HCLSoftware

# HCL **AION**

**Security Trust Center Document**

Version 1.0

# Table of Contents

# 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_Security_Trust_Center_Document_v1.0 |

# 1 Our Commitment to Security and Trust

At HCLSoftware, security is not an afterthought — it is a foundational design principle embedded into every aspect of HCL AION. Our mission is to enable organizations to unlock the full potential of AI and machine learning while maintaining the highest levels of security, privacy, and compliance.

We understand that trust is earned through **transparency, rigorous security controls, and consistent performance**, and this document outlines the architecture, principles, and operational practices that form the security backbone of HCL AION.

This document provides a comprehensive overview of the security architecture, features, and practices supporting the HCL AION platform. It describes how the platform ensures secure practices throughout the workflow. It details the continuous vulnerability scanning process, prioritization of remediation activities, and robust change management workflows designed to reduce exposure.

# 2 HCL AION Overview

HCL AION is a comprehensive AI lifecycle management platform designed to tackle real-world machine learning challenges. It spans the entire spectrum from raw data ingestion to the deployment of scalable machine learning models, including the orchestration and management of data pipelines. This enables seamless integration and operation with compatibility across diverse data sources.

HCL AION is operated entirely within the customer's infrastructure — whether in an on-premises Data Center or public cloud environment. This ensures that all operational control, data governance, and compliance measures are under the customer's direct supervision, allowing seamless alignment with internal policies and geographic data residency requirements.

HCL AION itself does not transmit customer data or ML models outside the customer's specified environment unless specifically configured by the customer for any external integration requirements.

# 3 Deployment Security Model

HCL AION provides a secure and flexible framework for deploying trained models into production environments, ensuring that the final, customer-facing asset is protected. In a customer-controlled deployment, **HCL does not host, manage,** or **access your operational environment**.

## 3.1 Containerized Deployment

The platform's architecture is built on a foundation of containerization, using technologies like Docker and orchestrators such as Kubernetes. Every component of HCL AION, including the model serving endpoints, is deployed as a container. This approach offers significant security benefits:

- **Isolation:** Containers provide process and filesystem isolation, ensuring that a vulnerability in one model or application does not affect others running on the same host.
- **Consistency:** Container images package the application with all its dependencies, creating a consistent and reproducible environment that behaves identically from development to production.

## 3.2 Customer Responsibilities

- Full control over **physical security** of hardware assets
- Configuration and enforcement of **network security controls**, including firewalls, network segmentation, intrusion detection/prevention systems (IDS/IPS), and VPN configurations
- Operating system **hardening** and patching for underlying infrastructure.
- Management of **identity and access** to infrastructure and the AION platform
- Execution of **backup, disaster recovery (DR), and business continuity (BCP)** plans aided by HCL AION's modular design that supports cross-site replication.

### 3.2.1 HCL Responsibilities

- Delivery of secure application packages
- Provision of security guidance recommended configuration baselines, and deployment of best practices.
- Timely security updates, patches, and advisory notifications for HCL AION
- Support for integration with enterprise security tooling

This **shared responsibility model** ensures that the customer maintains operational control, while HCL ensures the application software itself is secure, resilient, and compliant-ready.

# 4 Core Security Principles

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

## 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.

## 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.

## 4.3 Zero Trust

In a Zero Trust architecture, no user, device, or application is trusted by default, regardless of whether it is inside or outside the corporate network.

- All access requests require **strong authentication** via Federated Authentication or SSO, **fine-grained authorization** with RBAC, and **encryption in transit with** TLS 1.2+ for all external communications.

# 5 AI/ML Security Controls

## 5.1 AI Governance & Risk

- **AIMS scope and policy:** Follows HCLSW AI policy and scope for the AI Management System (AIMS) with defined accountable roles and human oversight for safety-critical use cases.
- **AI risk register:** Risk register covering data poisoning, model theft, inference leakage, and over-reliance risks.

## 5.2 Data & training

- **Provenance and licensing:** Enforced dataset lineage and consent checks, attesting training sets before use.
- **Poisoning defences:** Schema/semantic validation, canary datasets, outlier, and label-flip detectors, treating as a pre-train gate.

## 5.3 Model integrity & robustness

- **Model registry security:** Stores models in an access-controlled registry, **signed** model artifacts and serving images.
- **Adversarial testing:** Testing against evasion, extraction, inversion, and backdoors.

## 5.4 Operations & runtime

- **Service-to-service trust:** Mutual TLS and fine-grained authorization between data pipelines and model serving.
- **Security telemetry for AI:** Correlating drift with security anomalies (e.g., atypical input distribution).

# 6 Identity and Access Management (IAM)

- **Federated Authentication**: Integration with enterprise Identity Providers
- **Single Sign-On (SSO)**: Integrate with existing customer SSO.
- **Role-Based Access Control (RBAC)**: Predefined roles (Data Scientist, ML Engineer, Administrator) enforce the principle of least privilege.
- **Privileged Access Reviews**: Periodic access recertifications to prevent "privilege creep"
- **Immutable Audit Logging**: Every authentication and authorization event logged for compliance and forensic investigation.

# 7 Logging & Monitoring

The system generates and maintains Container, Application, Model, and Access logs within the application itself.

- The underlying application containers generate logs, which you can access through the hosting Docker or Kubernetes environment.
- The Application logs are part of the application and provide specific access to those respective logs.
- Logs are retained for the customer-defined duration and accessible only to authorized users.
- The application maintains access logs to track user activity, which can be monitored by the customer administrator.
- Model logs record activity for a specific model, accessible in the application UI and reviewable by the model trainer.

The system uses built-in features and container cluster mechanisms to monitor models and platforms regularly.

# 8 Data and Asset Protection

HCL AION employs a multi-layered strategy to protect the confidentiality and integrity of all data and intellectual property managed by the platform. This includes customer data, platform configurations, and the ML models themselves treated as first-class assets. Protection ensures whether the data is in transit across networks, at rest on storage media, or managed within the application.

## 8.1 Credential and Secret Security

- Service credentials stored using hashing with unique salts.
- The use of a unique salt for each secret is critical. It ensures that even if two services had the same password, their stored hashes would be completely different, preventing attackers from identifying duplicate passwords.

## 8.2 ML Model Protection

- Encryption of all trained model files
- **Model Obfuscation**: Complexity added to make reverse-engineering difficult.
- **Digital Watermarking**: Embedded identifiers for intellectual property protection

## 8.3 API & Model Access Security

- **API Key Management**: The system securely generates and rotates API keys.
- **Scoped Access**: The system protects APIs with least-privileged permissions and rate limiting to prevent abuse or overuse.

## 8.4 Responsible AI and Ethical Safeguards

- **Explainability**: The system traces all outputs with logs and explains model outputs to users.
- **Bias Detection:** Models provide statistics for fairness and bias regularly.

# 9 Secure MLOps Lifecycle Management

## 9.1 Data Sourcing & Preparation

- Access control to data sources via RBAC
- Automated data provenance and lineage tracking
- Data poisoning prevention through schema validation and anomaly detection

## 9.2 Model Training

- Encrypted communication for distributed training nodes
- Adversarial robustness testing to pre-empt evasion attacks.

## 9.3 Model Deployment

- All deployments containerized with minimal base images
- Deployment policies enforce runtime security constraints.

# 10 Model Governance

## 10.1 Model Lifecycle Management

- **Version control**: Tracking different versions of models, datasets, and features.
- **Lineage tracking**: Recording data sources, transformations, training parameters.

## 10.2 Data Governance

- **Data quality monitoring**: Checking for missing, corrupt, or skewed data.
- **Data lineage**: Traceability from raw sources to model inputs.
- **Privacy compliance**: Ensures privacy (e.g., through anonymization, consent logging).

## 10.3 Model Explainability & Interpretability

- Use of integrated model-specific explainers.
- Model Cards

## 10.4 Fairness & Bias Monitoring

- **Pre- and post-deployment bias checks** across attributes (gender, race, etc.).
- **Fairness metrics**: e.g., equal opportunity difference, demographic parity.

## 10.5 Model Performance Monitoring

- **Drift detection**: Concept and data drift.
- **Performance degradation tracking**: E.g., AUC, F1-score over time.

## 10.6 Auditability & Traceability

- **Comprehensive logging**: Who trained what, when, using which data and parameters.
- **Audit trails**: For external compliance and internal accountability.

## 10.7 Access Control & Security

- **RBAC**: Role-based access control on models and datasets.
- **Model encryption and secure deployment**.
- Secure APIs and access tokens for inference.

## 10.8 Retraining & Lifecycle Refresh

- Triggered by performance decay or data drift.
- **Appropriate role needed for retraining**.
- Tracking retraining logic, triggers, and audit logs.

# 11 Supply Chain Integrity

Signing of all application and model artifacts

## 11.1 Open source & model supply chain

- **SBOM everywhere.** Generation of SBOMs for libraries.
- **Signed provenance & admission.** Signed images/models and enforced verification.
- **Reproducible training.** Detailed logs available for each step and action taken.
- **Runtime verification.** Image scanning in registry and cluster.

# 12 Threat Management & Incident Response

- Protection against AI-specific attacks such as model theft, inference attacks, and supply chain manipulation
- Dedicated **Product Security Incident Response Team (PSIRT)** with defined responsibilities:
  - **Receipt and Triage:** Serving as the single point of contact for receiving vulnerability reports from customers, security researchers, and internal teams.
  - **Investigation and Coordination:** Investigating all reported vulnerabilities and coordinating with the HCL AION development team to analyze the root cause and impact.
  - **Remediation Planning:** Working with the product team to develop, validate, and release patches or mitigation strategies in a timely manner.
  - **Transparent Communication:** Communicating verified vulnerabilities and their resolutions to customers through official **HCLSoftware Security Bulletins**. These bulletins provide a description of the issue, reference relevant Common Vulnerabilities and Exposures (CVEs), and offer clear, actionable guidance for remediation.

This mature and structured process ensures that security vulnerabilities are managed systematically, effectively, and with the transparency our customers expect. Details on the process can be found here https://www.hcl-software.com/resources/psirt

# 13 Governance, Risk, and Compliance

HCLSoftware's commitment to security is operationalized through a comprehensive Governance, Risk, and Compliance (GRC) program. This program ensures that HCL AION is developed, maintained, and operated in accordance with rigorous internal standards and is validated against globally recognized external benchmarks.

Security is not a separate phase but is deeply integrated into every stage of the HCL AION development lifecycle. This Secure SDLC is aligned with industry best practices and standards such as the NIST Secure Software Development Framework (SSDF). Key security activities and gates are embedded throughout the process:

- **Requirements & Planning:** A Data Privacy Assessment is conducted for all new features to ensure privacy-by-design principles are met.
- **Design:** Formal threat modelling exercises are conducted to identify and prioritize potential threats to the system's architecture. A secure design review is performed by the HCLSoftware Security team to assess the proposed architecture and controls against industry standards.
- **Development:** Developers utilize secure OWASP coding practices, supported by regular training. Code is subjected to multiple layers of automated analysis, including Static Application Security Testing (SAST), Dynamic Application Security Testing (DAST) to detect vulnerabilities in custom code and open-source dependencies.
- **Testing & Validation:** In addition to quality assurance testing, every major release of HCL AION undergoes internal penetration testing. Annually, the platform is subjected to a comprehensive penetration test conducted by an independent, third-party security firm to provide an unbiased assessment of its security posture.
- **Maintenance:** The HCL PSIRT manages ongoing vulnerability monitoring and response.

# 14 Conclusion

HCLSoftware's commitment to security is unwavering. The HCL AION platform developed from the ground up with a deep understanding of the unique security, privacy, and compliance challenges of the modern MLOps landscape. Through a multi-layered strategy rooted in the principles of defense-in-depth, secure-by-design, and Zero Trust, we provide a robust and resilient environment for our customers' most critical AI initiatives.

Every aspect of HCL AION is to build and maintain trust. Our mature governance programs, provide the assurance that our security practices are both comprehensive and effective.

Our valued clients can rest assured that we keep security at the forefront as we develop, evaluate, and deliver powerful and secure MLOps solutions. We view security not as a destination but as a continuous journey of partnership with our customers to navigate the evolving threat landscape.

# 15 Support

For more information and related queries drop an email to: hcl-aion@hcl-software.com

HCLSoftware

hcltechsw.com