

H2O.ai

H2O AI Cloud - Fully Managed

Technical Whitepaper

Contents

Abstract 3

Executive Summary 3

High level architecture 4

Single Tenant Architecture

Region Selection

Zero Trust

Terminology

Infrastructure security 5

Shared Responsibility Model 5

General

Division of responsibility

Defense in Depth 7

Network security 8

Communications Encryption

Inbound Traffic Filtering

Network Security – IP restriction

H2O.ai Support Access

Identity & Access 9

Two-Factor Authentication (2FA)

Secret Management

Encryption 9

Managed Key

Certificate Management

Data Protection 10

Infrastructure redundancy

Metadata Collection

Physical Security 10

Data Governance 11

Region Selection

Role Based Access Control

Application Security 12

Path to production controls

Runtime Protection

Continuous monitoring and incident response 13

Third-party service providers 13

Compliance 13



Abstract

H2O AI Cloud - Fully Managed - Is a managed offering with a complete set of H2O Cloud capabilities on a fully managed cloud environment that handles infrastructure provisioning, scale, and software updates. H2O AI Cloud - Fully Managed is designed for high availability and dependability, providing tools to develop, deploy, manage and maintain AI based applications. Helping to protect the confidentiality, integrity, and availability of your systems and data is our top priority. This document is intended to provide an introduction to H2O.ai's approach to security, including the environment's architecture and controls.

Executive Summary

H2O AI Cloud - Fully Managed approach to cyber security is a multi-layer strategy to provide multiple independent lines of defence against cyber attacks.

Each layer is taking into account the following security pillars: Identify, Protect, Detect, Respond and Recover.

Key security features:

- Single tenant architecture. Each customer is running in a dedicated account, and networking zone, completely isolated from other customers accounts.
- Customer data never leaves the region selected.
- Data is encrypted at-rest and in-transit
- Access to the production environment is enforced with a 'zero trust' policy.
- Access to production is constantly logged and monitored.

High Level Architecture

Single Tenant Architecture

H2O AI Cloud - Fully Managed is designed as a single tenant per account architecture, meaning every customer is provided with a dedicated account, completely isolated from other customers accounts.

Region Selection

Upon onboarding, the customer is asked to choose a region from the H2O available region list (page 11) from that point on, customer data will never leave that region selected.

Zero Trust

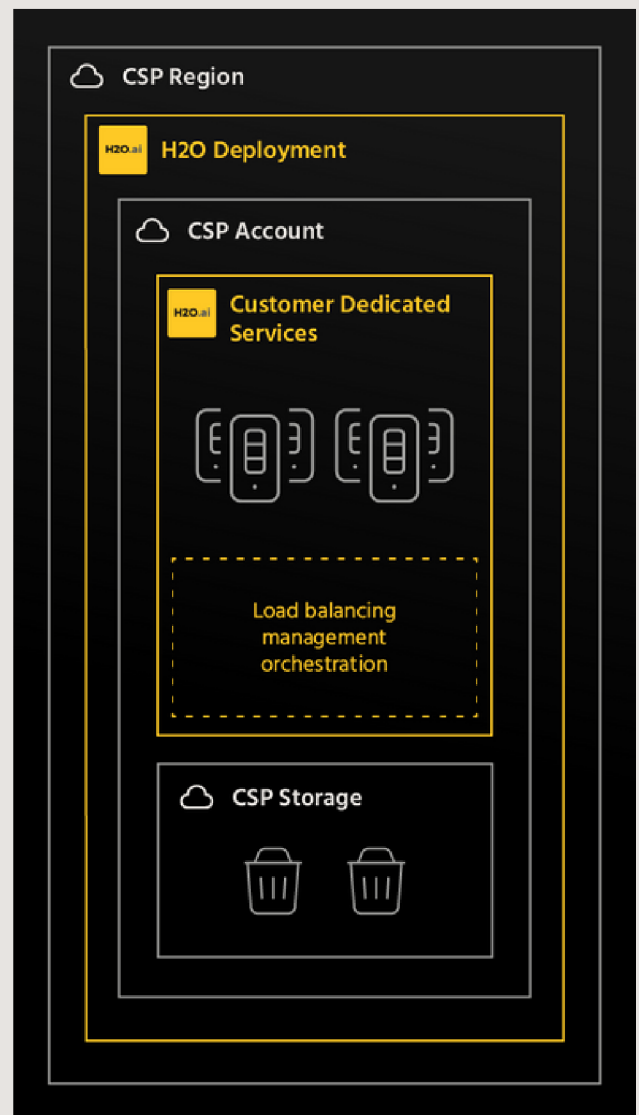
All access to the production environment of H2O AI Cloud - Fully Managed is enforced with a 'zero trust' policy.

Terminology

CSP
Cloud service provider

CSP Region
The region the customer selected to deploy his workload.

CSP Account
The account that was created upon customer's registration.



Infrastructure Security

H2O AI Cloud - Fully Managed has been architected to be a flexible and secure cloud computing environment. It is designed to provide an extremely scalable, highly reliable platform that enables customers to develop and deploy fully fledged AI solutions securely. This platform is built and managed not only according to security best practices and standards, but also with the unique needs of the cloud in mind. H2O.ai uses redundant and layered controls, continuous validation and testing, and a substantial amount of automation to ensure that the underlying platform is monitored and protected 24x7. This means that you get a resilient infrastructure, designed for high security, without the capital outlay and operational overhead of a traditional data center. H2O.ai operates under the shared security responsibility model.

Shared Responsibility Model


General

As you consider and evaluate our Fully Managed Cloud offering, it's critical to understand the shared responsibility model and which security tasks are handled by the H2O and which tasks are handled by you. The workload responsibilities vary depending on whether the workload is hosted on H2O AI Cloud - Fully Managed or in an H2O AI Cloud - Hybrid installation.

Division of Responsibility

In an H2O AI Cloud - Hybrid cloud installation, you own the whole stack. As you move to the Fully Managed cloud environment some responsibilities transfer to H2O. The following diagram illustrates the areas of responsibility between you and H2O, according to the type of deployment of your stack.

For the H2O AI Cloud - Managed deployment type, you own your data and identities, on-premises resources, and the cloud components you control.



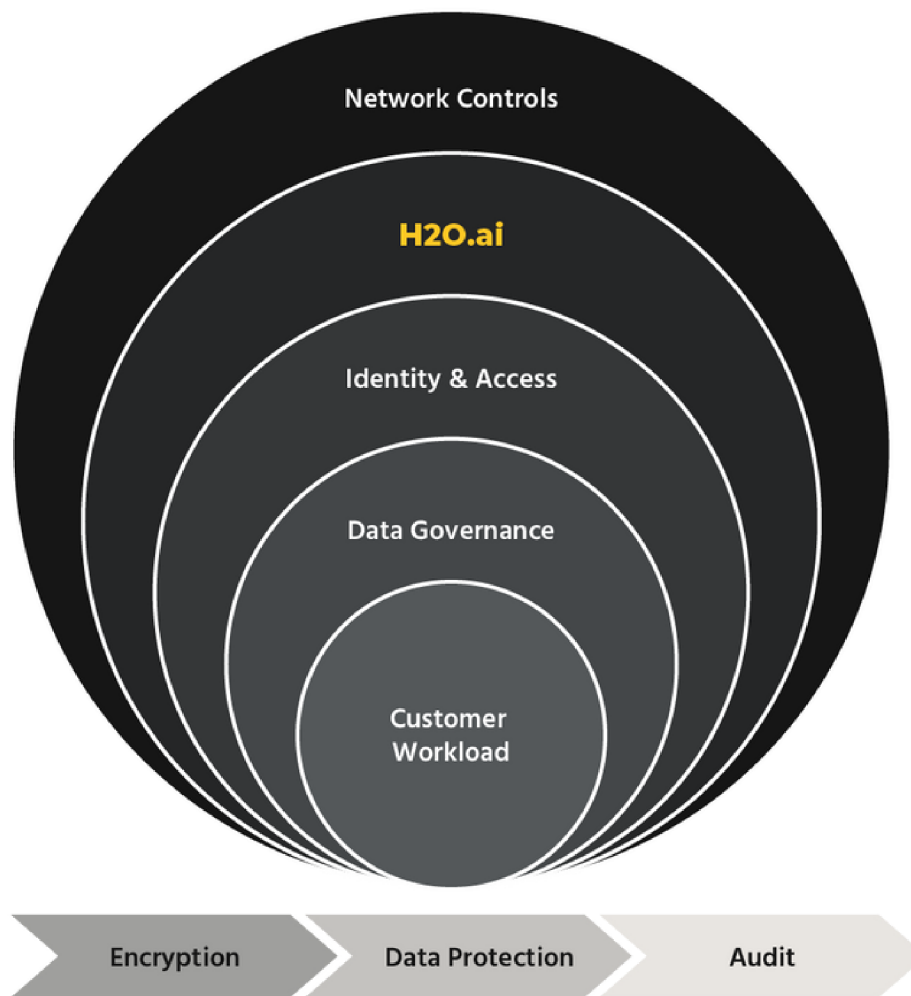
Regardless of the type of deployment, the following responsibilities are always retained by the customer:

- Information and data
- Devices
- Account and identities

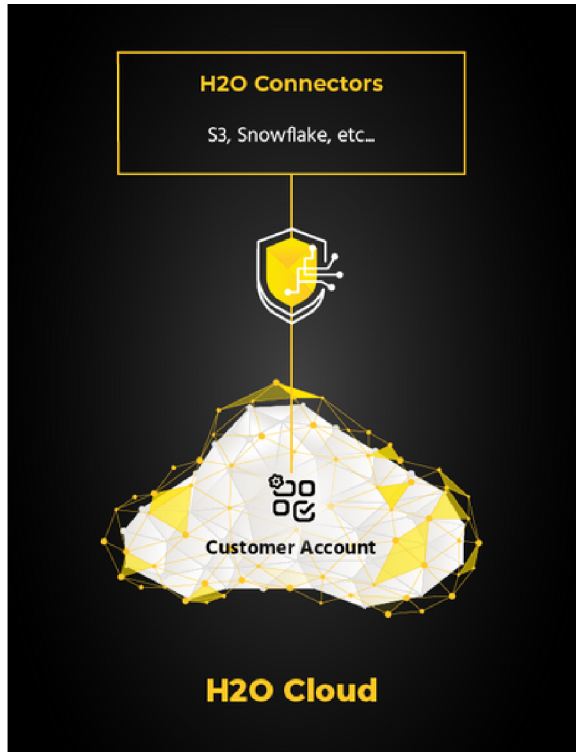
H2O.ai responsible	Identity & access management	<ul style="list-style-type: none"> • Provides a secure and redundant solution for the following capabilities • Multi factor authentication and role-based access control • Enforce the Use of two-factor authentication
	Application	<ul style="list-style-type: none"> • Promote the secure design and implementation of H2O.ai infrastructure, platform, and applications developed by H2O.ai.
	Identity & access management	<ul style="list-style-type: none"> • Manage inbound firewall rules • Data encryption at-rest • Data encryption in-transit for all public-facing interfaces and cross component traffic
Customer responsible	Information and data	<ul style="list-style-type: none"> • Uploading the proper data in accordance to organization policy business needs • Manage data lifecycle in H2O managed cloud • Delete data that is no longer in accordance with organization policy business needs • Data classification
	Devices	<ul style="list-style-type: none"> • Manage and protect the devices/endpoints that have access to H2O Managed Cloud
	Identity & access management	<ul style="list-style-type: none"> • Establishing and configuration of user identities • Assign proper roles and permissions
	Application	<ul style="list-style-type: none"> • Promote the secure design and implementation of custom code • Promote the secure sourcing, deployment, and maintenance of third-party integrations and dependencies • Comply with relevant security standards and regulations. • Consume and analyze audit logs in a timely manner • Deploy anti-abuse, fraud detection, and prevention measures • Ensure continuous monitoring and incident response on customer and custom third-party integration assets
	Network control	<ul style="list-style-type: none"> • Restrict access controls by using IP restriction

Defense in Depth

H2O AI Cloud - Fully Managed implements a defense in depth strategy using multiple layers of security controls to defend the platform against cyber attacks using several independent layers. Each layer is taking into account the following security pillars Identify, Protect, Detect, Respond, and Recover.



Network Security



Communications Encryption

All client communications are encrypted using TLS 1.2. Successful authentication is required for all connections by default.

Inbound Traffic Filtering

H2O.ai uses a variety of tools and techniques to implement ingress traffic inspection to prevent DDoS attacks and common application layer attacks.

Network Security

IP Restriction

Allows customers to restrict access to the H2O.ai control plane (inbound) in their H2O.ai cloud account to only clients originating from a list of approved IP ranges that the organization has defined.

URL Blacklist

Domains known to be malicious or have a high risk of malware infection. H2O.ai continuously maintains and updates the URL blacklist and keeps it up to date with recent malicious domains.

H2O.ai Support Access

For support purposes H2O.ai personnel may require access to customers' accounts. H2O.ai personnel with access to the client production environment is limited to specific engineers all using 2FA to achieve zero trust. Access is controlled in a JIT (just in time) manner with an internal approval process and limited to the necessary permissions needed. All access to the production environment is logged and monitored.



Identity & Access

Two-Factor Authentication (2FA)

H2O.ai provides a local user directory. Login to the H2O AI Cloud - Fully Managed environment enforces 2FA by default for an extra layer of security. 2FA helps protect against common security attacks like credential stuffing and brute-force attacks on passwords.

Secret Management

An app's environment-specific configuration, such as credentials, should not be stored in version control systems, but should instead be provided via environmental variables. H2O provides customers to store these sensitive environment-specific configurations.

Encryption

Managed Key

By default, all files and/or volumes stored in internal storage are encrypted at rest.

Certificate Management

H2O AI Cloud - Fully Managed automatically manages TLS certificates for all infrastructure including customer apps.

Data Protection



Infrastructure Redundancy

H2O AI Cloud - Fully Managed relies on the cloud provider's infrastructure for redundancy and utilizes database replication and application failover across 2 availability zones in the same region.

File storage is replicated across all of the cloud provider's availability zones in the region at a minimum of 2 availability zones per region.

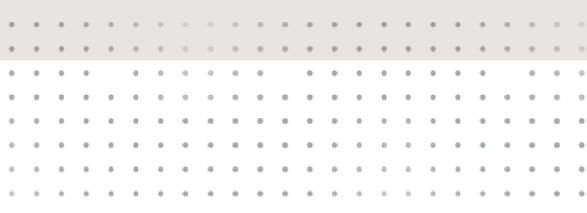
*Currently Steam and parts of H2O MLOps applications, along with certain individual dynamic workloads such as DAI/Wave app instances, are not highly available

Metadata Collection

H2O AI Cloud - Fully Managed collects metadata ("Systems Information"). Systems Information may be used by H2O to operate and maintain (including to improve) the Service.

Physical Security

H2O AI Cloud - Fully Managed relies on the cloud provider's datacenter infrastructure for physical security controls.



Data Governance

Region Selection

You can specify a region from the H2O region list. Customer data never leaves the region selected.

- **North America - N. Virginia, Oregon**
- **South America - São Paulo**
- **Asia Pacific - Sydney, Mumbai, Singapore**
- **Europe - Ireland**



Role Based Access Control

H2O AI Cloud - Fully Managed authorizes all access based on the OIDC identity of the user. The authorization is either governed by implicit rules based on the "role" of each user (admin vs regular user) and their relationship to a protected resource (resource owner vs not an owner) or via custom IAM policies. The customer is in control of user roles, as well as custom IAM policies.

Application Security


Path to Production Controls

H2O implements the following controls as part of the software development lifecycle.

- Code review
- Dependency scanning for vulnerabilities
- Dependency licenses check
- Final product artifact scan for vulnerabilities (docker images)
- Pen-testing in the integrated environment.
- Security scan of infrastructure.
- Penetration testing is performed on an annual basis.

Runtime Protection

- H2O AI Cloud - Fully Managed is evaluated quarterly for new vulnerabilities. All new vulnerabilities will be addressed and prioritized based on H2O risk management policy.
- H2O Implement a CSPM solution for detecting and remediating misconfigurations and threats that lead to data breaches and compliance violations.
- H2O Uses a variety of tools to achieve a defense in depth security strategy.





Continuous Monitoring and Incident Response

H2O.ai ensures continuous monitoring and incident responses on its assets.

H2O.ai will notify customers only in the event of a security breach with a direct impact to the customers confidentiality, integrity or availability. H2O.ai will notify only the affected customers via email notification to all users in the administrator role.

Third-party service providers

- H2O.ai 3rd party providers undergo an annual security review and risk assessment process.
- H2O.ai uses only well known 3rd party providers.
- Every new service provider will undergo an H2O.ai internal vendor approval process.

Compliance

H2O AI Cloud - Fully Managed is SOC2 Type2 certified. As mentioned above our cloud was designed with security as our top priority.