

# Managed Services Shared Responsibility Model

---

Databricks  
January 2023





# Databricks Managed Services Shared Responsibility Model

Security and compliance are a shared responsibility between Databricks, the Databricks customer, and the cloud service provider (CSP) such as AWS, GCP, or Azure. For their part, the CSPs also have formalized their shared responsibility models ([AWS](#), [Azure](#), [GCP](#)).

## Databricks Responsibilities

### Databricks Platform and Services

- Secure the Databricks Control Plane
- Utilize industry standards to harden images and operating systems deployed under our control
- Maintain a public bug bounty program
- Maintain the Databricks Control Plane with updated code and images

### Databricks Managed Resources

- Securely deploy and terminate Databricks managed systems
- Track security configurations against industry standard baselines for systems under Databricks control
- Deploy the latest applicable source code and system images upon launch of customer Data Plane hosts

### Identity and Access Management

- Authenticate Databricks personnel using industry best practices
- Set employee privileges consistent with least privilege principles
- Limit access to systems processing customer data to employees with roles that warrant access
- Restricts access to customer content based on the principle of least privilege and segregation of duties
- Secure interactions with the customer-managed cloud account
- Secure storage and policy enforcement of secrets scope

## Customer Responsibilities

### Account and Workspace Management

- Manage account responsibilities, including account setup and administration, subscription management and cloud resources ([AWS](#), [Azure](#), [GCP](#))
- Workspace management, including workspace creation, update, and deletion, and workspace resource access ([AWS](#), [Azure](#), [GCP](#))

### Cluster Policies

- Configure cluster management policies, personal compute policies ([AWS](#), [Azure](#), [GCP](#))

### VM Instance Management

- Restart workspace cluster VMs to deploy the latest patched images and code in accordance with patch management policy ([AWS](#), [Azure](#), [GCP](#))

### Identity and Access Management

- Setup Single Sign-on and password access controls for Databricks account and workspace(s) ([AWS](#))
- Enable multi-factor authentication via your SSO provider
- Enable System for Cross-domain Identity Management (SCIM) integration with your identity provider ([AWS](#), [Azure](#), [GCP](#))

### Identity, Service Principal and Access Management

- Manage users, groups, personal access tokens, and service principals ([AWS](#), [Azure](#), [GCP](#))
- Set Access Control Lists to restrict resource access (such as workspace objects, clusters, pools, jobs, tables) ([AWS](#), [Azure](#), [GCP](#))
- Set up the cross-account IAM role using least-privilege principles ([AWS](#))
- Secure management and use of secret scopes ([AWS](#), [Azure](#), [GCP](#))

## Cloud Responsibilities

### Cloud Service Platform and Services

- Maintain security of the cloud service infrastructure
- Maintain a security management program that maintains reasonable security measures to protect customer data and services

### Identity and Access Management

- Maintain access controls required to restrict access to authorized customer resources
- Restrict employee access to customer resources



## Platform Security



## IAM Security





# Databricks Managed Services Shared Responsibility Model

Security and compliance are a shared responsibility between Databricks, the Databricks customer, and the cloud service provider (CSP) such as AWS, GCP, or Azure. For their part, the CSPs also have formalized their shared responsibility models ([AWS](#), [Azure](#), [GCP](#)).

## Databricks Responsibilities

### Databricks Managed Data

- Transmit customer content using TLS 1.2 or higher between the Customer and the Databricks Control Plane and the Databricks Control Plane and the Data Plane
- Encrypt customer data-at-rest within the Databricks Control Plane using AES-128 bit equivalent or higher
- Delete customer content contained within a customer workspace within thirty (30) days of the workspace cancellation



Data Security

### Secure Network Communications

- Separate the Databricks Control Plane from the Customer Data Plane and workspaces within the Databricks Data Plane using multiple layers of network security controls
- Deploy local firewalls or security groups within the Customer Data Plane to isolate clusters
- Enable secure defaults for network access controls and security groups within the Control Plane



Network Security

## Customer Responsibilities

### Data Governance

- Enable [Unity Catalog](#) within your Databricks account
- Follow [data governance](#) best practices, as per your organization's requirements ([AWS](#), [Azure](#), [GCP](#))

### Customer-managed Data

- Secure management of data infrastructure ([AWS](#), [Azure](#), [GCP](#)):
  - Secure connectivity to customer-managed resources
  - Secure service integration with Databricks ([AWS](#), [Azure](#), [GCP](#))
  - [Azure] Enable Data Plane [local disk encryption](#) or [Spark inter-cluster encryption](#)

### Customer-managed Encryption Keys

- Deploy customer-managed encryption keys (CMK) ([AWS](#), [Azure](#))
  - Enable CMK for managed services
  - Enable CMK for workspace storage

### Cloud Network Security

- Configure Secure Cluster Connectivity ([AWS](#), [Azure](#), [GCP](#))
- Enable customer-managed networks ([AWS VPC](#), [Azure VNet](#), [GCP VPC](#))
- Configure Data Exfiltration Protection according to your organization's data protection policy ([AWS](#), [Azure](#), [GCP](#))

### IP Access Control Lists and Private Link

- Configure Databricks workspace IP access lists ([AWS](#), [Azure](#), [GCP](#))
- Configure Private Link access for Users → Control Plane and Control Plane → Data Plane connections ([AWS](#), [Azure](#))

## Cloud Responsibilities

### Cloud Service Managed Data

- Maintain encryption hardware and services
- Encrypt data in transit and at rest, where configured
- Maintain the confidentiality, integrity and availability of data stored on CSP services
- Enable Spark inter-cluster encryption (GCP, Subset of AWS Nitro that [support in-transit encryption](#))
- Enable Data Plane local disk encryption (GCP, AWS Nitro or NVMe)

### Secure Network Communications

- Secure the physical and logical security of cloud service networking
- Maintain secure network communications for cloud services, including APIs





# Databricks Managed Services Shared Responsibility Model

Security and compliance are a shared responsibility between Databricks, the Databricks customer, and the cloud service provider (CSP) such as AWS, GCP, or Azure. For their part, the CSPs also have formalized their shared responsibility models ([AWS](#), [Azure](#), [GCP](#)).

## Databricks Responsibilities

### Security Monitoring

- Deploy security detection capabilities, including those provided natively by Cloud Service Providers
- Generate audit logs from customer’s use of the platform services and retain them for at least one year
- Deliver audit logs from the customer’s use of the platform services based on the customer’s configuration (Premium subscriptions and above)
- Deploy a dedicated Detection engineering team that develops intrusion detection monitoring across its computing resources
- Employ an incident response framework to manage and minimize the effects of unplanned security events
- Inform customers of security breaches in accordance with data protection laws

### Secure Code Execution

- Maintain availability and security of the job scheduler
- Secure delivery of customer code (such as notebooks, repos and models, queries) from the control plane to the data plane

## Customer Responsibilities

### Audit Log Configuration

- Configure Databricks [audit log delivery](#) to your cloud storage ([AWS](#), [Azure](#), [GCP](#))
- Configure [verbose](#) audit logs for your workspace(s) ([AWS](#), [Azure](#), [GCP](#))

### Account and Workspace Security Monitoring

- Deploy account and workspace [security monitoring](#)
- Deploy cloud service security monitoring
- Investigate and respond to potential security incidents related to customer-managed features, services and resources

### Application Security

- Perform security reviews of your code, libraries and jobs, such as notebooks ([AWS](#), [Azure](#), [GCP](#)), [Terraform](#), and third-party libraries ([AWS](#), [Azure](#), [GCP](#))

### CI/CD Pipeline and Repo Integration

- Integrate Git with Databricks repos ([AWS](#), [Azure](#), [GCP](#))
- Manage CI/CD Pipeline integration with Databricks ([AWS](#), [Azure](#), [GCP](#))

## Cloud Responsibilities

### Security Monitoring

- Monitor for security violations of the underlying cloud service infrastructure and services
- Deliver audit logs for cloud service events based on customer configurations
- Employ an incident response framework
- Notify customer of a security breach for which that customer is impacted

### Secure Code Execution

- Maintain cloud infrastructure



Security Monitoring



Code Execution / Jobs





# Databricks Managed Services Shared Responsibility Model

Security and compliance are a shared responsibility between Databricks, the Databricks customer, and the cloud service provider (CSP) such as AWS, GCP, or Azure. For their part, the CSPs also have formalized their shared responsibility models ([AWS](#), [Azure](#), [GCP](#)).

## Databricks Responsibilities

## Customer Responsibilities

## Cloud Responsibilities



### Core Compliance

#### Standards and Compliance

- Maintain independent third-party audits, standards, and certifications that apply to all customer environments:
  - ISO 27001, 27017, 27018
  - SOC 2 Type II, SOC 1 Type II, SOC 3
- Provide tools and configurations that enable use of services in compliance with applicable laws (such as GDPR and CCPA)


*\* Additional compliance standards covered later, such as HIPAA, FedRAMP, PCI*

#### Maintain Adherence to Relevant Compliance and Standards:

- When using Databricks to process sensitive data such as PII, adhere to relevant privacy regulations such as the GDPR and CCPA
- Review your compliance needs and add optional compliance service offering where required (such as for FedRAMP, PCI-DSS, HIPAA)
- Comply with applicable laws when using Databricks, including by implementing any required configurations in accordance with Databricks documentation

#### Standards and Compliance

- Maintain independent third party audit, standards and certifications
- Maintain relevant independent third-party audits, standards, and certifications
- Maintain relevant compliant services



### Disaster Recovery

#### Maintain Disaster Recovery Capabilities\* For:

- Review Business Continuity and Disaster Recovery plans annually
- Conduct Business Continuity and Disaster Recovery drills annually
- Conduct periodic backups of the Databricks Control Plane\*

#### Data Backups

- Backup of your organization's [account and workspace](#)
- Set [Recovery Point Objectives](#) (RPO) and [Recovery Time Objectives](#) (RTO) using best practices ([AWS](#), [Azure](#), [GCP](#))

#### Disaster Recovery capabilities

- Cloud service capacity
- Review Business Continuity and Disaster Recovery plans annually
- Conduct Business Continuity and Disaster Recovery drills annually
- Conduct periodic backups



### Security Best Practices

#### Employ Security Best Practices

- Periodically review cryptographic standards to select and update technologies and ciphers in accordance with assessed risk and market acceptance of new standards
- Maintain a vulnerability management program that follows industry best practices
- Conduct third-party penetration tests at least annually
- Employ an in-house offensive security team

#### Multi-region Workspace Deployment

- Adopt Databricks security best practices based on the organization's cybersecurity requirements ([AWS](#), [Azure](#), [GCP](#))
- Follow security best practices for the customer's cloud environment ([AWS](#), [Azure](#), [GCP](#))

#### Employ Security Best Practices

- Review cryptographic standards
- Regularly run authenticated vulnerability scans
- Address vulnerabilities within SLAs
- Conduct third-party penetration tests
- Employ in-house offensive security personnel





# Databricks ESM/CSP Shared Responsibility Model



# Databricks Managed Services Shared Responsibility Model

Security and compliance are a shared responsibility between Databricks, the Databricks customer, and the cloud service provider (CSP) such as AWS, GCP, or Azure. For their part, the CSPs also have formalized their shared responsibility models ([AWS](#), [Azure](#), [GCP](#)).

## Databricks Responsibilities

## Customer Responsibilities

## Cloud Responsibilities



### Enhanced Security Monitoring

#### Databricks Enhanced Security Monitoring (ESM) Responsibilities

- Deploy ESM VMs with enhanced CIS Level 1 hardening
- Deploy antivirus, behavior-based malware monitoring and file integrity monitoring
- Provide Qualys vulnerability reports of the host OS upon request
- Enable FIPS 140-2 Level 1 encryption modules where available

#### Customer Enhanced Security Monitoring Responsibilities

- Enable [Enhanced Security Monitoring](#) on relevant workspace(s)
- Monitor [enhanced event logs](#) for security incidents
- Restart ESM clusters to deploy the latest patched VM versions and agent signatures
- Provide the destination Email for vulnerability reports delivery

#### CSP ESM Responsibilities

- Maintain security of the cloud service infrastructure



### Compliance Security Profile

#### Databricks Compliance Security Profile (CSP) Responsibilities

- Enable ESM security enhancements (listed above)
- Enforcement of AWS Nitro VMs on CSP workspace(s)
- Cluster update enforcement (auto-restart clusters after 25 days)
- Enumerate preview features that are usable within HIPAA, PCI, FedRAMP

#### Customer Compliance Security Responsibilities

- [Prepare](#) a workspace for the compliance security profile
- Enable the Compliance Security Profile on relevant workspace(s) ([AWS](#))
- All customer responsibilities from ESM

#### CSP Compliance Responsibilities

- Maintain security of the cloud service infrastructure



### HIPAA, PCI, DoD and FedRAMP

#### Databricks HIPAA, PCI and FedRAMP Responsibilities

- Complete annual HIPAA, PCI-DSS, FedRAMP audits ([region and cloud specific](#))
- Implement HIPAA, PCI and FedRAMP (Moderate on AWS, High on Azure Government) compliant services
- Maintain DoD Impact Level 2, 4 and 5 (IL2, IL4, IL5) (Azure Gov. only)
- Enforce Enterprise Security Monitoring and Compliance Security Profile features

#### Customer HIPAA, PCI, FedRAMP Responsibilities

- Enable [CSP](#) on relevant workspaces (AWS)
- Use only supported preview features (AWS: [PCI](#), [HIPAA](#), GCP: [HIPAA](#))
- Follow compliance-specific prerequisites:
  - Detailed docs: AWS: [HIPAA](#), [PCI](#), [FedRAMP](#), GCP: [HIPAA](#), Azure: [IL5](#)
  - Maintain a Business Associate Agreement w/ Databricks & your cloud provider (HIPAA)
  - Obtain entitlement to process regulated data from Databricks
  - Review the PCI Shared Responsibility Model requirements (PCI)
  - Follow FedRAMP PMO documentation requirements ([FedRAMP](#))

#### CSP HIPAA, PCI and FedRAMP Responsibilities

- Complete annual HIPAA, PCI-DSS, FedRAMP audits
- Implement HIPAA, PCI and FedRAMP compliant services
- Maintain DoD Impact Level 2, 4 and 5 (IL2, IL4, IL5)



### GDPR/CCPA

#### Databricks GDPR/CCPA Service Responsibilities

- Provide GDPR/CCPA compliance services
- Define and adhere to a [Data Processing Agreement](#) with GDPR [Standard Contractual Clauses](#)

#### Customer GDPR/CCPA Service Responsibilities

- Maintain GDPR/CCPA compliance usage of Databricks services

#### CSP GDPR/CCPA Responsibilities

- Provide GDPR/CCPA compliance services

