ALGORITHMIA ENTERPRISE

# SECURITY & COMPLIANCE OVERVIEW

ALGORITHMIA

Algorithmia Enterprise is the foundation layer for intelligent software. It turns complex services and machine learning models into REST APIs, centralizes them for ease of discoverability, and monitors them from a single dashboard. Companies use Algorithmia Enterprise to reduce duplication of effort between siloed teams and accelerate go-to-market for AI-driven products.

Algorithmia Enterprise is designed from the ground-up to meet the most stringent regulatory and corporate compliance requirements. Our customers' needs span a wide range of complexity, from online retail stores with internet-facing endpoints, to government and finance entities with locked-down classified environments.

This document gives a brief overview on the following topics:

- ➔ Cluster Configuration
- ➔ Three-Layer Permissions
- ➔ Organizations & QA Workflows
- ➔ Audit Trails

## Cluster Configuration

Many use cases require code to run inside a private network, behind complex firewall rules, and for data to never leave specific boundaries. Algorithmia Enterprise is designed with a flexible infrastructure that is compatible with government-grade security requirements.


**Public Cloud vs. On-premises**

Algorithmia Enterprise is cloud-agnostic and can be deployed on any of the major public clouds (AWS, Google, Azure) as well as on-premises via OpenStack. Al deployments will be isolated within their own VPC or combined with an existing VPC.


**Compute and Data Sovereignty**

Algorithmia Enterprise can be configured to work as a multi-region cluster to comply with data sovereignty requirements within regulated industries. A typical setup allows two compute pools, one for each region, and a URL prefix for each region, such as:

https://<region>.api.acme-codex.internal/api/user/algo


**Session Isolation**

A job (or API calls) on Algorithmia Enterprise will always operate in its own memory space and will never share or leak memory to other jobs. Each API call instantiates a dedicated Docker container which is destroyed after execution, for perfect isolation at a peta-scale performance.

## Three-Layer Permissions

Users want to understand and avoid unnecessary risk to their data when building an integration with another service, even if it was developed by their colleagues. To that end, Algorithmia Enterprise implements permissions on three layers: API keys, algorithms, and data sources.

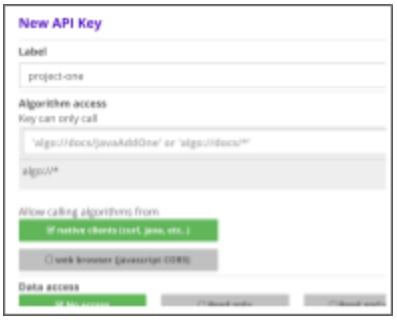


**API Key Permissions**

Users are expected to create multiple API keys under their profile, one for each project or experiment. Each API key is individually auditable, revocable, limited to access specific algorithms and given explicit read/write permissions over data sources.


**Algorithm Permissions**

Authors must specify whether their algorithm requires (1) access to network, and (2) access to call other algorithms. Algorithms created without those permissions will be executed in sandboxed Docker containers that do not have access to those resources. Permissions are clearly displayed on the algorithm page and made available to the consumer.


**Data Permissions**

Similar to algorithms, a data source (or a data collection can be configured to allow or disallow Read/Write access from other users. This is especially valuable when using Algorithmia Enterprise within a multi-business line environment where internal data regulations enforce a specific compliance scheme.

## Organizations & QA Workflows

Large organizations are broken down to teams with different mandates and access rights. Algorithmia Enterprise is designed to adhere to that design while at the same time bring an unprecedented level of collaboration between siloed groups.

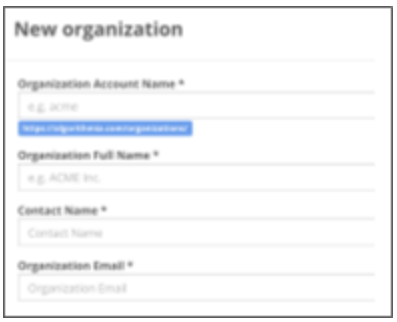


**Teams & Organizations**

Users can create or join teams, which gives them the ability to create algorithms under that team's name instead of their own. This allows algorithms to belong to and be maintained by a group of people (department o business line) instead of an individual (i.e. an engineer).


**Private vs. Public Algorithms**

Algorithms can be marked as private or public. Private algorithms are only accessible to their owners, which could be an individual or a team. Public algorithms are accessible to all users on the Algorithmia Enterprise platform.


**QA Publishing Workflow**

Organizations can choose to enforce a publishing workflow whereby algorithms can only be made accessible to outsiders after it is approved by a compliance officer. This ensures a minimum-level of quality assurance in the public pool of algorithms, while still allowing individuals to experiment privately.

## Audit Trails

Monitoring, tracing, and preventing suspicious activity is critical to all companies, especially in industries with high level of compliance and regulation. Algorithmia Enterprise can be configured to capture and trace any activity.


**API Logging and Verbosity**

Algorithmia Enterprise logs every API call, showing what user used what API key to call what algorithm and version. Admins can configure the logging module to capture more or less (such as full or partial inputs and outputs) depending on the use case.


**Usage Attribution**

Algorithmia Enterprise measures how much each user in the platform contributes to the total cluster utilization, allowing administrators to understand how different teams and specific individuals are contributing to total computing costs.

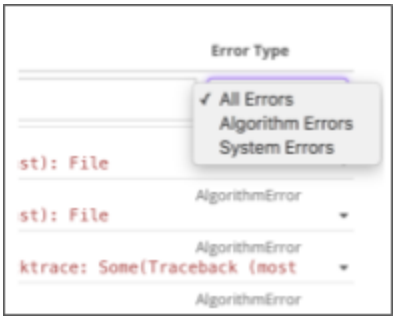


**Error Logs**

Every exception anywhere in the platform is captured, including errors from user-generated algorithms. Exceptions can be configured to capture full or partial input, error message, stack trace, and be connected to a notification trigger to take corrective action.