IMMUTA®

WHITE PAPER

# Best Practices for Securing Sensitive Data

## A Guide for Teams of Any Data Management Maturity

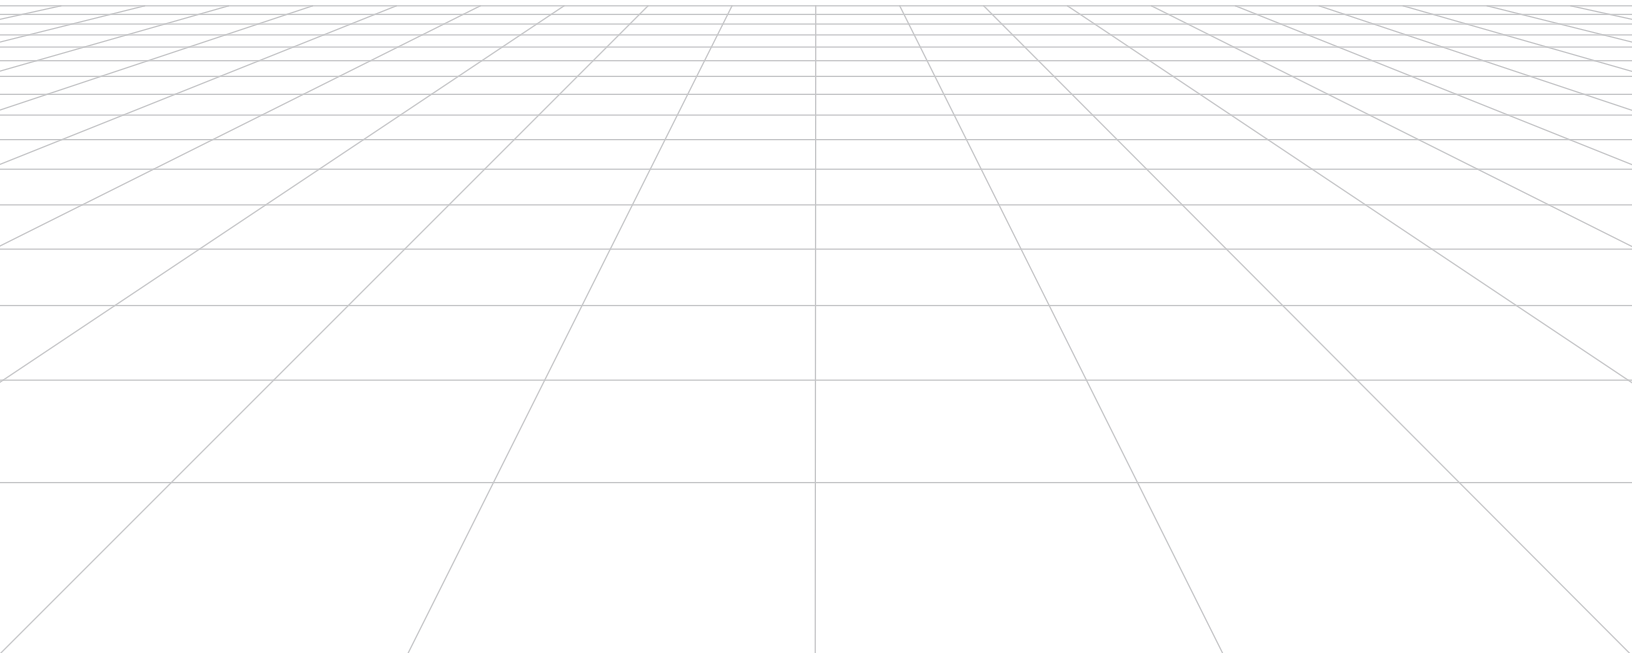

# Table of Contents

# Introduction

Most data-driven organizations have embarked on data and analytics modernization initiatives in an effort to achieve business-driving benefits, such as the ability to leverage their data assets for richer, proactive intelligence and to drive competitive advantages.

Such initiatives are possible only when organizations can unlock the vast troves of multi-structured data within the corporate firewall or external sources. Organizations have become increasingly comfortable exposing internal data in a secure and governed manner, and are now ready to experiment with more advanced use cases like data sharing and data mesh.

However, there is a problem. The data and analytics industry has been mainly focused on enabling access to the data through cloud data lakes, warehouses, and analytical engines, while overlooking essential data governance disciplines, such as data quality and protection. With new, stringent privacy regulations now shifting focus to safeguarding personal data, data teams must revisit their practices and frameworks to ensure security and compliance.

A primary area of focus is on legacy data governance platforms, many of which were overly restrictive. They prevented access to entire segments of data, which substantially reduced data utility. On the other hand, the lack of such a platform often made data accessibility overly permissive. Newer products provide the ability to write and automate policies for fine-grained access control at the row-, column-, and cell-level so organizations can accelerate data access, but in a secure and governed manner. These policies determine who is authorized to see sensitive data, including personally identifiable information (PII), protected health information (PHI), non-personal commercial data, and contractually protected or third-party data.

Examples of "fine-grained" access policies include:

- **Limiting access to sensitive data attributes**. For example, access to credit card numbers in clear text should be restricted to not display the column at all, randomize the content, obfuscate it completely, or show only a partial string. This data protection is called column-level security (CLS).
- **Limiting access to the entire record**. This data access restriction is commonly used to satisfy data sovereignty or data residency requirements where row access is restricted to only rows that fall in the region(s) to which the user has been granted access permissions. It is also known as row-level security (RLS).

The data and analytics industry is adopting the term "modern data stack" to categorize an architecture comprising best-of-breed tools meant to ingest data from producers and prepare it for data consumers. As the very definition is amorphous and applied inconsistently, data access and governance capabilities have historically not been well defined. There is no single, prescribed manner in which to secure and protect sensitive data, yet data security and access management are critical to an organization's investment in cloud data platforms.

This document aims to share best practices to enable trusted access to your corporate data. These tips vary based on the maturity of your organization's existing data management practices, specifically in the realms of data governance and access. Let's first establish a baseline for your organization's data protection maturity.

# Data Management Maturity Models

Several data management maturity models have been developed over time to help organizations assess their current state and identify gaps. These include Enterprise Data Management (EDM) Council's Data Management Capability Assessment Model (DCAM), and Data Management (DAMA) International's Data Management Book of Knowledge (DMBOK), among others. For this document, we'll focus on a subset of data management: an organization's ability to identify, secure, and protect sensitive or contractual data elements.

Maturity models help establish metrics to measure the effectiveness of current processes and establish progressive improvements. To achieve successful modernization initiatives like cloud data adoption, organizations must model best practices based on the maturity of their approaches to securing sensitive data. This graduated approach leads to continuous improvements without overwhelming data producers, enablers, and consumers.

The figure below describes the three stages of data access management maturity.

## Data Access Management Maturity Levels

| Nascent | Emerging | Mature |
|---|---|---|
| Use cases unidentified; Lacking a data access strategy, including RACI | Awareness of data's strategic value with basic data management quidelines | Some data privacy regulation, contractual data protection implemented |
| No data ownership and limited information on sensitive data location | Limited metadata management capabilities and awareness of sensitive data location | Risk assessment of data quality and trust exists, but not always automated |
| Business lacks policies, data usage knowledge, data governance, catalog | Basic policies and role-based access control but not integrated with IAM | Data access management is integrated with the rest of data governance |
| Technical team lacks lineage from producer to consumer and impact of data masking downstream | Technical teams lead data authorization with little business inputs. Deployment is ad hoc and departmental. | Technical and business teams are integrated. |
| Processes are manual, time-consuming, brittle (e.g., multiple copies of data) | Department-level processes are not known across the organization | Well documented, centralized policies, and processes to manage sensitive data |
| Organization lacks technical roles and skills, is weak in security and policy skills | Skills such as Snowflake or Databricks admin exist but not focused on security | Policy group defines corporate standards and users are trained in security |

The maturity model lays a blueprint for continous improvements.
To graduate to the next level, you must meet the previous levels' requirements.

Figure 1. Three stages of data access management maturity

Best practices of the maturity model fall into three categories: nascent, emerging, and mature stages.

- **Nascent**
  Organizations at the nascent stage do not have a data security and governance program suitable for modern data stacks in place, and therefore typically have not documented where sensitive data is within their cloud data architecture. Their access management processes are typically manual and create bottlenecks. Organizations in this stage also lack a formal data protection strategy, including a process for classification and tagging across cloud data assets.

- **Emerging**
  Organizations at the emerging maturity stage have some data protection mechanisms. Often, these are custom-developed and maintained, requiring considerable effort to keep up-to-date and compliant with external regulations. Other organizations have deployed one or more data access platforms, like Immuta. In either case, the scope of data protection is at the department level and is not cohesive. At this stage, organizations lack automated policy enforcement.

- **Mature**
  Organizations at the mature stage have an advanced data protection strategy and deployment process, as well as a strong understanding of business needs and compliance requirements. However, they are looking to scale their modern analytical architectures to innovate in the cloud, and are continuously refining and modernizing their data access strategy.

# Best Practices for Securing Sensitive Data

Data protection deployment should be comprehensive in order to ensure that there is no backdoor access to sensitive data. However, "big bang" approaches are hard to execute and rarely succeed. This is because while protecting sensitive data can be time-consuming and expensive, data utility for permissible consumption should not be compromised.

Following proven best practices ensures organizations derive the highest value possible from their software tools to safeguard sensitive data, while minimizing governance overhead. Best practices also provide a structured approach for organizations looking to achieve a successful deployment of their data access platform.

The next section describes best practices for each of the maturity stages. One common thread, however, is that every customer must *first define their ideal use cases and success criteria*. An example of a use case is enabling dynamic, self-service access and authorization to sensitive data in operational or analytical applications, like Snowflake or Databricks.

Use cases must be business strategy-driven and stated in business terms. They should not be technology-driven. The top data access use cases are:

- **Cloud Modernization**
  Most organizations are adopting cloud data platforms, or are in the process of expanding their ecosystem. However, it's important to not start a cloud migration process until you have identified where your sensitive data is located and how it is used. This applies even for lift-and-shift IaaS workloads. Compared to on-premises deployments, cloud deployments work on a "shared security model" that puts a greater responsibility on the users to ensure data protection. Having a grasp on sensitive data prior to modernization initiatives reduces the risk of reactionary governance efforts or data slipping through the cracks.
- **Data Privacy**
  Data privacy refers to approaches used to control and manage access to sensitive data, including PII and PHI. Such approaches include techniques like differential privacy, k-anonymization, and pseudonymization, many of which are mandated by data rules and regulations. In addition to ensuring data privacy controls satisfy legal requirements, it's important to enforce them consistently across all data platforms and consumption approaches. For example, a business user may use Microsoft Excel to access data, while an analyst may use SQL, and a data scientist may use a Python code from a notebook – but the same data privacy controls should be in place for each user.

- **Data Sharing**
  The interconnectedness of today's marketplace makes data sharing essential for organizations of every size and across all industries. Whether sharing data internally across business units or externally, such as with third parties, the need to ensure it is exchanged securely is a growing priority for many data teams. Data sharing is also used to assure the terms of data licensing and data use agreements are adhered to, and ultimately enables organizations to monetize their data products through secure data exchange platforms.

- **Regulatory Compliance & Auditing**
  Any sensitive or private data must be regulated – it is non-negotiable. Achieving compliance with external compliance laws and regulations, as well as with internal guidelines and standards, involves knowing what type of sensitive data you have, where your data consumers are accessing that data, and the specific requirements that apply to your data, such as data residency or the right to be forgotten. The legal team should align with the business team responsible for authoring policies, and with the data platform team responsible for implementing data policies.

- **Data Access Control**
  Controlling who can access data to what data in a rapidly evolving ecosystem isn't easy – but it is necessary. As with data privacy, data access control becomes more complex when you consider the array of users, technologies, and regulations involved, and the need for consistent policy enforcement to manage any combination of access requests across any compute platform. When organizations' data needs grow, their access controls must scale proportionately – ideally without putting an additional tax on data teams responsible for policy implementation and data access management.

- **Data Security**
  Data security is designed to mitigate the threat of unauthorized data access. As data-driven organizations adopt increasingly decentralized and flexible cloud data architectures, such as a data mesh or data lakehouse, data security has become an even more critical use case. To reduce the threat of data leaks, breaches, or losses from both internal and external sources, and to reap the full benefits of innovative architecture frameworks, data teams must have a strategy and resources for maintaining data security.

Understanding your use cases and having them inform best practices is an excellent starting point for embarking on a data access and security journey. This implementation approach provides an agile roadmap to demonstrate incremental success.

# Best Practices for the Nascent Maturity Stage

A key best practice for the nascent maturity stage is to get the user comfortable with the data access platform's capabilities while demonstrating its effectiveness, such as satisfying the applicable compliance regulations like the GDPR or HIPAA. As organizations mature, they will likely develop more sophisticated use cases, like data sharing.

Users new to the data access control and security space should adopt a 'systems thinking' approach comprising business, technology, and process pillars. Best practices enable a structured implementation of the first data access use case in the most cost effective and agile manner. An iterative and phased process allows organizations to start small and show value quickly.

The figure below lists best practices for organizations that are in the nascent stage of the maturity model.

## Nascent Use Case-Based Best Practices

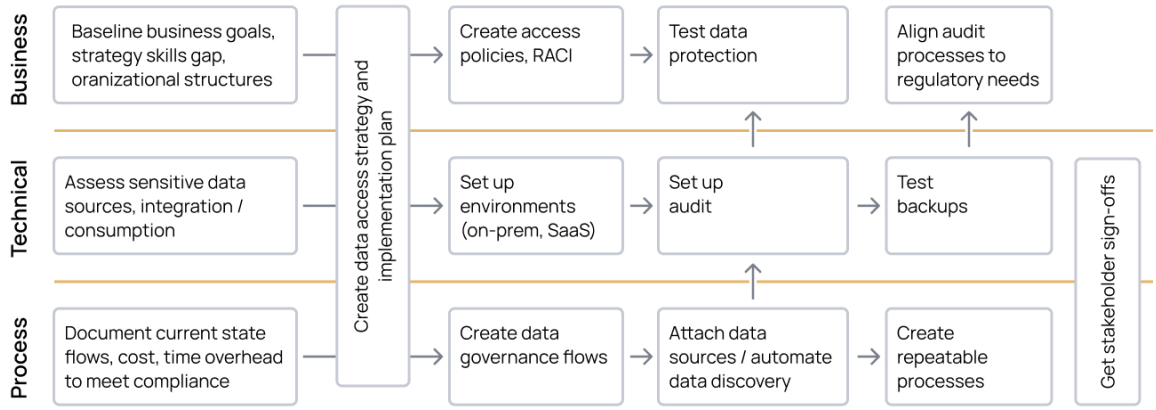


Figure 2. Nascent maturity stage best practices introduce the concept of data access control and data protection within an organization.

Once the starter use case is defined, assess the current state across all the pillars, and embark on developing an access strategy that aligns with the overall business goals, imperatives, and strategy. The rationale of the access strategy and implementation plan is to avoid scope creep, get business and IT stakeholders to sign-off on deliverables, and align on common business goals. This best practice identifies whether the selected use case is a good business and technical (e.g. data sources) fit for the first data protection iteration.

# Business Best Practices

Implementation of data policies varies across startups, small- or medium-sized businesses (SMBs), and large enterprises. Startups have small team sizes that share multiple responsibilities, while larger companies have dedicated teams and are often process-oriented. As a result, the implementation plans differ.

The business best practices for the nascent maturity stage include:

- The *data owners* should **author their first set of policies** as they understand the business domains and their needs. However, as data is a corporate asset and is often shared among domains, the task of defining corporate policies should be centralized under a shared team, most likely run by the chief data officer (CDO). The modern best practice for authoring policies is to use a low or a no-code product that generates easy-to-read and declarative policies.
- The *CDO or data platform owner* should **develop a RACI** chart with all stakeholders. Start with a workshop to align all stakeholders and assign clear roles and responsibilities, such as data owners and producers, data consumers, and data enablers and stewards. This task should cover business, infrastructure, data, legal, and security teams. The clarity of roles and responsibilities helps to ensure faster, more successful deployment of data protection initiatives.
- The *CDO or data platform owner* should **build a center of excellence (CoE)** to develop standards, catalog process changes, and train data professionals such as owners, stewards, and analysts. If there is already a data security CoE, then build a data access "special interest group" (SIG) within the CoE.

The best practices in this section focus on building, testing, and aligning data access policies to the goals of the stated use cases. These policies should be decoupled from the underlying technologies that are used to enforce them. This separation of policy authoring from policy enforcement is important to making policies fungible and reusable.

# Technical Best Practices

Technical best practices start with identification of critical and high-impact data sources, where sensitive data must be protected prior to analytics. The best practices vary depending upon the deployment mode of data access control products, such as SaaS, on-premises, or multi-cloud. On-premises deployments involve a long training cycle, which SaaS products typically do not require.

The technical best practices for the nascent maturity stage include:

- The *data platform architect or the CIO team* should **develop an end-to-end data access control architecture** that is comprehensive and has a centralized authorization subsystem independent of the execution engines. Fine-grained access controls, such as attribute-based access control (ABAC), dynamic policy authoring, and data monitoring, are some of the essential components of such an architecture.
- The *data platform owner* should **determine the technology stack** needed to achieve the goal of making the policies executable. In the old school approach, policies were documented in prose in a standalone text document and archived on a document store. It was easy for these documents to get ignored or become quickly outdated as the shape of data evolved. In today's era of data use, policy implementation must be dynamic, flexible, and scalable.

- The *data platform engineer* should set up crawlers to **discover and tag sensitive data** in data sources, transformations engines, and other persistence layers in the architecture relative to the scope of the identified use case. The reason for taking a holistic view is to prevent de-identified data elements from becoming re-identified at various points in the pipeline.
- The *data platform engineer* should **enable development, test, and production environments**, and engage with the DevOps team to **integrate with the continuous integration and delivery (CI/CD)** tools.
- The *data platform engineer* should enable **privacy enhancing techniques (PETs) and attribute-based access control (ABAC)** to ensure policies achieve column-level security through processes like data masking, encryption, and tokenization.

While the business and the technical teams work in tandem on the two sides of the coin – policy definition and policy enforcement – many other teams are involved in the overall operationalization of the data access control.

# Process Best Practices

Securing, governing, and protecting access to sensitive data impacts the current state processes. The tool used to enable data access control is a change agent. Unfettered access to sensitive data will be curtailed, which may not always be a welcome change inside organizations. Hence, users should be aware of the process changes. The best practices in this area are:

- The *data platform owner* should **start every project with a workshop** that defines the baseline and current state processes, identifies gaps, and documents proposed future state processes.
- The *data platform architect* should be tasked with **identifying related data governance gaps**, such as data quality programs, master data management initiatives, and data catalogs. These programs may be prerequisites to starting data protection efforts.
- The *data platform engineer* should be responsible for **implementing new secure processes iteratively** team-by-team, use case by use case, or platform-by-platform. The big bang approach does not work for organizations that are at the nascent maturity stage.
- The *data platform engineer* should **incorporate any relevant technical enhancements**, such as CI/CD, that will ensure all aspects of the data stack function cohesively and efficiently.
- The *data platform engineer* should **create a repeatable, scalable process** by introducing automation into data access workflows.

When best practices for all the three of the business, technical, and process pillars are applied, organizations graduate from the nascent to the emerging maturity stage.

# Best Practices for the Emerging Maturity Stage

At this maturity stage, organizations should have basic data access controls for key data sources. Their goal is now to apply data access control policies consistently across end-to-end pipelines. This requires formalizing cross-functional standards, applying lessons learned across teams, and ensuring data access control processes are repeatable and scalable.

The figure below lists best practices for organizations that are in the emerging maturity stage.

## Emerging Use Case-Based Best Practices

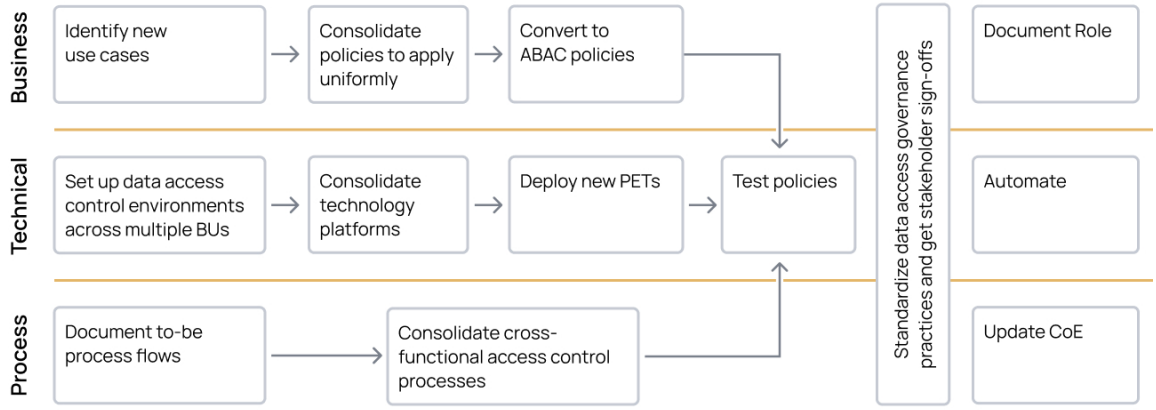


Figure 3. Emerging maturity stage best practices extend data access control and protection from a single department to cross-functional using standardized policies, products, and processes.

An outcome of the best practices for the emerging maturity stage should be that data consumers' access policies and permissions are applied uniformly, regardless of where the data is stored or processed.

# Business Best Practices

One of businesses' biggest pain points is writing redundant policies. Not only does this reduce the data team's productivity, but it also contributes to a sprawling and complex web of policies that may lack systematic logic. Large or growing enterprises may experience the impacts of repetitious policies more acutely than smaller organizations, but the bottom line is that any unnecessary policies create additional complexity, which should be avoided regardless of size.

Legacy access control frameworks, namely role-based access control (RBAC) are a major driver of policy redundancy. RBAC is static in nature, and determines access permissions based on user roles. This means that as new users or data sources are added and policy needs evolve, new policies must be created to absorb those changes. As a result, data teams are left to manage substantial role bloat, which can become unmanageable and unscalable – data is likely to either be locked down and overly restricted, or exposed due to slipping through the cracks.

Dynamic access controls are one critical way to avoid this, as well see in the best practices below. Enforcing policies at query time based on attributes about users, data objects, intended actions, and the data environment removes the need to predetermine and manually maintain policies. As a result, data teams can reduce the number of policies they must manage by 75x and scale policy enforcement with no additional overhead required.

The business best practices for the emerging maturity stage include:

- The *CDO or data platform owner* should **define common standards** that can be used to onboard new use cases efficiently in order to avoid potential policy bloat.
- The *data platform owner or architect* should **update the current state data access strategy** based on lessons learned from the previous data access governance approaches. The updated strategy should also quantify benefits in order to help make a business case for the expansion of the data access control initiative.
- The *data owner* should **identify new opportunities** to expand access control, with a focus on critical data elements (CDE) or data sources that deliver the highest value and are aligned to the business imperatives. Often, organizations tend to look at the "low hanging fruit" use cases, but these may not demonstrate the investment into the program.
- The *data platform engineer* should **develop new ABAC policies** and **convert any existing RBAC** policies into ABAC. ABAC offers the most scalable option, especially as the data and its usage change dynamically.

# Technical Best Practices

If the business teams are loath to write redundant policies, data engineers are also reluctant to manually deploy those policies across multiple execution infrastructures and keep them in sync. With data engineering skills in short supply, a best practice is to install a universal policy enforcement engine.

Data engineers must also work with Governance, Risk, and Compliance (GRC) stakeholders, such as data governors, to ensure that all platforms and data sources in scope are protected. This includes modern products as well as legacy systems, since most organizations' data sources span hybrid multi-cloud locations.

The technical best practices for the emerging maturity stage include:

- The *data platform architect* should **evaluate and consolidate existing technology platforms** in the context of the access control strategy.
- The *data platform engineer* should **optimize and automate** existing data access control implementation. The goal is to standardize data access control into common tools and approaches to achieve economies of scale and better resource utilization.
- The *data platform architect* should **add sophisticated privacy enhancing technologies (PETs)** to support the additional use cases. As mentioned in the nascent maturity stage, basic data access control relies on PETs like encryption, masking, and tokenization. In the emerging stage, advanced PETs may include differential privacy, k-anonymization, and randomized response.
- The *compliance officer* should **test policies** to ensure they are authored and enforced in compliance with all pertinent data use rules and regulations.

# Process Best Practices

With regard to process-oriented best practices at the emerging maturity stage, consolidating department-centric data access process flows into a standard enterprise-wide process flow enables centralized policy management and comprehensive audit logging. Simplifying and standardizing these processes across teams is a critical step in scaling secure self-service data use without adding complexity.

Best practices in this area include:

- The *data platform architect* should **document cross-functional data access control flows** with the goal of developing standards that reduce the costs of expanding the initiative to other data sources. This will help identify areas of opportunity and simplify consistent access control implementation throughout the organization.
- The *data platform architect* should work with the *data platform owner* to **socialize these access control flows across teams** to ensure data users across lines of business are adhering to the same standards.
- The *data platform engineer or architect* should **enhance the data access control center of excellence** with the goals of building common processes, establishing an environment for departments to prototype data access solutions, and training the organization on best practices.

Once these objectives are achieved, organizations are able to graduate to the final maturity stage.

# Best Practices for the Mature Stage

The goal of the mature stage is to extend sensitive data protection beyond internal departments and business units, and into all data sources, including external, third-party data sources.

Most organizations leverage data from multiple sources and need to adhere to contractual data sharing and data use agreements. Therefore, best practices in the mature stage are meant to help secure access to all internal and external data sources. This includes expanding data entitlements to third-party data sources, such as Bloomberg and FactSet. Points of guidance for the mature stage pertain to deploying end-to-end modern data stacks.

## Mature Use Case-Based Best Practices

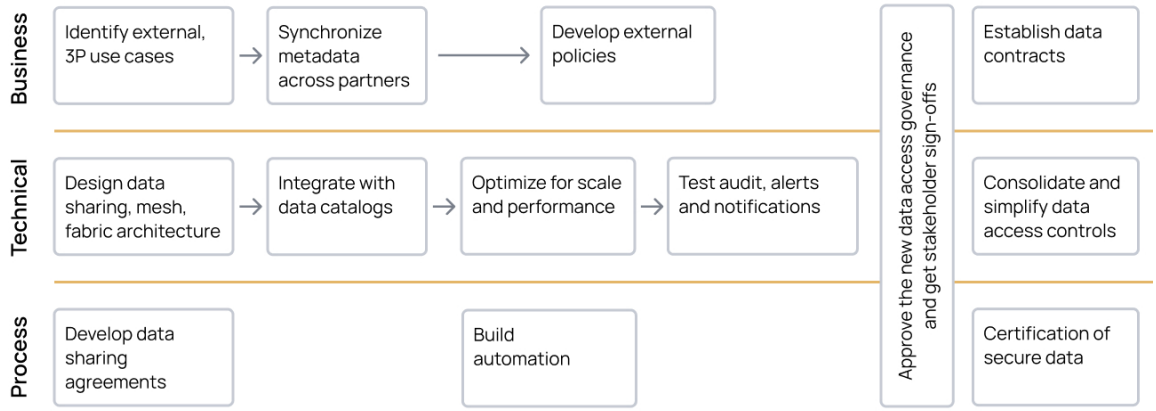


Figure 4. Mature stage best practices extend internal data access control best practices to the wider ecosystem, comprising third-party data sources in more advanced architectures, including data mesh and data fabric.

The best practices in this stage are particularly pertinent as organizations increasingly decentralize modern architectures using techniques such as data fabrics and data mesh. These fresh approaches have alleviated various bottlenecks in the data supply chain, such as constrained data engineering resources, and have afforded more accountability to data quality issues. However, this 'distributed' approach adds to overall complexity, increasing the need for a more advanced level of data usage and access control.

# Business Best Practices

The primary focus for the business functions in mature organizations is to extend the data access and security strategy into all internal and external sources in order to enable consistent data access policy enforcement. These best practices will open data pipelines and allow a more efficient flow of data both within and outside of the organization.

The business best practices for the mature stage are:

- The *data platform engineer* should work with data owners to **survey and identify all of the data sources and data consumption targets** within an organization, including use cases across lines of business. Oftentimes, this information may be available in a configuration management database (CMDB).
- The *data platform engineer* should **synchronize metadata across platforms and partners** in order to collect external user information that can be incorporated into data access policies.
- The *data platform architect* should **develop an actionable timeline** to secure and protect sensitive data in all sources and targets. The timeline should be comprehensive and include all data sources.
- The *data platform architect* should **create external policies** that incorporate partner metadata and can be implemented across all data sources.
- The *privacy or security officer* should **author data contracts and data usage agreements** for third-party data sources. This step may require collaboration with additional legal stakeholders.

By the time organizations reach the mature stage, they should be able to ensure that the policies are compliant with data use agreements, in addition to any relevant data privacy use cases.

# Technical Best Practices

As global regions and countries introduce new privacy legislations, businesses should enable policy portability and enforcement across distributed architectures using multiple compute engines.

Metadata regarding sensitive data elements is usually kept in metadata catalogs, so the data engineering team should enable integration across metadata catalogs, data orchestration, and other components of the modern data stack. Another focus of this stage is ensuring enterprise-level scalability, performance, availability, and reliability of the data access platform.

The technical best practices for the mature stage are:

- The *data platform architect* should **design data access controls in data mesh, data fabric, and data sharing** architectures, in accordance with the data platform owner's data infrastructure roadmap. Data access control should be part of the design process, and not an afterthought.
- The *data platform engineer* should **develop a bi-directional integration with data catalog(s)** by ingesting metadata tags from existing catalogs into the data access platform, which can then be absorbed into access control policies.

- The *data platform architect* should **identify opportunities to consolidate and simplify data access controls** in an ongoing effort to remove barriers, optimize implementation, and scale for performance.
- The *data platform engineer* should **set up automated auditing, alerts, notifications, and logging** to proactively handle data access issues. Any such issues that arise should be escalated to the compliance officer to be addressed immediately.

# Process Best Practices

Data access control is not just about data security and compliance with data privacy regulations. In our increasingly interconnected data environment, contractual agreements between first- and third-parties sharing data must also be adequately maintained. As this task requires cooperation of multiple organizations, standardizing processes is the most straightforward way to reduce manual effort and improve resource utilization.

The best practices for the process teams in the mature stage are:

- The *data platform architect* should **develop repeatable processes** for every internal department and external organization involved in decentralization or data sharing initiatives. These processes should be communicated to any and all teams that are involved with data, so as to ensure enterprise-wide adoption.
- The *privacy, security, or compliance officer* should **develop a process that "certifies" data**, and attests to the data quality level, its usage agreements, and other legalities on how to ensure sensitive data guidelines are met by all parties.
- The *data platform engineer* should **create an automated onboarding process to add new partners**. This self-service step should reduce the time to value for sharing data from new sources.

Following these best practices will ensure that even organizations with mature data use practices are able to continuously optimize data workflows, which in turn will increase their agility and ability to innovate with data, without compromising security.

# Summary

This white paper has covered the best practices that align with the primary stages of data security and privacy maturity. It is understood that organizations, and their data security and access control needs, are not static – therefore, we can expect their approaches to progress as internal processes, teams, and technologies evolve and improve. Examining the best practices across the business, technical, and process pillars for each of the maturity model stages helps cultivate robust, dynamic, and resilient data access control strategies.

The figure below shows the state of an organization's capabilities before it enters a mature stage and upon exiting it.

## Quantifiable progression of data access control and security

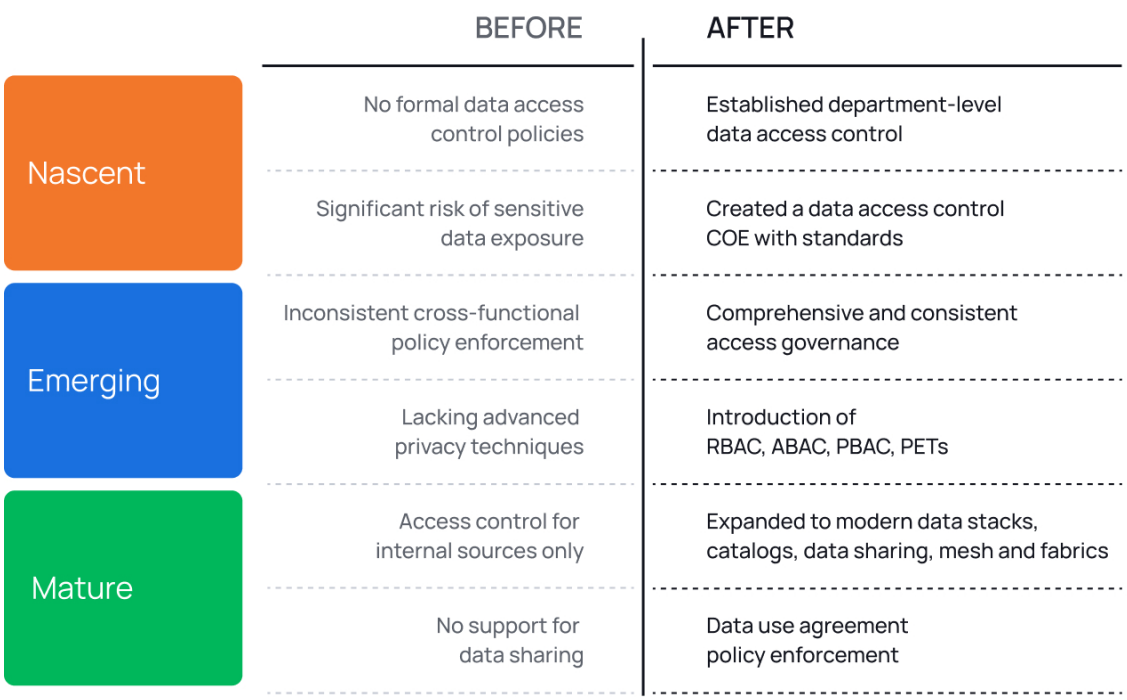


Figure 5. Data access governance is a journey that progresses as an organization matures its internal data access control capabilities.

Data access control as a key facet of data management strategy is becoming mainstream, as the number of use cases that require access control continues to increase. This space came into prominence as organizations began to analyze the increasing amounts of data they were generating, and has gained traction as satisfying data privacy regulations has become an essential requirement. Most recently, enforcing data use agreements on data shared among trusted partners is reinforcing and driving forward the need for this discipline.

Leading cloud data platforms, including Snowflake and Databricks have added native data protection features, providing customers with foundational access controls. These developments validate the need to protect sensitive data. However, most organizations have multiple operational databases, analytical data stores, and compute engines, which can make consistent policy enforcement more complicated and difficult to scale.

In the modern data environment, organizations that wish to scale secure data use need a holistic solution that protects sensitive data across the complex end-to-end data stack. Data access platforms like Immuta separate policy from platform, allowing for centralized policy creation and implementation across systems. Regardless of your current state of maturity, pairing the best practices introduced in this guide with a dynamic data access platform will help simplify sensitive data security and unlock more data for more use cases.

**To see how Immuta can help you achieve your data security and access management goals – regardless of which stage of maturity your organization is in.**

**Schedule a demo with our team today.**

## About Immuta

Immuta is the **market leader in Data Access**, providing data teams one universal platform to control access to analytical data sets in the cloud. Only Immuta can automate access to data by discovering, protecting, and monitoring data. Data-driven organizations around the world trust Immuta to speed time to data, safely share more data with more users, and mitigate the risk of data leaks and breaches. Founded in 2015, Immuta is headquartered in Boston, MA.

IMMUTA®

25 Thomson Place, 4th Floor, Boston, MA 02210 immuta.com (800) 655-0982