databricks

CUSTOMER COMPLIANCE MAPPING

Mapping FISC Security Guidelines to Databricks' Customer Capabilities

Version	
1.0	FISC Security Guidelines on Computer Systems for Financial Institutions, Thirteenth Edition (March 2025)

Contents

The FISC Security Guidelines on Computer Systems for Financial Institutions are organized into four primary divisions. This document walks each guideline unit and maps Databricks Data Intelligence Platform capabilities that help in-scope financial institutions meet the relevant control objectives.

C	DIVISION Control Guidelines	C1 - C28
P	DIVISION Practice Guidelines	P1 - P153
F	DIVISION Facilities Guidelines	F1 - F138
A	DIVISION Audit Guidelines	A1 - A1-1

Three-column reading guide. Each row pairs a FISC guideline number and category with the corresponding Databricks customer-capability mapping.

Shared responsibility. Refer to the Databricks Shared Responsibility Model alongside this guide. Customers remain accountable for organizational policies, contracts, and ongoing oversight.

N/A entries. Some guidelines fall outside the Databricks platform boundary (physical premises, branch operations, customer-managed procedures). Those rows are noted for completeness.

Cybersecurity & AI. The Thirteenth Edition strengthens cybersecurity requirements and introduces dedicated AI guidelines; both are mapped to Databricks platform capabilities.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
C Control Guidelines		
C1	POLICY / PLAN	N/A — establishing the organization's own security regulations is the financial institution's responsibility. The cybersecurity-specific sub-units C1-1 and C1-2 are addressed in the rows below.
C1-1	POLICY / PLAN	▪ Policy ownership sits with the customer. Databricks supports operationalizing cybersecurity policy through the Databricks Security & Trust Center, the Shared Responsibility Model, and documented platform security baselines. ▪ Enforce policy technically with account- and workspace-level security settings, the Compliance Security Profile (CSP), and Enhanced Security Monitoring (ESM) for regulated workloads.
C1-2	POLICY / PLAN	▪ Codify cybersecurity rules as enforceable configuration: compute policies, Unity Catalog governance rules, IP access lists, and network/egress policies. ▪ Manage security configuration as code (Databricks Terraform provider and Declarative Automation Bundles (DABs)) so rules are version-controlled, peer-reviewed, and applied consistently across workspaces.
C2	POLICY / PLAN	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C3	POLICY / PLAN	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C4	ORGANIZATIONAL STRUCTURE	N/A — establishing the security management system is the financial institution's responsibility. The cybersecurity-specific sub-units C4-1 and C4-2 are addressed in the rows below.
C4-1	ORGANIZATIONAL STRUCTURE	▪ Structure cybersecurity accountability using Databricks role separation — account admins, workspace admins, and Unity Catalog metastore admins — mapped to IdP groups via SCIM. ▪ Management and HR planning remain customer responsibilities; Databricks provides the role-based entitlements and admin governance to support them.
C4-2	ORGANIZATIONAL STRUCTURE	▪ Continuously check platform security posture with system tables (e.g., system.access.audit and system.compute.clusters) and the Databricks Security Analysis Tool (SAT), which benchmarks workspace configuration against security best practices. ▪ Forward audit and configuration telemetry to your SIEM for ongoing control-effectiveness monitoring and alerting.
C5	ORGANIZATIONAL STRUCTURE	▪ Missing number — C5 was removed in the 13th Edition and replaced by C5-1 to C5-5 (below).

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
C5-1	ORGANIZATIONAL STRUCTURE	▪ Unity Catalog provides a central, governed inventory of data and AI assets (catalogs, schemas, tables, volumes, models, functions) with automated lineage and discovery. ▪ Use tags, Data Classification, and system.information_schema to identify and label sensitive assets so cybersecurity risk can be assessed against a complete asset inventory.
C5-2	ORGANIZATIONAL STRUCTURE	▪ Risk assessment and response planning are customer-owned processes; Databricks supplies the inputs — asset inventory and lineage (Unity Catalog), posture scoring (SAT), and threat telemetry (system tables, ESM). ▪ Response actions can be executed on-platform: revoke tokens and grants, disable principals, rotate customer-managed keys, and restrict or revoke access to affected data via Unity Catalog.
C5-3	ORGANIZATIONAL STRUCTURE	▪ Databricks manages platform vulnerabilities: runtimes are regularly patched, and the Compliance Security Profile enforces CIS-hardened, FIPS 140-validated images with automatic security updates and automatic cluster update (VM restart). ▪ Constrain customer-introduced software with compute policies and allowlisted libraries; monitor Databricks security bulletins and use ESM host-level telemetry to detect vulnerable components.
C5-4	ORGANIZATIONAL STRUCTURE	▪ Cybersecurity exercises (tabletop and disaster-recovery drills) are customer-owned. Databricks supports failover and DR rehearsals through multi-region workspace deployments and cross-region replication of the underlying cloud storage (e.g., S3 CRR) together with Delta Deep Clone, configured by the customer.
C5-5	ORGANIZATIONAL STRUCTURE	▪ Security-awareness education for the customer's workforce is a customer responsibility. Databricks provides product security documentation, the Security & Trust Center, and administrator training resources.
C6	ORGANIZATIONAL STRUCTURE	▪ Databricks supports system management through account and workspace administration, compute policies, and configuration-as-code (Databricks Terraform provider and Declarative Automation Bundles (DABs)). Unity Catalog centralizes governance of data and AI assets across workspaces. ▪ System tables and audit logs give a single operational view of compute, jobs, and configuration for ongoing management. Establishing the management structure and policies remains the customer's responsibility.
C7	ORGANIZATIONAL STRUCTURE	▪ Unity Catalog is Databricks' data management system: a central catalog for data and AI assets with fine-grained access control, automated lineage, tagging, and Data Classification across catalogs, schemas, tables, volumes, models, and functions. ▪ Delta Lake adds ACID transactions, schema enforcement, and versioning; system tables provide centralized visibility. The customer defines its own data management policies on top of these capabilities.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
C8	ORGANIZATIONAL STRUCTURE	▪ Databricks supports network management through customer-managed VPC/VNet, PrivateLink and private endpoints, IP access lists, and network policies for serverless compute — all configurable as code. ▪ Network configuration lives in the customer's cloud account; Databricks provides the controls and connectivity options, while the network management structure and policies are owned by the customer.
C9	ORGANIZATIONAL STRUCTURE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C10	ORGANIZATIONAL STRUCTURE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C11	ORGANIZATIONAL STRUCTURE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C12	ORGANIZATIONAL STRUCTURE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C13	EVALUATION OF MANAGEMENT STATUS	▪ Confirm ongoing security observance using audit logs and system tables (sign-ins, grants, configuration changes) streamed to a SIEM, and the Databricks Security Analysis Tool (SAT), which benchmarks workspace configuration against security best practices. ▪ Unity Catalog access reviews, the token report, and independent SOC/ISO/PCI attestations support periodic verification. The review cadence and process are defined by the customer.
C14	HUMAN RESOURCES (STAFFING / TRAINING)	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C15	HUMAN RESOURCES (STAFFING / TRAINING)	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C16	HUMAN RESOURCES (STAFFING / TRAINING)	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C17	HUMAN RESOURCES (STAFFING / TRAINING)	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C18	HUMAN RESOURCES (STAFFING / TRAINING)	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
C19	HUMAN RESOURCES (STAFFING / TRAINING)	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
C20	OUTSOURCING MANAGEMENT	N/A — there is no Databricks platform feature for this outsourcing-management requirement; the financial institution addresses it through its own procurement and vendor-selection governance.
C21	OUTSOURCING MANAGEMENT	N/A — there is no Databricks platform feature for this outsourcing-contract requirement; it is addressed in the customer's own contracts (see the Databricks MCSA and Security Addendum).
C22	OUTSOURCING MANAGEMENT	N/A — there is no Databricks platform feature for this requirement; the financial institution manages its own outsourcing, risk assessments, and ongoing oversight.
C23	OUTSOURCING MANAGEMENT	N/A — there is no Databricks platform feature for this outsourcing-management-structure requirement.
C24	USE OF CLOUD SERVICES	▪ Configure Databricks with cloud-specific controls: customer-managed VPC/VNet, PrivateLink, CMK, and cloud-native IAM roles for least-privilege access. ▪ Unity Catalog unifies governance and workspace isolation across multi-cloud deployments. For regulated workloads, enable ESM/CSP compute profiles.
C25	SHARED DATA CENTER	▪ Databricks does not own or operate the underlying physical data centers; production services run on AWS, Azure, and GCP. Physical security (biometric access, CCTV, power redundancy) is managed by each CSP under ISO 27001, SOC 1/2, and PCI-DSS; Databricks reviews CSP audit reports on a regular cadence. ▪ Databricks security commitments are set out in the MCSA and Security Addendum, independently verified annually. See the Databricks Trust Center.
C26	SERVICES ON FINANCIAL INSTITUTIONS' MUTUAL SYSTEM NETWORK	▪ Connect mutual financial network data via private VPC/VNet with full audit logging. Unity Catalog enforces access control and lineage for mutual-network data. ▪ Use Delta Sharing and Clean Rooms for secure cross-institution analytics with strict row/column- level permissions and no data copying. ▪ For tighter runtime security, run workloads on ESM/CSP-enabled clusters.
C27	COOPERATION WITH FINTECH COMPANIES	N/A — there is no Databricks platform feature for this FinTech-cooperation requirement.
C28	SUPPLY CHAIN MANAGEMENT	▪ Databricks manages its own supply chain (sub-processors, third-party components) and evidences it through SOC 1/2, ISO 27001/27017/27018, and PCI-DSS attestations in the Trust Center; customers inherit these assurances under the shared responsibility model. ▪ Govern your own data and AI supply chain on-platform: control third-party data sources and models through Unity Catalog, Lakehouse Federation, Partner Connect, and Marketplace, with lineage and access controls applied uniformly.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P	**Practice Guidelines**	
P1	DATA PROTECTION	• Integrate with your IdP via SSO with MFA enforced at the IdP. Use SCIM for user lifecycle management. • Replace long-lived PATs with OAuth token federation for both users and service principals. Unity Catalog enforces least-privilege data access based on these identities.
P2	DATA PROTECTION	N/A — there is no Databricks platform feature for this requirement; Databricks does not interact directly with terminals on public networks.
P3	DATA PROTECTION	• Enable at-rest encryption on cloud storage (S3/ADLS/GCS); use CMK for workspace and storage where required. Unity Catalog enforces RBAC, row filters and column masking, with catalog-to- workspace binding to restrict sensitive data to approved workspaces. • Register external databases (PostgreSQL, Snowflake, etc.) as Lakehouse Federation foreign catalogs to apply the same UC policies without copying data. Use Data Classification and ABAC to automatically tag and protect sensitive columns.
P4	DATA PROTECTION	• Enforce TLS 1.2+ for all connections; use PrivateLink and private endpoints to keep traffic off the public internet. Apply IP access lists and context-based ingress policies. • Use network policies (Serverless Egress Control) to enforce a deny-by-default outbound posture across serverless compute (notebooks, jobs, SQL, and model serving); some deny-by-default controls are in preview. Store credentials in Databricks Secrets rather than in code. • Apply Unity Catalog least-privilege and row/column-level controls. Forward audit logs to your SIEM and alert on anomalous download patterns or egress denials.
P5	DATA PROTECTION	• Use Unity Catalog as the primary access layer: expose data as tables and volumes, granting fine- grained privileges (SELECT, MODIFY, OWNERSHIP) to approved users and service principals only. • Combine with workspace ACLs and cloud-storage IAM. Enable audit logs and system tables to record object access. Run regulated workloads on ESM/CSP-enabled clusters.
P6	DATA PROTECTION	• Use Lakeflow Declarative Pipelines (formerly Delta Live Tables) expectations to drop, quarantine, or flag invalid records during ingestion. Delta Lake constraints (NOT NULL, CHECK) catch schema violations at write time. • Lakehouse Monitoring continuously profiles tables for freshness, volume, and quality anomalies. Unity Catalog lineage traces defective data to upstream sources.
P7	DATA PROTECTION	• Use Delta Lake's ACID transaction log and time travel to detect unauthorized changes. Route audit logs to your SIEM with detection rules for anomalous write events. • Register models in the Unity Catalog Model Registry (MLflow) for version-controlled promotion workflows (via model aliases) with audit trails per model version. Enable CSP for FIPS 140-validated encryption and CIS-hardened OS images on compute.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P8	PREVENTION OF UNAUTHORIZED USE	▪ Integrate with your enterprise IdP via SSO; enforce MFA, conditional access, and device posture at the IdP. Manage identities via SCIM. ▪ Use OAuth token federation so users and service principals exchange IdP tokens for short-lived Databricks OAuth tokens. Enable single-use refresh tokens for U2M flows. Review and revoke PATs via the account-level token report. ▪ Run regulated workloads on ESM/CSP compute to correlate identity events with host-level telemetry.
P9	PREVENTION OF UNAUTHORIZED USE	▪ Configure MFA, risk-based conditional access, and prohibit shared credentials via your IdP. Use SCIM to deprovision identities promptly on role changes or departures. ▪ Replace PATs with OAuth token federation; enable single-use refresh tokens to prevent replay attacks. Regularly review the token report to revoke unused or non-expiring tokens. ▪ Unity Catalog limits blast radius via least-privilege access even when credentials are compromised.
P10	PREVENTION OF UNAUTHORIZED USE	▪ Enable account-level audit log delivery and the audit log system table to capture sign-ins, token use, cluster/job operations, SQL queries, and permission changes. Forward to cloud storage and your SIEM. ▪ Use the token report to audit PAT inventory across workspaces. Unity Catalog provides detailed data access records and lineage for full traceability.
P11	PREVENTION OF UNAUTHORIZED USE	▪ Use Unity Catalog privileges to limit operations per role: grant SELECT for read-only access; restrict INSERT, UPDATE, DELETE, and OWNERSHIP to controlled groups. Apply row filters and column masking to limit visible data per user. ▪ Use compute policies to restrict compute creation by role and SQL warehouse permissions to control warehouse access.
P12	PREVENTION OF UNAUTHORIZED USE	▪ Cancel active job runs and stop clusters via the Jobs and Clusters APIs during incidents. Use Unity Catalog to revoke INSERT/UPDATE/DELETE/MODIFY privileges on affected tables. ▪ Use audit logs to identify which principals performed actions at incident time. Use Delta Lake time travel to restore data to a pre-incident state.
P13	PREVENTION OF UNAUTHORIZED USE	▪ Databricks delegates key storage to cloud KMS (AWS KMS, Azure Key Vault, GCP Cloud KMS). Configure CMK with least-privilege IAM roles so only designated workspaces can perform wrap/unwrap operations. ▪ Unity Catalog external locations govern access to CMK-protected storage without distributing keys to users. Use Lakehouse Federation to connect to external sources without sharing credentials.
P14	PREVENTION OF UNAUTHORIZED ACCESS FROM EXTERNAL NETWORKS	▪ Enforce SSO with MFA and conditional access via your IdP. Manage identities centrally via SCIM and prefer OAuth token federation over long-lived PATs. ▪ Monitor logins, token usage, and admin changes via audit logs. Revoke unused credentials regularly. Run regulated workloads on CSP/ESM-enabled compute.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P14-1	PREVENTION OF UNAUTHORIZED ACCESS FROM EXTERNAL NETWORKS	▪ Enhanced Security Monitoring (ESM) adds host-level monitoring agents — antivirus, file-integrity monitoring, and vulnerability scanning — plus advanced threat detection on hardened classic compute for regulated workloads. ▪ Stream audit logs and system tables to your SIEM and apply correlation rules to detect early indicators of compromise — privilege escalation, anomalous geographies, unusual service-principal activity, and abnormal data egress.
P14-2	PREVENTION OF UNAUTHORIZED ACCESS FROM EXTERNAL NETWORKS	▪ The financial institution is responsible for vulnerability assessment and penetration testing of its own in-scope environment. Customers may run their own vulnerability scans and penetration tests of their workloads in line with the Databricks penetration testing policy, and use CSP-hardened images to reduce the attack surface. ▪ As supporting assurance, Databricks performs regular independent penetration tests and vulnerability assessments of the platform itself; results are summarized in SOC 2 and available under NDA via the Trust Center.
P15	PREVENTION OF UNAUTHORIZED ACCESS FROM EXTERNAL NETWORKS	▪ Restrict access to approved networks via PrivateLink, IP access lists, and context-based ingress. Use customer-managed VPCs and compute policies to avoid public IPs on compute. ▪ Replace long-lived PATs with OAuth token federation to reduce credential exposure. Bind sensitive Unity Catalog catalogs to designated workspaces. Enable ESM/CSP for host-level security on exposed workloads.
P16	MEASURES TO DETECT UNAUTHORIZED ACCESS	▪ Enforce PrivateLink, compute policies, and context-based network policies to prevent direct internet egress. Use network policies (Serverless Egress Control) for a deny-by-default egress posture across serverless compute; some controls are in preview. ▪ Route audit logs and system tables to your SIEM for anomaly detection. Unity Catalog constrains which data objects are accessible from externally-connected workspaces. Enable ESM/CSP for host-level telemetry.
P17	MEASURES TO DETECT UNAUTHORIZED ACCESS	▪ Use Databricks SQL and Lakehouse Monitoring to detect anomalous transaction patterns (e.g., isolation-forest models on streaming or batch data). Configure SIEM alerts from audit logs. ▪ Unity Catalog lineage attributes unusual transactions to upstream jobs or users. Enable ESM/CSP for additional host-level behavioral telemetry.
P18	MEASURES TO DETECT UNAUTHORIZED ACCESS	▪ Build exception-monitoring pipelines (SQL queries or streaming jobs filtering threshold breaches) and expose results via dashboards and alerts. Use audit logs and system tables with your SIEM to detect bulk exports or mass updates. ▪ Unity Catalog fine-grained access control and lineage help focus monitoring on high-sensitivity catalogs. Enable ESM/CSP for host-level telemetry.
P19	RESPONSE MEASURES FOR UNAUTHORIZED ACCESS	▪ On detection, revoke tokens, disable accounts, rotate credentials, and reconfigure workspace or network settings. Sync users and groups via SCIM. ▪ Use Delta Lake time travel to restore data to a pre-incident state. Unity Catalog lets you revoke all privileges for compromised users and rebuild tables from trusted sources.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P20	MEASURES AGAINST MALICIOUS PROGRAMS	▪ Follow secure-coding guidelines in notebooks and jobs. Unity Catalog limits blast radius by restricting what data a compromised workload can access or write. ▪ Enable ESM/CSP, which adds hardened OS images, antivirus/malware detection, file-integrity monitoring, and host-level security telemetry forwarded to your SIEM.
P21	MEASURES AGAINST MALICIOUS PROGRAMS	▪ Use compute policies to restrict allowed compute images, libraries, and init scripts; limit package installation to approved sources. Unity Catalog enforces least-privilege access to tables, volumes, and external locations. ▪ Enable ESM/CSP for antivirus/malware detection, behavioral monitoring, file-integrity monitoring, and vulnerability reporting on compute nodes.
P22	MEASURES AGAINST MALICIOUS PROGRAMS	▪ Restrict runtimes and libraries via compute policies; minimize outbound internet access. Apply Unity Catalog least-privilege permissions to limit data exposure. Enable ESM/CSP for antivirus/malware detection and file-integrity monitoring. ▪ Recover corrupted data using Delta Lake time travel and versioned tables, complemented by backups in your cloud account.
P23	DOCUMENTATION	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P24	DOCUMENTATION	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P25	MANAGEMENT OF ACCESS RIGHTS	▪ Assign account roles, workspace entitlements, and Unity Catalog privileges via IdP-synced SCIM groups. Use OAuth token federation so users and service principals exchange IdP tokens for short-lived Databricks OAuth tokens. ▪ Enable Databricks-native MFA only when SSO is not configured. Apply CSP on workspaces handling regulated data. Use Lakehouse Federation for governed access to external data systems.
P26	MANAGEMENT OF ACCESS RIGHTS	▪ Centralize authentication via IdP SSO and SCIM; avoid local Databricks credentials. Use OAuth token federation for service principals to remove the need for personal access token (PAT) management. Enable single-use refresh tokens and configurable session lifetimes for U2M flows. ▪ Store credentials in Databricks Secrets rather than in notebooks or job configurations. Review and revoke unused PATs via the token report.
P27	MANAGEMENT OF ACCESS RIGHTS	▪ Implement joiner/mover/leaver workflows by syncing your IdP to Databricks via SCIM, assigning roles and entitlements per approved access requests. Use audit logs and system tables to evidence authentication and privilege changes for periodic access reviews. ▪ Centralize data authorization via Unity Catalog grants; include Lakehouse Federation foreign catalog privileges in the same review process. Enable CSP on workspaces handling regulated data.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P28	MANAGEMENT OF ACCESS RIGHTS	• Use Unity Catalog to organize and govern data as catalogs, schemas, tables, views, and volumes, managing permissions via grants. Use SQL warehouses as the governed query interface. • Register external JDBC databases as Lakehouse Federation foreign catalogs and manage access through Unity Catalog grants. Use Delta Sharing to govern external data sharing without copying underlying data.
P29	DATA MANAGEMENT	• Use Delta Lake's versioned tables, ACID transaction log, and time travel as the technical basis for data revision control. Use Git Folders and Declarative Automation Bundles (DABs) for controlled ETL promotion across environments. • Unity Catalog centralizes ownership and lineage for data-level revision. Register models in the Unity Catalog Model Registry (MLflow) for version-controlled promotion workflows with full audit trails.
P30	DATA MANAGEMENT	• Configure CMK with cloud KMS for workspace and storage encryption. Use audit logs and cloud KMS logs together to verify keys are used per operational procedures. • Point Unity Catalog external locations at CMK-protected storage so sensitive catalogs use encrypted storage. Enable CSP for stricter baselines on cryptographic operations.
P31	EDUCATION AND TRAINING	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P32	MEASURES AGAINST COMPUTER VIRUSES	• Enforce compute policies to constrain runtime versions, images, and init scripts; restrict package installation to curated private repositories. Unity Catalog limits blast radius via fine-grained permissions on catalogs, schemas, tables, volumes, and external locations. • Enable ESM/CSP for antivirus/malware detection, file-integrity monitoring, and vulnerability scanning on compute nodes, with security events forwarded to your SIEM.
P33	EXTERNAL CONNECTION MANAGEMENT	• Implement external connection conditions through customer-managed VPCs, PrivateLink or private endpoints, and IP access lists so only approved networks reach Databricks. • Bind sensitive Unity Catalog catalogs to specific workspaces to align data boundaries with network rules. Use audit logs and system tables to demonstrate that external access follows agreed-upon conditions. Mandate ESM/CSP clusters for externally-connected workloads.
P34	EXTERNAL CONNECTION MANAGEMENT	• Monitor approved external connections by reviewing audit logs and system tables for anomalous access patterns. Periodically review IP access lists and PrivateLink endpoint configurations. • Validate via the token report that only approved principals hold credentials with external-access scope.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P35	MANAGEMENT OF OPERATIONS	▪ Tie high-privilege roles (account admin, workspace admin, compute policy admin) to specific IdP groups so only qualified staff can be added. Require all operator credentials to be issued via OAuth token federation, gating service principal access on IdP group membership. ▪ Use the token report in periodic operator-qualification reviews alongside IdP group membership audits. Unity Catalog ownership and privileges distinguish data stewards from general users.
P36	MANAGEMENT OF OPERATIONS	▪ Govern Databricks role and group assignments through ITSM/IdP change-management workflows, validated via audit logs and system tables. Include the token report as a mandatory step in both provisioning (verify new operator token appears with expected scope/expiry) and deprovisioning (verify tokens are revoked). ▪ Unity Catalog supports procedure-driven operations via SQL scripts or automated pipelines for grants and ownership changes.
P37	MANAGEMENT OF OPERATIONS	▪ The operations organization itself is defined by the financial institution; Databricks supports it by structuring operations teams using account admin, workspace admin, and metastore admin roles mapped to SCIM groups. Unity Catalog formalizes data governance responsibilities via ownership and privilege control at catalog, schema, and table levels.
P38	MANAGEMENT OF OPERATIONS	▪ Audit logs and system tables record supported, auditable operations such as cluster/job management, SQL queries, permission changes, and workspace configuration (coverage varies by service and configuration). Forward to cloud storage or SIEM for long-term retention and alerting. ▪ Schedule regular token report exports as a supplementary credential-layer operational record. Unity Catalog provides detailed records of grants, revocations, and data-access operations.
P39	DATA FILE MANAGEMENT	▪ Databricks supports data-file backup primarily through Delta Deep Clone (full or incremental copies to a separate, independently protected location). Delta Lake versioning, time travel, and RESTORE give point-in-time recovery within a table's retention window but operate on the same underlying storage, so they are not an independent backup on their own. ▪ Underlying data in cloud storage (S3/ADLS/GCS) can be replicated cross-region using cloud-native storage replication configured by the customer. The backup schedule, retention, and recovery strategy are defined and owned by the customer to meet their RPO/RTO.
P40	PROGRAM FILE MANAGEMENT	▪ Manage notebooks, jobs, and libraries using Git Folders with workspace folder permissions. Use compute policies to define approved library sources and versions, and branch-based workflows with CI/CD for dev/test/prod promotion.
P41	PROGRAM FILE MANAGEMENT	▪ Manage program files (notebooks, jobs, libraries) under version control in Git via Git Folders, backed by your Git provider. Version control supports change management and recovery but is not, by itself, an independently retained backup. ▪ For retained backups, export workspace objects (notebooks, job and pipeline definitions) on a schedule via the Databricks APIs or Declarative Automation Bundles (DABs) to versioned cloud storage, and back up your Git repositories per your retention policy.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P42	NETWORK CONFIGURATION MANAGEMENT	▪ Manage VPCs/VNets, subnets, security groups, and PrivateLink endpoints in your cloud account; track configuration in your CMDB or IaC tooling. Unity Catalog catalog bindings help map data domains to specific network segments.
P43	NETWORK CONFIGURATION MANAGEMENT	▪ Use IaC (Terraform/CloudFormation/ARM) and version control as the primary backup for cloud network configuration. Export workspace networking settings (IP access lists) via APIs or IaC templates. Maintain catalog-to-workspace bindings under version control as part of your network-to-data mapping.
P44	DOCUMENT MANAGEMENT DURING SYSTEM OPERATION	▪ Store operational runbooks, diagrams, and configuration baselines in Git or Confluence. Host technical runbooks as notebooks in Databricks Git Folders. Use Unity Catalog to attach documentation to data assets for discoverability alongside the datasets they support.
P45	DOCUMENT MANAGEMENT DURING SYSTEM OPERATION	▪ Back up operational documents in redundant Git repositories or document management systems. Sync Databricks notebooks and queries to Git via Git Folders. Export Unity Catalog metadata via APIs/SQL for backup and DR planning.
P46	MONITORING OF OPERATIONS	▪ Query system tables (e.g., system.billing.usage and system.access.audit) plus compute, job-run, and data-quality logs via Databricks SQL to build operational dashboards and alerts. Export insights to observability or SIEM platforms. ▪ Enable CSP/ESM on critical workspaces to ensure dashboards reflect a secure and compliant infrastructure baseline.
P47	RESOURCE MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P48	RESOURCE MANAGEMENT	▪ Databricks is a PaaS deployed entirely on AWS, Azure, and GCP. Physical hardware and device maintenance are governed by those CSPs under ISO 27001, SOC 1/2, and PCI-DSS: Databricks reviews CSP audit reports on a regular cadence. ▪ At the software layer, Databricks maintains an ISMS, documented change management, a secure SDLC (SAST/DAST, vulnerability scanning), and customer commitments under the MCSA and Security Addendum (SOC 1/2 Type II, ISO 27001/27017/27018, HIPAA, PCI-DSS). Details: Databricks Trust Center.
P49	DEVICE MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P50	DEVICE MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P51	DEVICE MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P52	DEVICE MANAGEMENT	N/A — there is no Databricks platform feature for this physical-facility requirement; Databricks does not own or operate the underlying physical data centers, which are provided by the cloud providers (AWS, Azure, GCP).

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P53	DEVICE MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P54	DEVICE MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P55	PHYSICAL ACCESS CONTROL	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P56	PHYSICAL ACCESS CONTROL	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P57	PHYSICAL ACCESS CONTROL	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P58	PHYSICAL ACCESS CONTROL	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P59	FACILITY MONITORING	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P60	FACILITY MONITORING	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P61	TRANSACTION MANAGEMENT	▪ Define operational authority using RBAC (account/workspace admin roles, Unity Catalog privileges, group-based entitlements). Treat admin credentials as operator-level access governed by formal approval, stored in Databricks Secrets, and periodically recertified via the token report. ▪ Enable audit logging and system tables; forward to SIEM with alerts for high-risk actions (bulk exports, privilege escalations). Validate data input via Unity Catalog privileges, Delta Lake constraints, and Lakehouse Monitoring.
P62	TRANSACTION MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P63	TRANSACTION MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P64	TRANSACTION MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P65	INPUT/OUTPUT MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P66	INPUT/OUTPUT MANAGEMENT	• Databricks does not provide full DLP for off-platform exports (email, print); customers must enforce downstream controls for exported or printed data. Databricks compute policies restrict allowed images, libraries, and init scripts, while workspace export controls (notebook exports, result downloads, dashboard creation) limit the introduction and movement of files. • Unity Catalog enforces least-privilege access with fine-grained row/column filtering, masking, and centralized audit of read operations. Use Databricks Data Classification to detect and tag sensitive columns, feeding Attribute-Based Access Control (ABAC) policies for classification-based row filters, column masks, and workspace bindings. • For high-risk or regulated workloads, run compute on ESM/CSP-enabled clusters (hardened images, host-based antivirus/malware detection, file-integrity monitoring, vulnerability scanning). Stream audit logs and system tables to your SIEM, and complement with network egress controls such as PrivateLink and IP access lists.
P67	FORMS MANAGEMENT	• Management of physical printed forms is outside the Databricks platform and must be governed by your organization's records management policy. For digital equivalents (report templates, notebook-generated reports), manage in Git Folders and restrict access via workspace permissions. • Unity Catalog controls access to underlying datasets referenced by such forms.
P68	FORMS MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P69	CUSTOMER DATA PROTECTION	• Enforce SSO/MFA, least-privilege groups, and workspace security baselines (PrivateLink, compute policies, Databricks Secrets, controlled export paths). Enable encryption at rest and in transit; route audit logs to a SIEM. • Unity Catalog provides centralized access control, row filters/column masking, data classification, lineage, and system tables for audit and usage tracking. Run regulated workloads on ESM/CSP- enabled clusters.
P70	MEASURES IN CASE OF FAILURE AND DISASTER	• Databricks does not prescribe incident communication runbooks, but customers can use platform signals as inputs: subscribe to the Databricks status page, configure job/cluster/SQL warehouse alerts, and integrate audit logs and system tables into your SIEM. • Unity Catalog maps data products to their owners and stewards, facilitating rapid identification of responsible parties during incidents. For platform incidents, use Databricks Support and the Trust Center.
P71	MEASURES IN CASE OF FAILURE AND DISASTER	• Deploy primary and secondary Databricks workspaces in different AWS regions with cloud-native storage replication meeting your RPO/RTO. Back up critical notebooks, workflows, and configuration artifacts using Git Folders and DABs. • Unity Catalog supports DR via per-region metastores, enabling reconstruction of catalogs, access controls, and data locations during failover. Enable CSP/ESM on both primary and DR workspaces for a consistent hardened baseline.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P72	MEASURES IN CASE OF FAILURE AND DISASTER	▪ Use cluster, job, and SQL warehouse logs, system tables, and audit logs to diagnose failures (configuration errors, resource exhaustion, code bugs). Forward events to observability tools for cross-system correlation. ▪ Unity Catalog lineage identifies upstream causes and blast radius in data pipeline failures. Run DR investigation workloads on CSP/ESM-enabled workspaces.
P73	DEVELOPMENT OF CONTINGENCY PLANS	▪ Define a contingency plan covering DR triggers, roles, communication paths, and failback steps. Leverage secondary workspaces in different regions, Delta Lake backup/restore procedures, and Jobs/workflow failover runbooks. ▪ Use audit logs and system tables to validate the plan and refine RPO/RTO through DR tests. Unity Catalog per-region metastores and Delta Sharing enable structured data access continuity. Run DR procedures on CSP/ESM-enabled workspaces.
P73-1	DEVELOPMENT OF CONTINGENCY PLANS	▪ The financial institution is responsible for developing incident-response and contingency plans that assume cyberattacks for its own in-scope environment. Build these on-platform using cross-region workspace/DR failover, Delta Lake time travel and RESTORE for ransomware or tampering recovery, and SIEM-driven detection (audit logs and system tables) feeding your incident-response runbooks. ▪ As supporting context, Databricks operates its own documented security incident-response program and notifies affected customers in accordance with the Security Addendum.
P74	BACKUP CENTERS	▪ Deploy separate Databricks workspaces in multiple AWS regions as logical primary and disaster recovery (DR) centers. Configure cross-region replication of the underlying cloud storage (e.g., S3 CRR) plus Delta Deep Clone or periodic backups of tables and workspace configuration (via IaC and Databricks APIs). ▪ Unity Catalog requires a separate metastore per region, each acting as an independent governance center. Enable CSP/ESM on both primary and backup workspaces for consistent hardened baselines.
P75	SYSTEM DEVELOPMENT AND MODIFICATION MANAGEMENT	▪ Define SDLC using Git Folders (branching, pull requests, code review) and Declarative Automation Bundles (DABs) to promote notebooks, workflows, compute policies, and SQL objects across dev/test/prod workspaces. ▪ Enforce separation of duties via workspace permissions (restrict production deployments to service principals) and compute policies. Unity Catalog supports governed SDLC via separate catalogs per environment and workspace-catalog bindings.
P76	SYSTEM DEVELOPMENT AND MODIFICATION MANAGEMENT	▪ Establish separate dev, test, and staging workspaces from production. Use Git Folders plus DABs to promote code and configurations only after automated tests and approvals. Unity Catalog provides environment isolation via separate catalogs and workspace-catalog bindings. ▪ Manage production workspaces via IaC with audit logs and system tables providing deployment evidence. Enable CSP/ESM on all SDLC workspaces that process regulated data.
P77	SYSTEM DEVELOPMENT AND MODIFICATION MANAGEMENT	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P78	DOCUMENT MANAGEMENT DURING DEVELOPMENT AND MODIFICATION	▪ Manage development documentation via Git Folders, which version notebooks, SQL files, and pipeline definitions in an external Git provider. Use DABs for infrastructure-as-code artifact versioning and lifecycle management.
P79	DOCUMENT MANAGEMENT DURING DEVELOPMENT AND MODIFICATION	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P80	PACKAGE INSTALLATION	▪ Restrict library installation to an approved set using compute policies that auto-install allowlisted libraries and block ad-hoc user installs. Connect to private package repositories via Databricks Secrets and workspace configuration. ▪ Unity Catalog workspace-catalog bindings isolate evaluation clusters to non-sensitive catalogs. Enable ESM/CSP on evaluation and production clusters for runtime antivirus and file-integrity monitoring.
P81	PACKAGE INSTALLATION	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P82	DISPOSAL OF SYSTEMS	▪ Execute disposal plans using admin UI, APIs, and CLI to delete workspaces, clusters, jobs, and workspace objects. Use Delta Lake lifecycle controls (VACUUM, retention configuration) and cloud-storage lifecycle policies to physically remove retired data. ▪ Unity Catalog provides governed disposal via catalog/schema/table drops and privilege revocation through SQL, CLI, or Terraform: preventing post-disposal access via orphaned objects or stale permissions. Enable ESM/CSP on final disposal/migration workloads.
P83	DISPOSAL OF SYSTEMS	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P84	BACKUP FOR HARDWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P85	BACKUP FOR HARDWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P86	BACKUP FOR HARDWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P87	BACKUP FOR HARDWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P88	BACKUP FOR HARDWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P89	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	▪ Implement a secure SDLC using Git Folders, DABs, and Unity Catalog governance across all environments. In design, define architecture as code with PR-based review and least-privilege policies built in. In implementation, use protected branches with required reviews and SAST/DAST in CI. ▪ Deploy to non-production workspaces via DABs with automated tests (unit, integration, data- quality) before production promotion. Standardize on Databricks Runtime LTS releases. Restrict library sources via compute policies. ▪ Enable ESM/CSP on all environments handling regulated workloads for consistent antivirus, file- integrity, and vulnerability scanning across the full SDLC.
P90	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P91	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P92	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P93	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P94	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P95	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P96	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P97	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	▪ Use workspace file ACLs (VIEW/RUN/EDIT/MANAGE), cloud-storage IAM, PrivateLink, and IP access lists to restrict workspace and storage access. Unity Catalog enforces exclusive data access via fine-grained privileges, row/column-level policies, and Delta Lake ACID transaction log with schema enforcement and lineage. ▪ Enable ESM/CSP for hardened, auditable compute on backup, matching, and reconciliation workloads.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P98	MEASURES TO IMPROVE THE QUALITY OF SOFTWARE	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P99	MEASURES TO IMPROVE OPERATIONAL RELIABILITY	▪ The Databricks control plane spans multiple availability zones with automatic zonal failover. Clusters and SQL warehouses autoscale — with idle auto-stop, and scale-to-zero on SQL warehouses and serverless — and Lakeflow Jobs handles scheduling, retries with exponential backoff, and self-recovery. ▪ System tables and Lakehouse Monitoring provide data quality, lineage, and SLA tracking (retention varies by system table; system.access.audit retains up to 365 days). Enable ESM/CSP for regulated workloads.
P100	MEASURES TO IMPROVE OPERATIONAL RELIABILITY	N/A — there is no Databricks platform feature that customers can leverage to address this requirement. The financial institution meets it through its own controls and procedures.
P101	MEASURES TO IMPROVE OPERATIONAL RELIABILITY	▪ Monitor load with compute metrics, the compute and query system tables, and Databricks SQL dashboards; set alerts on utilization and queue depth. ▪ Control load through autoscaling and idle auto-termination (with auto-stop and, on SQL warehouses and serverless, scale-to-zero), SQL warehouse scaling, compute policies that cap cluster size and cost, and Lakeflow Jobs concurrency and retry controls. Enable CSP/ESM on regulated workloads.
P102	EARLY FAILURE DETECTION AND RECOVERY FUNCTION	▪ Lakeflow Jobs and Lakeflow Declarative Pipelines (formerly Delta Live Tables) emit structured run status, error codes, and event logs per task to trigger alerts. Lakehouse Monitoring catches freshness, volume, and quality anomalies before bad data propagates downstream. ▪ System tables and audit logs unify job runs, cluster events, SQL activity, and configuration changes for SIEM root-cause isolation. Unity Catalog lineage identifies affected downstream assets. Delta Lake time travel enables rollback to a known-good state.
P103	EARLY FAILURE DETECTION AND RECOVERY FUNCTION	▪ Configure job-level alerts, event-log monitoring, and dashboards to detect failures in jobs, clusters, and SQL workloads. Drill into job, cluster, and system table logs to isolate the failing component. ▪ With Unity Catalog enabled, use object-level audit logs, access history, and data lineage to identify which user, service principal, or upstream job caused a pipeline failure. Enable ESM/CSP for additional host-level telemetry.
P103-1	EARLY FAILURE DETECTION AND RECOVERY FUNCTION	▪ Validate that redundant and backup structures actually work. The Databricks control plane runs across multiple availability zones with automatic failover; compute recovery behavior (auto-restart, job retries, SQL warehouse restart) varies by compute type, configuration, failure domain, and cloud — so test it rather than assume it. ▪ Exercise DR regularly using secondary workspaces in another region, Delta Deep Clone and time-travel restore tests, and Lakeflow Jobs failover runbooks; confirm results via system tables and audit logs. Run DR tests on CSP/ESM-enabled workspaces.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P104	EARLY FAILURE DETECTION AND RECOVERY FUNCTION	▪ Reduce or shut down operations during failure by pausing/disabling job schedules, stopping clusters, and applying workspace-level restrictions via the Jobs and Clusters APIs. Redirect workloads to alternate workspaces or regions as needed. ▪ Unity Catalog enables rearrangement by reconfiguring access policies to restrict write access and switching consumers to read-only or backup data locations. Enable ESM/CSP on shutdown and recovery workloads.
P105	EARLY FAILURE DETECTION AND RECOVERY FUNCTION	▪ Implement read-only mode during incidents by pausing job runs and stopping clusters via the Jobs and Clusters APIs. Use Unity Catalog to revoke INSERT/UPDATE/DELETE privileges on critical tables. ▪ Enable ESM/CSP on failure-management workloads for hardened images and security telemetry.
P106	EARLY FAILURE DETECTION AND RECOVERY FUNCTION	▪ Use Delta Lake time travel, restore from cross-region backup storage, and rerun ETL/ELT pipelines to restore affected tables and jobs. Unity Catalog centralizes metadata and privileges for reconstructing the logical environment (catalogs, schemas, grants) during recovery. ▪ Enable ESM/CSP on recovery workloads for consistent hardened baselines and forensic telemetry.
P107 – P111	CARD TRANSACTION SERVICE	N/A — there is no Databricks platform feature for card-transaction-service requirements; these consumer-facing financial channels are outside the platform.
P112 – P118	INTERNET AND MOBILE SERVICES	N/A — there is no Databricks platform feature for Internet and mobile banking channel requirements; these are outside the platform.
P119 – P124	CD/ATM AND UNMANNED TERMINALS	N/A — there is no Databricks platform feature for CD/ATM and unmanned-terminal requirements; these are outside the platform.
P125 – P126	NONBANK / IN-STORE ATM LOCATIONS	N/A — there is no Databricks platform feature for nonbank/in-store ATM location and device requirements; these are outside the platform. (P127–P131 are Missing numbers in the 13th Edition.)
P132 – P137	DEBIT CARD AND PREPAID SERVICES	N/A — there is no Databricks platform feature for debit-card, prepaid-payment, and electronic-value requirements; these are outside the platform.
P138 – P139	USE OF E-MAIL AND INTRANET	N/A — there is no Databricks platform feature for e-mail and intranet usage requirements; Databricks is not an e-mail or web-browsing service.
P140 – P141	BIOMETRIC AUTHENTICATION	N/A — there is no Databricks platform feature for biometric-authentication requirements; Databricks integrates with the customer's identity provider via SSO but does not itself provide biometric authentication.
P142 – P144	QR CODE PAYMENT	N/A — there is no Databricks platform feature for QR-code payment requirements; these are outside the platform.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
P145 – P148	REMOTE WORKING	N/A — there is no Databricks platform feature for these remote-working device and channel requirements; they are managed on customer-controlled endpoints outside the platform.
P149	MOBILE-BRANCH VEHICLE SERVICES	N/A — there is no Databricks platform feature for mobile-branch-vehicle requirements.
P150	USE OF AI	▪ Govern AI as first-class assets in Unity Catalog: models, functions, features, and vector indexes carry the same access controls, ownership, and lineage as data — the technical foundation for AI-use policy and accountability. ▪ Use the MLflow Model Registry for governed, version-controlled model promotion and the Unity AI Gateway (formerly Mosaic AI Gateway) to centrally define and enforce AI usage policies (permitted models, rate limits, logging). Appoint asset owners in Unity Catalog to fulfill accountability; policy formulation itself remains a customer responsibility.
P151	USE OF AI	▪ Operate AI transparently: MLflow captures model versions, parameters, and lineage from training data to endpoint; Unity AI Gateway logs inference requests and responses; and inference tables retain payloads for review. ▪ Monitor quality and drift with Lakehouse Monitoring, including fairness and bias monitoring, and evaluate models and agents with MLflow 3 GenAI evaluation (mlflow.genai, formerly Mosaic AI Agent Evaluation), so unfair, incorrect, or biased AI behavior can be detected and remediated.
P152	USE OF AI	▪ Apply Unity AI Gateway service policies (guardrails) — PII detection and masking, safety and topic filters, and prompt-injection mitigations — on model-serving endpoints. ▪ Secure the AI stack with Unity Catalog access control over training data and models, private networking for model serving, and CSP/ESM for regulated AI workloads to prevent unauthorized operation, data leakage, and unintended behavior.
P153	USE OF AI	▪ User education on responsible AI use is a customer responsibility. Databricks supports it with product documentation and responsible-AI guidance, and with Unity AI Gateway service policies (guardrails) that can surface usage constraints and block unsafe prompts.

MAPPING FISC SECURITY GUIDELINES · DATABRICKS' CUSTOMER CAPABILITIES

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
F	**Facilities Guidelines**	
F1 – F83	DATA CENTER FACILITIES	N/A — there is no Databricks platform feature for these physical data-center facility requirements; Databricks does not own or operate the underlying physical data centers, which are provided by the cloud providers (AWS, Azure, GCP).
F84 – F138	HEAD & BRANCH OFFICE FACILITIES	N/A — there is no Databricks platform feature for these physical head/branch-office and ATM facility requirements.

01 GUIDELINE	02 CATEGORY	03 DATABRICKS CUSTOMER CAPABILITIES
A	**Audit Guidelines**	
A1	SYSTEM AUDITING	▪ Databricks provides a unified audit and observability framework that supports regulated and financial workloads by capturing control plane, data plane, and AI/BI activity across accounts and workspaces in system tables such as system.access.audit, lineage, compute, Lakeflow jobs, billing/usage, Apps, AI/BI Genie, and Vector Search. The system.access.audit table retains up to 365 days of authentication, authorization, configuration, and data-access activity, while customers may extend retention by exporting these logs to encrypted cloud storage and a SIEM with their own multi-year lifecycle, immutability, and integrity controls. ▪ Audit events from Databricks system tables and diagnostic logs can be configured to continuously ingest logs to your centralized SIEM, where correlation rules and alerting are implemented for high-risk actions such as privilege escalation (e.g., changes to high-privilege groups or Unity Catalog grants on sensitive schemas), creation or modification of model serving endpoints that touch regulated data, changes to workspace and account-level security configurations, and anomalous access patterns from unusual geographies or service principals. ▪ For regulated and financial workloads, the Compliance Security Profile (CSP) can be enabled — and is required for certain compliance configurations — with Enhanced Security Monitoring (ESM) on eligible classic compute, so those workloads run on hardened, monitored infrastructure with FIPS-validated cryptography, host-level telemetry, and advanced threat detection. Serverless and any other planes not yet covered by ESM are still in scope for auditing via system tables and log export, with documented compensating controls and monitoring in the SIEM to ensure consistent visibility across all Databricks services.
A1-1	CYBERSECURITY AUDITING	▪ Conducting the internal cybersecurity audit — risk-based audit planning, independence, reporting material issues to the CEO/board, and tracking remediation — is the financial institution's responsibility. Databricks supplies the evidence auditors examine: tamper-evident audit logs and system tables (e.g., system.access.audit) exportable to immutable cloud storage and a SIEM, Unity Catalog grant and lineage history, and the token report for privileged-access review. ▪ Map the audit's targets to Databricks evidence — maintenance/operation status (system tables, compute and job logs), response and recovery (disaster recovery configuration and SIEM-based incident detection), and third-party/platform risk (independent SOC 2, ISO 27001/27017/27018, and PCI-DSS attestations plus the Compliance Security Profile, via the Trust Center). See A1 for the underlying audit and observability framework.

databricks

END OF MAPPING

Mapping FISC Security Guidelines to Databricks' Customer Capabilities

REFERENCE EDITION

FISC Security Guidelines on Computer Systems for Financial Institutions, Thirteenth Edition (March 2025)

DOCUMENT VERSION

Version 1.0

SECURITY & COMPLIANCE

Databricks Security & Trust Center

PLATFORM DOCUMENTATION

docs.databricks.com