

WHITE PAPER

# Implementing & Auditing ICD 501 Using Immuta

Discovery and Dissemination or Retrieval of  
Information within the Intelligence Community

2020

 IMMUTA™

# The Intelligence Reform and Terrorism Prevention Act of 2004 requires Chief Data Officers (CDOs) to strengthen the sharing, integration, and management of information within the US intelligence community (IC).

Intelligence Community Directive (ICD) 501,<sup>1</sup> enacted in 2009, builds on the 2004 requirements by charging IC agencies with the responsibilities to provide, discover, and request information to fulfill their specific missions. ICD 501 goes a step further by requiring that these data sharing practices of providing, discovering, requesting are accomplished through automated means.

This means that the future of secure, efficient data sharing within the IC requires an advanced data access governance solution that can deliver data privacy and speed to data access. The implications of this – an improved ability to accurately detect and warn of threats to the US and its citizens, and to provide better, faster data analytics to the president to inform decision making – are wide reaching and directly impact national security.

Therefore, the IC needs a platform that enables fast, scalable, and secure access to data. Immuta is the automated data governance solution built to accelerate the development of data access and security across multiple data sources and platforms, without copying or moving data. Born out of the IC in 2015, Immuta tackles complex data access governance challenges and enables data teams to harness the value of their data without risking privacy and security.

<sup>1</sup> ICD 501, Office Directive of National Intelligence (DNI), [https://www.dni.gov/files/documents/ICD/ICD\\_501.pdf](https://www.dni.gov/files/documents/ICD/ICD_501.pdf)

# Data Stewardship

To understand how Immuta empowers data teams in the IC to protect and share critical, sensitive data, it's important to discuss the data governance and data consumer dynamic. In large organizations, it can be difficult – if not impossible – for data scientists to access all the data they need. Once they do get access, it is often still difficult to make sure they use the data compliantly.

Immuta aims to solve both problems by providing a single, unified access point for data across an organization and ensuring that all data access controls are dynamically and consistently enforced through the platform. Implemented properly, the Immuta platform can help ensure that only the right people see the right data under the right conditions.

Immuta uses the following titles to designate data use roles across an organization:

- **Data Owners** are responsible for the data that is connected to Immuta and can set policies on their data to restrict who accesses it. Once a data owner connects data to Immuta and creates a data source, they are able to set policies that restrict user access, down to the row- and column-levels. Data owners make their data source public and discoverable, or private so that only they and assigned subscribers know it exists.
- **Data Governors** manage global policies, tags, access controls, and acknowledgment statements to restrict how data is used across multiple projects and data sources. By default, governors can subscribe to data sources; however, this setting can be disabled in the Immuta Configuration, removing their ability to create or subscribe to data sources.
- **Data Users** access the data exposed by data owners for analytic purposes, can browse the Immuta UI for data sources, and connect their third-party data science tools to Immuta. This allows them to use data that has been made available through Immuta to build analytics, data science notebooks, etc.
- **Project Owners** are either data owners who want to use purpose-based access controls to restrict how their data will be utilized, or data users who want to efficiently organize their data sources.
- **System Administrators** manage the permissions, attributes, and groups that attach to each user in Immuta. Application admins manage Immuta's configuration, including connecting to external identity managers and catalogs, enabling or disabling data handlers, adjusting emails and cache settings, generating system API keys, and managing other advanced settings. User admins, on the other hand, manage the permissions, attributes, and groups that attach to individual users. Permissions are only managed locally within Immuta, but groups and attributes can also be derived from user management external frameworks such as Lightweight Directory Access Protocol (LDAP) or Active Directory.

By default, data governors and admins are able to subscribe to data sources; however, this setting can be disabled in the Immuta Configuration, removing their ability to create or subscribe to data sources. Additionally, users can default to being an admin and governor simultaneously, but this setting can also be changed in the Configuration Builder, rendering the roles mutually exclusive.

# Active Data Catalog and Platform Integrations

As the central point for data access, Immuta makes data available, discoverable, and secure, and ensures that policy can be created or changed rapidly. The Immuta active data catalog can act as a standalone data clearinghouse or synchronize with existing enterprise catalogs to integrate data and data owners. This serves as a mechanism to provide access to named users, groups of users, or individuals that request access for a specific purpose.

Immuta integrates with enterprise identity management solutions as well, including the LDAP, Active Directory, and Security Assertion Markup Language (SAML) or OpenID Connect (OIDC) based solutions to ensure that user attributes and statuses are authoritative. By default, all network communications with and within Immuta are encrypted via TLS, meaning data is protected while in transit.

This approach ensures that data is transparently and effortlessly made available to those who need it for legitimate purposes, and is kept out of the hands of those that do not. Additionally, this allows Immuta to build upon work your security team has already done to validate users, protect credentials, and define roles and attributes, without having to start from scratch.

By acting as an active data catalog and integrating with other on-premises and cloud-based platforms, Immuta enables ICD 501's directives that IC data teams provide and discover information through automated means. Sensitive data discovery capabilities also tag incoming sensitive data for human inspection, ensuring that it is only provided to and discovered by the appropriate data consumers.

Keeping a virtual metadata catalog rather than maintaining separate copies of data means Immuta's database is designed to remain small and responsive. Running replicated instances of this internal database allows the catalog to scale in support of any and all platforms an IC agency relies upon for data storage and analysis.

## Fine-Grained Access Control

Immuta implements attribute-based access control (ABAC), rather than traditional role-based access control (RBAC), to simplify policy orchestration and management across a wide variety of data sources. With RBAC, administrators "implicitly determine what the users will have access to by adding them to a role. Then... explicitly determine the privilege associated with each role."<sup>2</sup>

Most databases implement the RBAC approach, primarily because it simplifies the data access process — administrators do not have to add users one by one — and access patterns are relatively easy to establish. However, as the number of users and roles grows, RBAC is more difficult to manage for data teams, who are responsible for explicitly predetermining every possible role combination, as well as manually administer and manage each combination. Moreover, these processes are often done through antiquated methods, like spreadsheets, or systems requiring in-depth software development and code. So, despite its relative ease to establish, RBAC creates more backend work that quickly becomes difficult to manage efficiently.

<sup>2</sup> How Privacy Killed RBAC, by Steve Touw, <https://medium.com/immuta-engineering/how-privacy-killed-rbac-328edf6e4be7>

In contrast, according to the National Institute of Standards and Technology (NIST), “the ABAC engine can make an access control decision based on the assigned attributes of the requester, the assigned attributes of the object, environment conditions, and a set of policies that are specified in terms of those attributes and conditions. Under this arrangement, policies can be created and managed without direct reference to potentially numerous users and objects, and users and objects can be provisioned without reference to policy.”<sup>3</sup>

Put simply, within ABAC data engineers and architects define users, objects, and rules independently so that the rules — rather than an administrator — make access control decisions at the time a request is made. With ABAC, attributes are objectively assigned to users, and those attributes, rather than explicit roles, are used to define access.

Immuta implements ABAC with no data copying, no coding, and simple, plain English rules. In doing so, Immuta streamlines IC data team’s responsibilities to provide, discover, and request — ABAC allows data owners to specify who can find and access certain data, and grants or denies requests based on unique sets of attributes. This means any data or analysis is made available to authorized IC personnel based on security clearance and need to know, and removes friction between data owners, compliance professionals, and data analysts by creating regulation-compliant digital data exchanges.

## Audit Capabilities

The responsibilities to provide, discover, and request data are paramount for IC data teams, but comprehensive audits are necessary to measure whether they are being carried out securely.

Immuta allows authorized data consumers to be notified about the existence of information without revealing its contents or any analysis performed on it, which supports both the discovery and request policies. When IC personnel discover information that has the potential to contribute to analysis, optimize data collection, inform data collection strategies and priorities, or to otherwise advance an intelligence mission, they have a corresponding responsibility to request relevant information.

Immuta’s data monitoring and auditing capabilities help eliminate any questions about who is able to find specific data, who accessed that data, and who has requested access to it. This provides complete visibility and flexibility into how policies are enforced and monitored.

Additionally, Immuta allows complex joins and data transformations to be exposed as authoritative virtual data sources for downstream users. Updates or changes to data like these are monitored and audited consistently across data silos from within Immuta’s centralized, virtual layer. This increases the richness and simplifies analysis of audit logs.

3 NIST Guide to Attribute Based Access Control (ABAC) Definition and Consideration  
<https://www.nist.gov/publications/guide-attribute-based-access-control-abac-definition-and-considerations-1>

# Benefits for IC Data Teams and Users

## Centralized Control

The Immuta data control plane enables you to unify disparate data silos virtually, establishing a consistent point for data access for all data analysis. All data and metadata remain in a single location, but are virtually unified to achieve uniform policy orchestration and enforcement. The control plane dynamically protects your data with complex subscription and anonymization policies that are enforced based on the user accessing the data, the attributes of the data itself, environmental variables, and policy logic built from applicable regulations.

## Dynamic Data Policies

Data can be hidden, masked, redacted, and anonymized based on the attributes of the users accessing the data and the purpose under which they are acting. Additionally, the attributes of the data itself and the environment in which the data is stored and controlled can guide data policies. This can be done with fine-grained access controls, down to the cell-level within a database. Policies can be written globally across data sources or locally within them.

## Scalability

Immuta's standard deployment requires minimal administrative effort to scale beyond the addition of nodes to its system. The Immuta web service is stateless and horizontally scalable. As more organizations move to a cloud-based data ecosystem, this is critical for achieving consistent, secure data access and use. Running replicated instances of this internal database allows the active data catalog to scale in support of the web service.

Additionally, the Immuta SQL Query Engine can scale horizontally with user load. In scenarios where queries cannot be fully pushed down to business databases, individual queries are limited by the access controls allocated to an individual instance. This reduces the burden on IC personnel to manually grant access to individual users, which is time consuming and delays speed to data access.

## High Availability

Because each of Immuta's components is designed to be horizontally scalable, Immuta can be configured for high availability. Upgrades and major configuration changes may require scheduled downtime, but even if Immuta's master internal database fails, recovery happens within seconds. With the addition of an external load balancer, Immuta's standard deployment comes preconfigured with these availability features.

# Conclusion

Immuta empowers compliance with ICD 501 for IC and Defense agencies, as well as with the Department of Defense Manual (DoDM) 5240.01 which concerns the discovery, governance, security, and auditing of information collected or analysis produced.

With Immuta, the U.S. Government can leverage an automated data access governance platform that protects our national security intelligence information, including the protection of U.S. Persons' constitutional legal rights, privacy, and civil liberties. Immuta helps to ensure proper implementation of policies regarding the dissemination of U.S. persons' information and other legal protections in the development and use of the information sharing environment.

