HCL AION

HCLSoftware

Security Trust Center Document

Table of Contents

List of Tables

Document Revision History

This document is updated with each release of the product or when necessary.

This table provides the revision history of this AION Security Trust Center document .

Version Date	Description
September, 2025	HCL_AION_ v2.5.1_Security_Trust_Center_Document
June, 2026	HCL_AION_v2.6.0_Security_Trust_Center_Document

1 Our Commitment to Security and Trust

At HCLSoftware, we are fundamentally committed to earning and maintaining your trust. Security, privacy, and compliance are the foundations of our products. The HCLSoftware security strategy covers all aspects of our business, including corporate and organizational security policies, incident management and response, business continuity and disaster recovery, secure software development processes, and privacy.

For HCL AION, security is a foundational design principle embedded into every aspect of the platform rather than an afterthought. Our objective is to enable organizations to unlock the value of artificial intelligence and machine learning while maintaining the highest levels of security, privacy, and compliance across the entire AI lifecycle.

2 Document Summary

HCL AION is a unified, enterprise AI platform that manages the complete lifecycle of both Classical Machine Learning (ML) and Generative AI (GenAI) workloads — from data ingestion and model development through deployment, inference, and ongoing monitoring. Its classical machine-learning capabilities span automated data pipelines, exploratory data analysis, feature engineering, AutoML and custom model training, prediction and serving, drift monitoring, and explainability. Its Generative AI capabilities include the NLP interface, LLM fine-tuning, a centralized Model Registry, a multi-layer guardrail framework, and enterprise governance with approval workflows.

This document provides an overview of the security architecture, controls, and practices of HCL AION. Specifically, it describes:

- How HCL AION is architected, deployed, and operated within the customer's own infrastructure, and the security model that applies to that customer-hosted deployment.
- The data privacy, identity, access, governance, auditability, and responsible-AI controls that protect customer data, models, and intellectual property across both the Classical ML platform and the GenAI subsystem.

3 Introduction

HCL AION is a comprehensive AI lifecycle management platform. It spans raw data ingestion through the deployment of scalable ML models, as well as the end-to-end Generative AI lifecycle, including LLM fine-tuning, deployment, inference, and monitoring.

The platform is delivered as a containerized, cloud-native application and is operated entirely within the customer's environment, whether an on-premises data center or the customer's own public-cloud tenancy.

3.1 Data Sources

3.1.1 User

Authenticated users interact with HCL AION through the web interface, the NLP Interface conversational interface, and the REST API. User identity, persona assignment, and session data are governed by the platform's identity and access controls (see Identity and Access Management).

3.1.2 Enterprise Data Sources

HCL AION ingests data that the customer provides from local files, external databases, URLs, Kafka, Apache Ni-fi and unstructured documents (including PDF, CSV, DOCX, JSON, and text). Ingested content is processed into compatible format for AION use cases. All ingested data remains within the customer-controlled deployment.

3.1.3 Knowledge and Training Data

Customer-supplied datasets used for model training, fine-tuning, and retrieval are bound to a specific use case, which acts as the container for the data, the resulting model versions, the applicable guardrail policies, and the deployment state. Access to these assets is governed by role-based access control and ownership rules.

3.2 HCL AION Deployment Model

HCL AION is deployed and operated within the customer's own infrastructure. HCL does not host, manage, or access the customer's operational environment or data. As a result, the customer retains full ownership of, and control over, all data, models, and configuration, and can align the deployment with internal policies and geographic data-residency requirements. The following table summarizes the division of responsibilities under this customer-hosted model.

Table 1 - HCL AION Deployment Responsibilities (Customer-Hosted Model)

Responsibility	Customer	HCLSoftware
Physical and infrastructure security (hardware, data center, host OS)	Owns	—
Network security controls (firewalls, segmentation, IDS/IPS, VPN)	Owns	—
OS hardening and patching of underlying infrastructure	Owns	—

Identity, authentication, and authorization configuration	Owns / operates	Provides product capability (Keycloak, RBAC)
Backup, disaster recovery (DR), and business continuity (BCP)	Owns / operates	Supports via modular design and cross-site replication
Secure application software, patches, and security advisories	Applies updates	Delivers and maintains
Application-layer vulnerability remediation in HCL AION	Applies updates	Develops and releases fixes (PSIRT)

3.3 Platform Security Highlights

HCL AION applies a defense-in-depth approach in which multiple, overlapping controls protect the platform, its data, and its models. Highlights include:

- Encryption of data in transit using TLS 1.2/1.3 for external communications, and encryption of data at rest.
- Standards-based identity via Keycloak with OAuth2/OIDC authentication, JWT-secured API requests, and enterprise SSO (including SAML 2.0) through identity brokering.
- Persona-based role-based access control (RBAC) enforcing least privilege, with resource ownership rules and read-only access for unassigned users.
- A three-layer guardrail framework governing GenAI input, output, and data interactions, with administrator-configured, versioned policies.
- Enterprise governance with approval workflows that gate sensitive actions (access, deployment, model and data upload, infrastructure configuration, and fine-tuning).
- Comprehensive audit logging and structured, centralized observability for forensic traceability and operational visibility.
- Internal microservice communication restricted to the internal network.

3.4 Core Security Principles

HCL AION's deployment adheres to three industry-standard principles:

3.4.1 Defence-in-Depth

This is a layered approach to security. This principle assumes that no single security control is infallible. Therefore, the platform is protected by multiple, redundant, and overlapping security measures across its entire technology stack.

- Multiple layers of protection
- Redundant mechanisms to prevent, detect, and respond to threats.

3.4.2 Secure-by-Design

The secure-by-design principle dictates that security considerations are integrated into the entire product lifecycle, from the earliest stages of conception and design through development, testing, deployment, and maintenance.

- Threat modelling during design phases
- Static and dynamic application security testing (SAST/DAST) during development

- Independent and rigorous penetration testing before release.

3.4.3 Zero Trust

No user, device, or application is trusted by default, whether inside or outside the network. All access requests require strong authentication, fine-grained authorization with RBAC, and encryption in transit using TLS 1.2 or higher for all external communications.

3.5 Responsible AI Controls

HCLSoftware promotes safe, transparent, and ethical AI usage. HCL AION provides the following responsible-AI controls:

Table 2 – Responsible AI Controls

Control	Description
Content Filtering & Prompt Validation	Multi-layer guardrails validate inputs and outputs and enforce data-handling policies, violations can block, warn, redact, or allow with logging.
Risk & Incident Management	AI-specific risks are tracked in a risk register, vulnerabilities are managed through the HCLSoftware PSIRT process.
Human Oversight	Workflows require a human decision (with recorded rationale) before sensitive actions such as deployment, fine-tuning, and data upload proceed.
PII Protection	Data-category guardrails support detection and redaction of sensitive data, data minimization is applied by design.
Accuracy & Transparency	Model evaluation metrics, explainability, model cards, and an Inference Playground support transparent assessment of model behavior.
Traceability, Auditability & Observability	All actions on governed assets are recorded in an audit log, structured logs and metrics provide end-to-end observability.
Guardrails	Administrator-configured, versioned guardrail policies are bound to use cases and take effect immediately for new requests.

4 Technical Architecture

4.1 Overview

HCL AION comprises two complementary subsystems that share a common identity, governance, and deployment foundation: the Classical ML platform, which manages the end-to-end MLOps lifecycle, and the GenAI subsystem, which manages the end-to-end Generative AI lifecycle. Both run as containerized microservices on Kubernetes within the customer's environment.

HCL AION is offered as both a Classical ML and a Generative AI AIOps platform. The two are delivered together on a shared identity, governance, and deployment foundation, and customers may adopt either capability set or both.

4.2 Architecture Foundations

The Classical ML platform is a containerized, cloud-native platform deployed as Kubernetes microservice pods, organized into functional engines spanning data ingestion and pipelines, data curation, the AI/ML engine core, and model operations and analytics.

The GenAI subsystem is organized into two integrated layers: a conversational and workflow layer, and a platform infrastructure layer. Its orchestration engine manages all AI workflows as a state machine, in which the workflow state object maintains the conversation history, the planned action queue, and the active data-source configuration.

4.3 Classical ML Platform and Lifestyle

The Classical ML platform supports the complete machine-learning lifecycle through a set of functional processing engines, available through low-code and no-code interfaces as well as the REST API:

- Data ingestion and pipelines: connectors to databases, object and file storage, distributed file systems, and streaming sources, with pipeline orchestration for repeatable data flows.
- Data curation: exploratory data analysis, data-quality checks, transformation, and feature engineering.
- Model development: AutoML and custom model training, including distributed training for large datasets, hyperparameter tuning, and a versioned model registry.
- Prediction and serving: real-time (instance), batch, and streaming inference, with deployable artifacts such as containers and generated code.
- Monitoring and explainability: input and performance drift detection that can trigger retraining, and model explainability for transparent, accountable predictions.
- Bring-Your-Own-Model (BYOM): import and govern externally trained models alongside platform-trained models.

All processing runs as isolated microservices within the customer's environment and is governed by the same identity, access-control, logging, and model-governance controls applied across the platform.

4.4 Containerized Deployment

HCL AION is built on containerization, using Docker images orchestrated by Kubernetes. Every component, including model-serving endpoints, is deployed as a container, which provides two core security benefits:

- Isolation: process- and filesystem-level isolation between components, so that a vulnerability in one model or service does not affect others running on the same host.
- Consistency: each image packages the application with all of its dependencies, producing a reproducible environment that is identical from development through to production.

In the customer-hosted model, HCL does not host, manage, or access the operational environment. HCL delivers secure, signed application and model-serving images together with recommended configuration baselines, while the customer operates the cluster, network controls, and identity integration. The platform's modular design supports cross-site replication to aid the customer's backup, disaster-recovery, and business-continuity plans.

4.5 Solution Capabilities

- Classical ML: automated data ingestion and pipelines, exploratory data analysis, feature engineering, AutoML model training, prediction/serving, drift monitoring, explainability, and model packaging/export.
- Generative AI: GenAI use-case lifecycle management, LLM fine-tuning, an Inference Playground, an LLM Recommender, a centralized Model Registry, and a multi-layer guardrail framework.
- Bring-Your-Own-Model (BYOM): import of external models (for example, ONNX and HDF5) or LLM into the platform registry.

4.6 Consumption Channels

- Web interface — a browser-based application whose navigation and options adapt to the signed-in user's persona.
- NLP Interface — a conversational interface that drives the complete lifecycle in natural language, with streaming responses, persistent conversation history, and persona-aware behavior.
- REST API and WebSocket — a programmatic interface for use-case management and event-driven interactions. Use-case state remains consistent across the chat and web interfaces.

4.7 Platform Features

The HCL AION Engine provides orchestration capabilities — Interpreter, Router, and Curator — that unify Classical ML and Generative AI workflows. Each GenAI use case acts as the container for its data, fine-tuned model versions, guardrail policies, and deployment state. Workflow orchestration supports asynchronous execution with notifications, and LLM routing is configurable per workflow node.

4.8 AI and Foundation LLM Integration

HCL AION supports LoRA, QLoRA, and full fine-tuning strategies. Any model can be referenced as a base model. For inference and composition, the platform is configured to use the customer's own foundation-

model provider account (Azure OpenAI or OpenAI), models are registered in the Model Registry with their provider, and deployment is driven by customer before production use. Because the provider account and its data-processing terms are owned and controlled by the customer, foundation-model usage remains within the customer's governance boundary.

4.9 Enterprise Integrations

HCL AION ships with a catalog of enterprise integrations. Each integration is configured by the customer and secured with customer-managed credentials.

Table 3 – Enterprise Integrations and Security Notes

Integration	Category	Security notes
Keycloak	Identity & security	OAuth2/OIDC authentication, SSO, RBAC, identity brokering to SAML 2.0.
Azure OpenAI / OpenAI	LLM provider	Customer's own provider account/API key, used for fine-tuning and inference. Deployment approval required.
Hugging Face Hub	LLM provider	Source for base models referenced by identifier, subject to model-provenance governance.
Custom extensions	Extensibility	Microservices implementing the AION contract, those handling data, or model artifacts and all calls are audited.

4.10 Hosting and Deployment

- Containerized and cloud-native: deployed as Kubernetes microservice pods, portable across AWS, Azure, GCP, and on-premises using the same container set.
- Prerequisites: Kubernetes 1.26+, Helm 3.12+, an ingress controller, and ReadWriteMany (RWX) shared storage (for example, Amazon EFS, Azure Files, or NFS). ReadWriteOnce volumes such as EBS or Azure Disk are not supported.
- Container images are pulled from HCLSoftware Harbor (or an equivalent registry) using customer-configured credentials.
- Horizontal pod autoscaling provides per-service elasticity.
- Schema migrations are managed with Alembic, post-deployment configuration imports the Keycloak realm and preloads guardrail evaluation models.

4.11 LLM Routing and Orchestration

GenAI workflows are coordinated by a state-machine orchestrator. The Interpreter parses the user's intent, the Router selects and routes to the appropriate model or workflow node (with configurable LLM routing per node), and the Curator manages the data and model context.

4.12 Model Fine-Tuning Approach

Fine-tuning adapts a base foundation model to a customer's domain using the customer's own training data. Jobs are submitted through the Web/NLP Interface or the API and tracked in the fine-tuning service. Each run creates a new model version record with a version name, training status, and evaluation metrics, enabling comparison and selection. Fine-tuned models must be registered in the Model Registry and certified by an administrator before deployment to production.

4.13 Conversation and Workflow State

The NLP Interface maintains persistent conversation history and persona-aware behavior. The orchestration engine persists the workflow state object — conversation history, the planned action queue, and the active data-source configuration — so that long-running, asynchronous workflows can execute reliably and emit notifications. Use cases, model versions, approval requests, notifications, and messages are persisted in the platform.

4.14 Data Storage

Within the customer's environment, HCL AION uses PostgreSQL and the RWX shared storage as the primary store for platform metadata, use cases, the model registry, approvals, notifications, and audit records, Redis for caching and session/state data, and ReadWriteMany shared storage for large artifacts referenced by path. All data at rest is protected, and all data stores reside within the customer-controlled deployment.

5 AI and ML Security Controls

Beyond conventional application security, HCL AION applies controls specific to the AI and ML lifecycle, across both the Classical ML platform and the Generative AI subsystem

5.1 AI Governance and Risk

HCL AION follows an AI Management System (AIMS) approach aligned with HCLSoftware AI policy, with defined accountable roles and human oversight for safety-critical use cases. An AI risk register tracks AI-specific risks — including data poisoning, model theft, inference leakage, and over-reliance — and feeds the broader risk-management program described under Information Security Governance and Risk Management.

5.2 Data and Training Integrity

Training and fine-tuning data is governed by enforced dataset lineage and consent or licensing checks before use. Poisoning defenses — schema and semantic validation, canary datasets, and outlier and label-flip detection — are applied as a pre-training gate, so that only validated data is admitted into model development.

5.3 Model Integrity and Robustness

Models are managed in an access-controlled registry, and model artifacts and serving images are signed.

5.4 Operations and Runtime

Service-to-service communication between data pipelines and model-serving components uses mutual TLS (mTLS) with fine-grained authorization. Security telemetry correlates model drift with security anomalies — for example, atypical input distributions — so that data-quality and security signals are assessed together.

6 Data Privacy

Data privacy (also called information privacy) refers to the right of individuals to control how their personal information is collected, used, shared, and stored. It is about ensuring that personal or sensitive data is handled in ways that protect people's rights and maintain their trust.

6.1 Data Roles and Responsibilities

Because HCL AION is deployed and operated entirely within the customer's own infrastructure, the customer is the data controller and retains full ownership and control over the data collected and the purpose of its use. HCL does not host, process, or access customer data in this model. HCL AION itself does not transmit customer data or models outside the customer's specified environment unless the customer explicitly configures an external integration (for example, a foundation-model provider account). For more information, visit HCL Software Privacy.

6.2 Purpose, Limitation and Data Minimization

The platform is designed to process data solely for the purpose of building, operating, and governing the customer's AI and ML use cases. We practice data minimization, and the platform is not designed to process special categories of sensitive personal data. Customers determine what data is ingested and are responsible for ensuring an appropriate lawful basis for its use.

6.3 PII Data

Sensitive and personally identifiable information (PII) is protected through data-category guardrails and data-handling controls. Guardrail policies can detect and act on sensitive content in inputs and outputs, including redaction.

6.3.1 PII Minimization and Consent Controls

Data-category guardrails (Layer 1 rule-based checks, Layer 2 library-backed evaluation, and Layer 3 LLM-as-Judge) can be configured to block, warn, or redact sensitive data. Customers, as data controllers, manage consent and any data-subject rights (such as erasure) within their environment, supported by the platform's access controls and audit trail.

6.4 Chat Data

6.4.1 User

Prompts and instructions that users submit through the NLP Interface are processed to drive workflows and are retained as conversation history to preserve context. Access to conversation data is restricted by persona and ownership.

6.5 Data Access

6.5.1 Access Management

Administrators manage users and persona assignments through the Access Management panel (Settings → Access Management), which updates the corresponding Keycloak realm roles. Administrators can create

users, assign personas, and deactivate accounts (blocking login until reactivation). Periodic access reviews are recommended to prevent privilege creep. Non-administrator users can view and modify only the use cases and resources they own.

6.5.2 SSO using OAuth2/OIDC and SAML 2.0

Keycloak provides OAuth2/OIDC identity management and supports identity brokering with enterprise identity providers, including SAML 2.0. Enterprise directory groups can be mapped to Keycloak realm roles to automate persona assignment.

7 Data and Asset Protection

HCL AION applies a multi-layered strategy to protect the confidentiality and integrity of customer data, platform configuration, and machine-learning models — treated as first-class assets — in transit, at rest, and in use within the customer's environment.

7.1 Credential and Secret Security

Service credentials and secrets are protected using strong hashing with unique, per-secret salts, so that identical inputs do not produce identical hashes and credentials cannot be correlated. Sensitive credentials such as database passwords, API keys, and any other credentials are never committed to source control.

7.2 Model Protection

Trained model files are encrypted, and additional safeguards protect model intellectual property: model obfuscation increases the difficulty of reverse-engineering, and digital watermarking embeds identifiers that support IP protection and provenance. Model and serving artifacts are signed, and access to the model registry is governed by role-based access control.

8 Information Life Cycle and Data Management

Because HCL AION runs within the customer's environment, the customer controls the retention, return, destruction, and backup of all data. HCLSoftware provides recommended practices and platform capabilities that support these activities.

8.1 Data Retention

Application, model, conversation, and audit data are retained for the period the customer configures. For audit and operational logs, recommended retention is 12 months online, with up to 7 years in cold storage for regulated industries, retention can be enforced through a scheduled cleanup job or database partitioning.

8.2 Data Return and Destruction

As the data controller operating the platform, the customer can export, return, or destroy data directly within its environment at any time, in line with its own policies and contractual or regulatory obligations.

8.3 Data Backup and Resiliency

HCL AION's modular, containerized design supports backup, restore, database migration, and cross-site replication. Backup frequency and retention are operated by the customer, the following values are recommended defaults.

Table 4 – Data Backup and Resiliency (Recommended, Customer-Operated)

Component	Recommended frequency	Recommended retention
PostgreSQL (platform metadata, registry, audit)	Daily	30 days
Vector store (embeddings)	Daily	30 days
Redis (cache / session state)	Daily	15 days
Object / shared storage (artifacts)	Daily (Twice)	60 days

9 Information Security Governance and Risk Management

9.1 Security Frameworks

HCLSoftware aligns its security program with recognized international standards, including ISO/IEC 27001 and 27002 for information security management, ISO/IEC 27034 for application security, and ISO 31000 and ISO/IEC 27005 for risk management. Secure development aligns with the NIST Secure Software Development Framework (SSDF).

9.2 Network Security

TLS is terminated at the ingress for all external endpoints, and TLS is required for ingress in production deployments. Internal microservice communication is restricted to the internal network, and microservice ports are not exposed externally. Cross-origin resource sharing is restricted to explicitly configured origins (wildcards are not used in production). Network policies isolate the data tier.

9.3 Infrastructure Security (Customer-Managed)

Physical, host, and network infrastructure security is owned by the customer, who controls the data center or cloud tenancy in which HCL AION runs. HCLSoftware provides recommended configuration baselines and hardening guidance, including the use of Kubernetes Secrets or a vault-backed solution for credentials and the principle of never committing secrets to version control.

9.4 Application Security

- Authentication via Keycloak (OAuth2/OIDC) with enterprise SSO support (via SAML).
- JWT-secured API requests (HS256) with a 30-minute access-token expiry and a 60-minute refresh-token validity, tokens are validated on every protected endpoint.
- Passwords are hashed before storage, no plaintext passwords are stored, and the JWT signing secret must be a strong random value.
- Persona-based RBAC enforces least privilege, unassigned users receive read-only access, administrators can intervene on any resource.
- Encryption in transit (TLS 1.2/1.3) and at rest.
- API security includes scoped permissions, approval gating of sensitive actions, and rate limiting.

9.5 Identity and Access Management

Identity is centralized through Keycloak, which provides federated authentication and single sign-on (SSO) with enterprise identity providers. Access is governed by role-based access control (RBAC) that enforces least privilege:

- Classical ML platform: roles such as administrator, data scientist, ML engineer, and operator gate use-case, training, and deployment operations.

- Generative AI subsystem: the built-in personas — domain expert, data scientist, data engineer, and administrator — control route visibility, API permissions, and operation scope, with read-only access for unassigned users (see Generative AI Data Security and Responsible AI Controls).

Privileged access is recertified through periodic access reviews to prevent privilege creep, and authentication and authorization events are recorded in an immutable audit log for compliance and forensic investigation.

9.6 Logging and Monitoring

HCL AION maintains multiple log types so that platform activity remains traceable and observable:

- Container logs: generated by the underlying containers and accessed through the hosting Docker or Kubernetes environment.
- Application logs: emitted by the application and retained for the period the customer configures, accessible only to customer-specified personnel.
- Model logs: maintained per model and reviewed by the user who trained the model, available through the application interface.
- Access logs: track access to the application and are administered by the customer's administrator.

Models and the platform are monitored through built-in functionality and the container cluster's native mechanisms. In the Generative AI subsystem, services additionally emit structured JSON logs that can be forwarded to a central log-management platform with alerting on error and critical events (see Audit and Logging).

9.7 Human Resources Security

Human resources security practices, background checks, and training processes are managed by the HCLSoftware team.

9.8 Risk Management

HCLSoftware has a formalized risk management program that aligns with ISO 31000 and ISO/IEC 27005 best practices, as well as ISO/IEC 27001 and 27002. Risk management processes are integrated with other management systems, such as the Information Security Management System (ISMS). Security controls are implemented in accordance with our ISMS to manage risk across the organization.

9.9 Responsibility for Risk Management

Risk management is owned at senior levels of the organization and operates on a regular cadence of reviews. Every staff member shares responsibility for identifying and managing risk within their area, supported by defined policies, standards, and the ISMS.

10 Secure Development Lifecycle Management (SDLC)

10.1 Secure Development

HCL AION is developed under a secure SDLC aligned with the NIST Secure Software Development Framework (SSDF), with security gates at each stage.

10.2 Requirements & Planning

A data privacy assessment is performed for new features, applying privacy-by-design principles.

10.3 Design

Formal threat modeling and a secure design review are conducted by the HCLSoftware Security team against industry standards.

10.4 Development

Developers follow OWASP secure-coding practices with regular training. Automated static (SAST) and dynamic (DAST) analysis, software composition analysis, and container image scanning are applied to custom code and open-source dependencies.

10.5 Penetration Testing

Internal penetration testing is performed on every major release, and a comprehensive penetration test is performed annually by an independent, third-party security firm.

10.6 Maintenance

The HCLSoftware Product Security Incident Response Team (PSIRT) manages ongoing vulnerability monitoring and response and publishes Security Bulletins referencing relevant CVEs with remediation guidance.

11 Generative AI Data Security and Responsible AI Controls

HCL AION applies a layered set of controls specifically to the Generative AI subsystem.

11.1 User Data Protection

Customer data ingested for GenAI use cases (files, database extracts, URLs, and unstructured documents) is processed and stored within the customer-controlled deployment. Data-category guardrails support detection and redaction of sensitive content, and access to use-case data is restricted by persona and ownership.

11.2 Access Control & Identity

Authentication is enforced through Keycloak OAuth2/OIDC. JWT tokens (HS256) secure all API requests, with a 30-minute access-token expiry and a 60-minute refresh-token validity and are validated on every protected endpoint. Passwords are hashed. Authorization uses persona-based access control across the domain expert, data scientist, data engineer, and administrator personas, users without an assigned persona receive read-only access.

Table 5 - Built-in Personas and Access Scope

Persona	Representative access scope
Domain expert	Create use cases, upload data, chat with the NLP Interface, and run inference, owns the use cases they create.
Data scientist	All domain-expert capabilities, plus fine-tuning configuration, model evaluation, and Playground testing.
Data engineer	Data upload, ingestion pipelines, and data-source configuration.
Administrator	Full platform access, manages users, roles, system configuration, and approval workflows, can act on any resource.
Unassigned	No access until a persona is assigned.

11.3 API & Model Access Security

API requests are authenticated and authorized on every protected endpoint, with scoped permissions and rate limiting. Sensitive actions are gated by approval workflows, and fine-tuned models must be registered and certified in the Model Registry before deployment to production.

11.4 Content Filtering & Prompt Validation (Multi-Layer Guardrails)

HCL AION provides a three-layer guardrail framework for safe, controlled, and policy-compliant GenAI usage. Guardrail policies are created and modified by administrators in the Guardrails section, are versioned, require approval, and take effect immediately for new requests. Each policy specifies a category (input, output, or data), an evaluation layer, and an action on violation.

Table 6 – Multi-Layer Guardrail Framework

Layer	Mechanism	Categories	Actions on violation
Layer 1	Rule-based checks (regex patterns or keyword lists)	Input / Output / Data	Block, Warn, Redact, or Allow
Layer 2	Library-backed evaluation using preloaded models	Input / Output / Data	Block, Warn, Redact, or Allow
Layer 3	LLM-as-Judge using a judge prompt template	Input / Output / Data	Block, Warn, Redact, or Allow

11.5 Governance and Approvals

Sensitive operations require administrator approval before they proceed, including access role changes, model deployment, data upload for sensitive use cases, infrastructure configuration changes, and fine-tuning job submissions. Each approval record captures the requester, the requested action, the justification, timestamps, and the approver's decision and rationale, and is retained in the audit log. Models must be certified before deployment, and pipelines are tracked in a data-pipeline registry with their own approval requirements.

11.6 Audit and Logging

Audit events for all actions on governed assets are recorded in the platform's PostgreSQL database and are accessible to administrators through the UI and API. Services emit structured JSON logs that can be collected and forwarded to a central platform with alerting on error and critical events.

Table 7 – Audit Event Categories

Category	Events captured
User access	Login, logout, registration, and role assignment/removal.
Use-case lifecycle	Creation, update, deletion, and status transitions.
Workflow execution	Start, completion, failure, and retry.
Model registry	Registration, certification, activation, deactivation, and deletion.
Data pipeline registry	Creation, configuration changes, execution, and deletion.
Approvals	Request creation, approval, rejection, and comments.

11.7 Secure Development Practices

Generative AI components are developed under the same secure SDLC described in Section 8, with attention to AI-specific threats. Practices include alignment with the OWASP Top 10 for LLM Applications and STRIDE-based threat modeling, supported by SAST/DAST tooling.

11.8 Responsible AI and Ethical Safeguards

HCL AION supports explainability, model evaluation metrics, model cards, and human oversight through approval workflows. Bias and fairness considerations are addressed through monitoring and evaluation, and outputs can be assessed in the Inference Playground before deployment.

12 Model Governance and Secure AIOps

HCL AION provides comprehensive model governance and a secure AIOps lifecycle across both Classical ML and Generative AI models.

12.1 Model Lifecycle Management

A centralized Model Registry provides versioning, lineage, comparison, approval workflows, and deployment tracking for both ML and LLM models. Each model version records training status and evaluation metrics, and lineage captures data sources, transformations, and training parameters.

12.2 Data Governance

Data quality is monitored for missing, corrupt, or skewed data, data lineage is tracked from raw sources to model inputs, and privacy compliance is supported through anonymization and consent logging.

12.3 Model Explainability and Interpretability

The platform integrates model-specific explainers (for example, SHAP and anchor-based explanations) and model cards so that model behavior and predictions can be understood and communicated.

12.4 Fairness and Bias Monitoring

Pre- and post-deployment bias checks are performed across sensitive attributes, using fairness metrics such as equal-opportunity difference and demographic parity.

12.5 Model Performance Monitoring

Concept drift and data drift are detected, and performance degradation is tracked over time (for example, AUC and F1-score), triggering retraining when appropriate.

12.6 Auditability and Traceability

Comprehensive logging records who trained which model, when, and with which data and parameters, providing audit trails for external compliance and internal accountability.

12.7 Retraining and Lifecycle Refresh

Models are refreshed when performance decay or data drift is detected. Retraining requires an appropriately privileged role, and the retraining logic, triggers, and outcomes are tracked in the audit trail to preserve lineage and accountability.

12.8 Secure AIOps Lifecycle

– Data sourcing and preparation: RBAC-controlled access to data sources, automated provenance/lineage tracking, and poisoning prevention via schema validation and anomaly detection.

– Model training: encrypted communication between distributed training nodes and adversarial-robustness testing.

– Model deployment: containerized deployments built on minimal base images, with deployment policies enforcing runtime security constraints.

12.9 Supply Chain Integrity

– Software Bills of Materials (SBOMs) are generated for libraries and dependencies.

– Application and model artifacts are signed, with signed provenance and admission enforced for images and models.

– Training is reproducible, with detailed step-by-step logs.

– Runtime verification includes image scanning in the registry and the cluster.

13 Threat Management and Incident Response

HCL AION is designed to protect against AI-specific attacks, including model theft, inference attacks, and supply-chain manipulation. The HCLSoftware Product Security Incident Response Team (PSIRT) provides:

- Receipt and triage: a single point of contact for vulnerability reports from customers, researchers, and internal teams.
- Investigation and coordination: investigation of reports and coordination with the AION development team on root cause and impact.
- Remediation planning: development, validation, and timely release of patches or mitigations.
- Transparent communication: verified vulnerabilities and resolutions are communicated through official HCLSoftware Security Bulletins, referencing relevant CVEs with actionable remediation guidance.

More information is available at the HCLSoftware PSIRT page: https://www.hcl-software.com/resources/psirt

14 Conclusion

HCL AION is architected to meet the security, privacy, and compliance challenges of operating Classical ML and Generative AI at enterprise scale. Its multi-layered strategy is rooted in defense-in-depth, secure-by-design, and Zero Trust principles, and is reinforced by strong identity and access control, multi-layer guardrails, enterprise governance with approvals, and comprehensive auditability — all operating within the customer's own environment. Our valued clients can rest assured that we keep security foremost in our minds as we develop, test, and deliver effective and secure solutions.

15 Support

For any product-related queries, please contact the HCL AION team.

HCLSoftware

hcl-software.com