, data is gathered for the solution. Data acquisition involves getting access to large amounts of data that are distributed. A data lake, along with a data catalog for search, are essential for efficient data acquisition. Moreover, **data management** tools greatly facilitate the creation and handling of complex **data pipelines** as compared to simply hand coding them. **Data security** and data quality are also important aspects and should ideally be done using appropriate tools.

Successful data acquisition requires capabilities across four solution areas:

- **Catalog:** An intelligent, enterprise-class **data catalog** enables business and IT users to unleash the power of their enterprise data assets by providing a unified metadata view that includes technical metadata, business context, user annotations, relationships, data quality and usage. It helps users discover the right datasets for modeling. Ideally your data catalog solution will integrate with a data governance tool. This will enable users to easily see definitional information (such as glossary terms) and key stakeholders directly in the catalog.

The right data governance solution will enable you to see the business context and processes the data is used in, and provide a holistic view on usage, quality levels and applicable policies.

- **Ingest:** Data acquisition requires efficient ingestion of data into on-premises systems, cloud repositories and messaging hubs like Apache Kafka. This ensures it's quickly available for real-time processing. In addition, your solution should provide support for streaming IoT and log data, large file sizes and **change data capture** (CDC) for databases. A solution that offers cloud-based services (database, file and application) to meet your specific data replication and ingestion needs is critical to your success. Even better is a solution that has authoring wizard tools to easily create data replication and ingestion pipelines and real-time monitoring with a comprehensive dashboard..



Operationalization in Data Engineering (continued)

- **Process:** Data engineers can help data scientists and data analysts by:

- Finding the right data and making it available in their environment
- Ensuring the data is trusted and sensitive data is masked
- Operationalizing data pipelines and helping everyone spend less time preparing data

The ideal solution is a comprehensive data engineering portfolio that provides virtually everything you need to process and prepare workloads to fuel AI/ML and analytics: robust data integration, data quality, streaming and **masking** capabilities.

- **Deliver:** Data scientists and data analysts need to rapidly discover, enrich, cleanse and govern data pipelines for faster insights. An AI-powered **data preparation** tool can help simplify self-service data preparation across cloud and hybrid data lakes.

3. Model Development

Model development is the core of the MLOps flow. Up until now, the data scientist has been in an advisory and approver role. Now that the problem and KPIs are clearly defined and high-quality datasets are readily available, the data scientist can leverage their expertise.

During model development, the data scientist iterates through multiple candidate models, validating them against test data and measuring KPIs until expectations are met. If more data is needed, the data scientist

can again coordinate with the data engineer on data acquisition and perform additional cleansing and standardization operations.

As a result, the data scientist can identify a model and provide performance metrics that can be used as benchmarks. These metrics may be quite different from the KPIs. For example, they may be more like what would be published by a data scientist working on a standalone project.

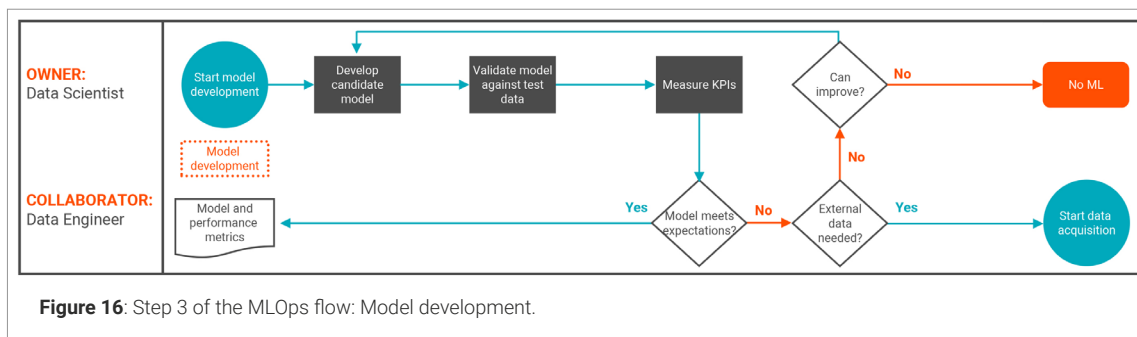


Figure 16: Step 3 of the MLOps flow: Model development.

Operationalization in Data Engineering (continued)

4. Model Deployment

The data engineer drives this phase, using the pipelines defined during the data acquisition phase as a starting point. The data engineer integrates the model developed by the data scientist and validates it against actual production data. Metrics and KPIs from previous phases are also validated. Invalidation means returning first to the pipeline and then to the model development process to determine the error source.

Once a validated pipeline is identified, a new pipeline that measures metrics for future monitoring must be established. This will allow continuous validation of the metrics identified to ensure that the model remains correct with time. Changing upstream data models and data distributions make this a critical step for any models that are expected to be used over time. This final pipeline is deployed in production with the help of the DataOps team for continuous use and monitoring.

The model deployment phase requires tools that allow easily reproducible and reliable pipeline deployment. You can use a server such as Jenkins for automating deployment jobs, REST APIs for communication between required modules and Docker for containerization.

A data engineering integration tool can help you deploy your ML model in the production pipeline. Ideally it provides out-of-the-box Python transformation or REST consumer transformations to deploy the ML model and integrates with a data quality tool. This ensures that the same quality operations are performed on the input data as were used during the model building phase, further increasing the efficiency of the process.

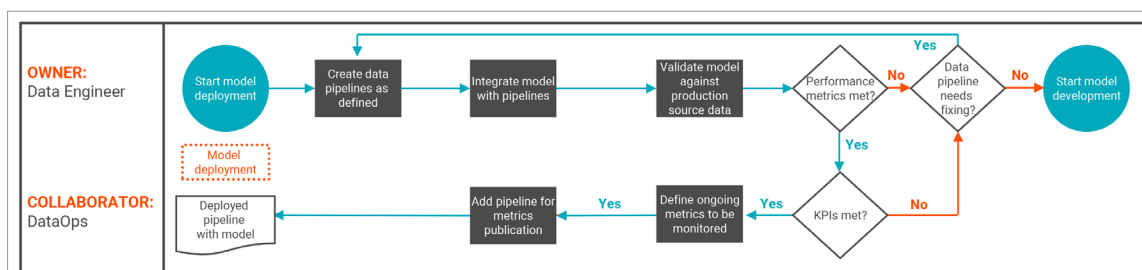


Figure 17: Step 4 of the MLOps flow: Model deployment.

Operationalization in Data Engineering (continued)

5. Model Monitoring

During the model monitoring phase, a deployed pipeline is integrated with a metrics monitoring mechanism. The DataOps team can then monitor the pipeline metrics, ensuring continued value and increasing confidence in ML. Alerts are generated any time performance metrics are not being met. Often, a change in the data flow is the cause and a minor data pipeline change is all that is needed.

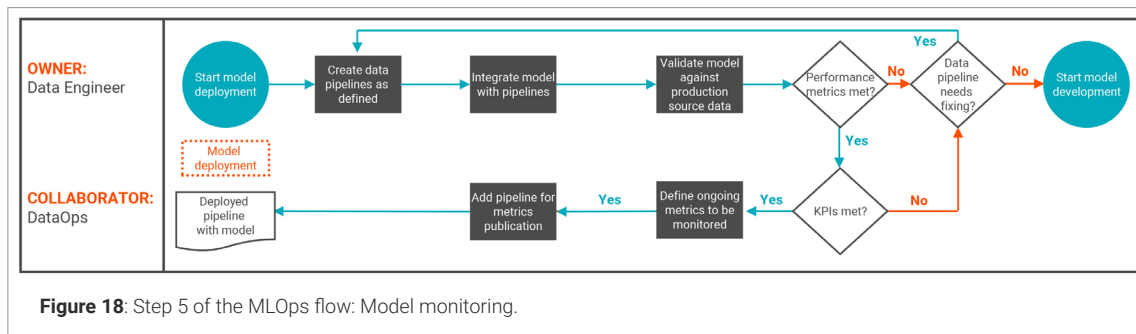
In some cases, there is a general change in the underlying data pattern and the model needs to be redeveloped. In extreme cases, a significant change in the business landscape requires going back to the business understanding phase to address a new business problem.

Continuous data profiling and quality scorecard evaluations help identify changes in data over time that require updated model training and evaluation.

DevOps and Continuous Integration/Continuous Development (CI/CD)

DevOps is a combination of software development (dev) and operations (ops). It's a culture shift or methodology that encourages communication and collaboration to build better-quality, more reliable software faster. DevOps organizations break down the barriers between operations and engineering by cross training the teams. This approach leads to higher quality collaboration and more frequent communication.

Continuous integration (CI) is a software development practice in which all developers merge code changes in a central repository multiple times a day. Continuous delivery (CD) adds the practice of automating the entire software release process.



Operationalization in Data Engineering (continued)

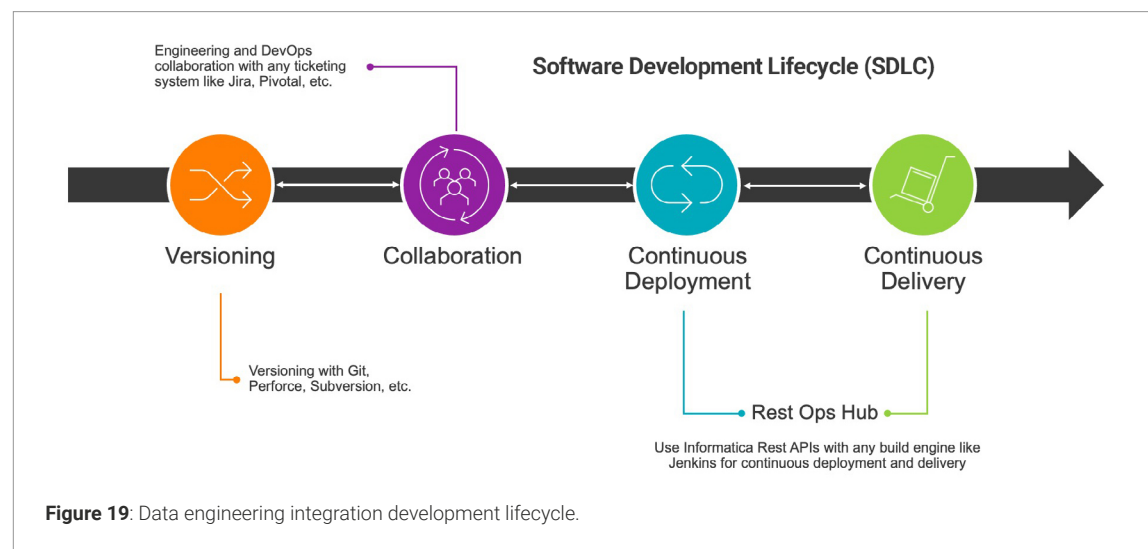
Key CI/CD Considerations for a Data Engineer

Here are some key considerations while implementing CI/CD:

- Versioning is the ability to check-in and check-out data pipelines, roll back to previous versions and trace who has made what changes on data pipelines.
- Collaboration enables multiple users to collaborate and work on the same sets of data pipelines (although not all at the same time).
- Build process flow is the set of steps that enables new/modified data pipelines to be versioned and the right changes to be checked in the code repository.
- Deployment process flow is the set of steps that enables new/modified data pipelines to be deployed to the next environment with the right sanity checks and tests performed on the data pipeline.

DataOps: Building and Expanding Data Pipelines and Analytics

DataOps is a set of practices, processes and technologies that combine an integrated and process-oriented perspective on data with automation and methods from agile software engineering to improve quality, speed and collaboration. DataOps provides a way to operationalize your data platform by extending the concepts of DevOps to the world of data. It also promotes a culture of continuous improvement in data analytics.¹⁸



¹⁸ <http://ceur-ws.org/Vol-2191/paper13.pdf>

Operationalization in Data Engineering (continued)

Three key principles of DataOps include:

1. **Continuous integration:** This process relates to how data engineers integrate, prepare, cleanse, master and release new data sources and data pipelines in a sustainable, automated way. Data engineers can get off to a quick start when data scientists, data analysts and data stewards collaborate using data catalog and data prep tools powered by AI/ML. With these tools they are able to automate **data discovery** and curation, facilitate search, recommend transformations and auto-provision data and data pipeline specifications. With streaming and CDC technology, data engineers can turn these data pipelines into real-time streams that feed predictive analytic algorithms.
2. **Continuous delivery:** This stage is about operationalizing data governance across your enterprise so that all of your consuming applications are using high-quality data. Data governance democratizes and frees your data so that data delivered across the enterprise is trusted, secured, protected and compliant with policies. In this stage, data curation is ongoing, and data is delivered in a collaborative fashion among all stakeholders (e.g., data engineers, data scientists and analysts, data stewards, data governance professionals, etc.).

For example, data scientists can rapidly iterate through the design and validation of predictive analytic models when data they can trust is easy to find. During development, testing and AI model training, it's critical to ensure data quality rules and data masking are applied in accordance with data governance policies so that analytic algorithms and ML models deliver positive business outcomes. Only a unified and intelligent data platform that integrates data governance with data cataloging, **data quality** and data privacy enforces that virtually all data is trusted and protected as it moves throughout the enterprise.

During development, testing and AI model training, it's critical to ensure data quality rules and data masking are applied in accordance with data governance policies so that analytic algorithms and ML models deliver positive business outcomes.

Operationalization in Data Engineering (continued)

- 3. Continuous deployment:** At this stage, you're enabling self-service and making trusted data available to a wide variety of users across your organization. Now every change that passes all stages of your data pipeline development is released to the consuming applications used by analysts and line of business users.

Data-driven apps have become mission critical to many business functions such as customer service, sales, ecommerce, fraud detection, supply chain management and more. This means the business expects faster access to fresh data. This is best achieved with scale-out and microservices-based architectures often deployed in the cloud for agility and flexibility.

AI and ML play a critical role in monitoring and managing the data pipelines, so they continuously run and are optimized for performance and capacity utilization. As DevOps deals with the delivery of software products, DataOps automates data orchestration by delivering data across an organization.

FinOps: The Model for Optimizing, Governing and Controlling Cloud Costs

As organizations continue to modernize to the cloud, finance and operations practitioners keep looking for recommendations from data engineers on cost optimization parameters. Data engineering teams play an active role in optimizing the value of cloud services to the enterprise. Before going further, let's look at the definition of FinOps.

¹⁹ <https://www.finops.org/introduction/what-is-finops/>

"FinOps is an evolving cloud financial management discipline and cultural practice that enables organizations to get maximum business value by helping engineering, finance, technology and business teams to collaborate on data-driven spending decisions."¹⁹

Having access to a large variety of cloud service offerings makes it difficult to choose a cost-effective model. At the same time, cloud service choices need to be optimized to meet the business need during the cloud modernization process. Data engineers can help FinOps practitioners design a **cost-optimized model** that meets business needs.

To help select the right infrastructure, many cloud service providers (CSPs) offer pricing calculators to compare and estimate costs based on specific infrastructure regions and services. To support FinOps effectively, data engineers should be well versed with various cost models so they can clearly articulate cost implications. Now let's look at how FinOps impacts a data engineer's design decision.

The Impact of FinOps on Design Decisions

Price is a key factor when designing data engineering processes. Data engineers should carefully consider the following factors to save costs and avoid future changes in design/ infrastructure components due to cost overruns:

Operationalization in Data Engineering (continued)

- **Track and tune data pipelines:** Track the cost per workload/jobs/project and list processes that incur more cost and try to tune them based on AI-powered cost-based calculators. In addition, enhance existing processes with automatic tuning capabilities based on load during specific date and time. This will be an iterative process rather than a one-time exercise.
- **Runtime and pipeline deployment:** When deploying a data pipeline, you should consider elastic and serverless deployments first. Only consider the alternatives if serverless does not meet your demands. Serverless helps organizations save on overhead costs since they no longer need to pay for an idle infrastructure. Instead, you only pay when a job is executed. Serverless also offers auto tuning, auto scaling and high availability, all without requiring a dedicated administrator to manage the environment.

Data engineers have come a long way from provisioning infrastructure and resources on a yearly basis to auto allocation based on demand and AI recommendations. With on-demand data processing you only pay when your elastic cloud data integration is in use. You don't need to pre-allocate resources and pay for idle time.

Elastic cloud data integration should be flexible enough to support different data integration patterns: from ETL to ELT and from data warehousing to **data fabric**.

Once you build mappings using a data engineering integration solution, you must get an option to run mappings in an existing cluster for on-premises deployment or serverless using the cluster auto-deployment option.

The role of the data engineer is to give clean and meaningful data to the business. Today a large portion of the data engineer's time is spent on tuning jobs or answering questions due to production down time.

Now let's take a look at how AI, ML and data science are used in real-world use cases.

"It would be incredibly difficult to tie together all our different datasets and get them into a cloud data lake without Informatica Cloud Data Integration. It's a huge step in our digital journey to build a new kind of health system."

Bruno Moura

Data Engineering and
Analytics Manager
SulAmerica



Case Studies | AI, ML and Data Science in Action

Accelerating the Credit Approval Process by 70% With Enhanced Analytics

Banco ABC Brasil offers deposit and commercial banking services in Brazil and the Cayman Islands. It also advises on underwriting activities and mergers and acquisitions, offers treasury services and provides international lines of credit.

Goal: Banco ABC Brasil wanted to deliver a better experience to clients, while improving data analytics capabilities and accelerating the credit application process.

Solution: The financial institution leveraged Informatica data integration and data cataloging capabilities with Google Cloud. This helped them build enhanced analytical models to enable better, quicker decision-making that supports the business and their customers. In addition, AI-powered Informatica cloud application integration capabilities automated Banco ABC Brasil's credit analysis process.

Results: The automated processes reduced predictive model design and maintenance time by up to 70%. This sharpened the accuracy of predictive models and insights with trusted, validated data. Banco ABC Brasil also enabled analysts to build predictive models 50% faster, accelerating credit application decisions by 70%.

"By accelerating our digital transformation with artificial intelligence and machine learning, we can use our data assets to the bank's competitive advantage."

Rafael Kataoka

Chief Data & Analytics Officer and Chief Privacy Officer, Banco ABC Brasil



Using Data to Improve Operations and Maintain a Competitive Edge

Maersk Line is an international container shipping company and the largest operating subsidiary of A.P. Møller – Mærsk A/S, a Danish business conglomerate with activities in the transport, logistics and energy sectors. Maersk has been the largest container ship and supply vessel operator in the world since 1996. With 374 offices in 116 countries, the company manages approximately 900 vessels, providing global coverage for shipping everything under the sun.

Goal: Maersk wanted to modernize their infrastructure with a cloud data lake to feed predictive analytics and ML models with timely, high-quality data from diverse sources, including telematics data from ships at sea

Solution: The shipping company uses Informatica data quality and master data management capabilities to integrate, verify and deduplicate data from ships and legacy systems.

Results: Maersk is now able to better track customer cargos, reduce costs and maintain its competitive edge in the global container shipping market.

"There are a lot of disruptors coming into freight forwarding, and we want to use AI and machine learning to help us maintain our leadership position in the industry. We've also seen that our users want to get their hands on the data, and they want to be able to generate their own insights without relying on IT. We want to democratize the data, get it out in front of people, and make it easy for them to discover and govern."

David Falder

Senior Technical Specialist, Maersk

How Informatica Intelligent Data Management Cloud (IDMC) Supports Modern Data Engineering

As technological advancements make overall data processing simpler, data requirements are getting more complex. So how can you as a data engineer help organizations grow fast, drive down costs and develop big ideas? And how can you do it while also quickly adapting to shifting business demands?

To be successful, you must take advantage of modern technologies, such as data mesh and data fabric, to make processes scalable, reusable and adaptable. Having tools that leverage AI and ML can help you automate and simplify tasks related to data management – across data discovery, data integration (e.g., ETL/ELT), data quality, data governance and data matching and enrichment for the mastering of data.

To do this well, you need a holistic and connected data management strategy that makes room for emerging cloud technologies and allows you to make informed decisions about your business much faster. This approach can separate the leaders from the laggards, empowering organizations to improve customer experience and get products to market more quickly.

Informatica offers a comprehensive, end-to-end data management platform with the **Intelligent Data Management Cloud™ (IDMC)**. Its capabilities are designed to meet virtually any data management need

of data engineers and simplify your tasks through automation using **CLAIRE®**, its AI-driven engine. IDMC can help empower you, as a data engineer, to:

- Build a foundation for analytics, AI, ML and data science initiatives that support large volumes of batch and streaming data, whether it is structured, unstructured or semi-structured
- Improve productivity by simplifying the development, deployment, tuning and maintenance of complex data pipelines
- Operationalize virtually any AI/ML model to put AI into action
- Support modern frameworks and data engineering trends like data observability, data mesh, data fabric, modern data stack, lakehouse and modernizing to super clouds
- Gain serverless and elastic scale to meet business demands and optimize, govern and control cloud costs with FinOps
- Choose virtually any cloud services at any time across IDMC as your requirements change with **Informatica Processing Units (IPU)**
- Solve data engineering problems that inform critical business decisions and accelerate innovation

How Informatica Intelligent Data Management Cloud (IDMC) Supports Modern Data Engineering (continued)

To continue to effectively convert the ongoing reams of data into meaningful insights, organizations need an intelligent data management cloud that is scalable and interoperable to meet modern data engineering needs. Informatica is a **leader in data engineering**. In fact, we have helped over 5,000 organizations succeed with their data-driven initiatives using IDMC.

Leverage the knowledge, skills and tools explained in this eBook to be an exceptional data engineer who drives real change across your organization. And get the data management capabilities you need to take advantage of today's — and tomorrow's — innovations.

Ready to get started on your data engineering journey with Informatica? Try the industry's only free, AI-powered solution to easily load, transform and integrate data with **Cloud Data Integration-Free**.

Join the **Informatica Data Engineer Central Community**

to connect, learn and collaborate with like-minded professionals.



As a member, get access to exclusive resources, best practices and events, all designed to help you advance your career and expand your knowledge.

SIGN UP NOW

About Us

Informatica (NYSE: INFA) brings data and AI to life by empowering businesses to realize the transformative power of their most critical assets. When properly unlocked, data becomes a living and trusted resource that is democratized across your organization, turning chaos into clarity. Through the Informatica Intelligent Data Management Cloud™, companies are breathing life into their data to drive bigger ideas, create improved processes, and reduce costs. Powered by CLAIRE®, our AI engine, it's the only cloud dedicated to managing data of any type, pattern, complexity, or workload across any location — all on a single platform.

IN19-4545-0824

© Copyright Informatica LLC 2024. Informatica and the Informatica logo are trademarks or registered trademarks of Informatica LLC in the United States and other countries. A current list of Informatica trademarks is available on the web at <https://www.informatica.com/trademarks.html>. Other company and product names may be trade names or trademarks of their respective owners. The information in this documentation is subject to change without notice and provided "AS IS" without warranty of any kind, express or implied.

[informatica.com](https://www.informatica.com)

Where data & AI come to



Worldwide Headquarters
2100 Seaport Blvd,
Redwood City, CA 94063, USA
Phone: 650.385.5000
Fax: 650.385.5500
Toll-free in the US: 1.800.653.3871

[informatica.com](https://www.informatica.com)
[linkedin.com/company/informatica](https://www.linkedin.com/company/informatica)
twitter.com/Informatica

CONTACT US