Mixpeek Information Security Policy

Version: 1.0 **Effective Date:** April 15, 2026 **Last Reviewed:** April 15, 2026
Owner: Mixpeek Engineering

1. Purpose

This policy establishes the information security requirements for Mixpeek, Inc. to protect the confidentiality, integrity, and availability of all information assets, including customer data, infrastructure, and intellectual property.

2. Scope

This policy applies to all Mixpeek employees, contractors, and third-party service providers who access Mixpeek systems, data, or infrastructure.

3. Data Classification

All data handled by Mixpeek is classified into the following categories:

Classification	Description	Examples	Controls
PUBLIC	Information intended for public disclosure	Marketing content, documentation, blog posts	No special controls
INTERNAL	Business information not intended for public release	Internal communications, meeting notes, architecture diagrams	Access restricted to employees
CONFIDENTIAL	Sensitive business or customer data	Customer API keys, source code, financial records	Encryption at rest and in transit, access logging

Classification	Description	Examples	Controls
PHI	Protected Health Information subject to HIPAA	Health records, medical data processed through Mixpeek platform	All CONFIDENTIAL controls plus HIPAA safeguards, BAA required

Namespace-level classification is enforced through the `data_classification` field on each namespace resource. PHI-classified namespaces trigger additional access logging and encryption controls.

4. Access Control

4.1 Authentication

- All user access to Mixpeek systems requires authentication via PropelAuth (Studio) or Clerk (Canvas).
- API access requires a valid API key issued through the Mixpeek API key management system.
- API keys are hashed using SHA-256 before storage; plaintext is displayed only once at creation time.
- Multi-factor authentication (MFA) is enforced for all organization members.

4.2 Authorization

Mixpeek implements a role-based access control (RBAC) model with four permission levels:

1. **READ** — View resources and data
2. **WRITE** — Create and modify resources (implies READ)
3. **DELETE** — Remove resources (implies WRITE, READ)
4. **ADMIN** — Full access including user and key management (implies all)

API keys support fine-grained scoping by namespace, resource type, and operation, including wildcard patterns.

4.3 Principle of Least Privilege

- Access is granted based on the minimum permissions required for the user's role.
- API keys should be scoped to the specific namespaces and operations needed.
- Administrative API keys are protected and cannot be edited or deleted by non-admin users.

4.4 Access Review

- User access is reviewed quarterly.
- API keys have configurable expiration dates.
- Unused API keys are flagged for review after 90 days of inactivity (tracked via `last_used_at`).

5. Encryption

5.1 Encryption in Transit

- All external communications use TLS 1.2 or higher.
- Internal service-to-service communication within the Kubernetes cluster uses mTLS where supported.
- MongoDB connections enforce `requireTLS` mode.
- HSTS headers are configured with `max-age=31536000; includeSubDomains; preload`.

5.2 Encryption at Rest

- Organization secrets are encrypted using Fernet symmetric encryption before storage.
- Storage connection credentials are protected using MongoDB Client-Side Field Level Encryption (CSFLE).
- Cloud storage (AWS S3) uses server-side encryption (SSE-S3).
- Encryption keys are managed through environment-level secrets with planned migration to GCP Cloud KMS.

6. Network Security

- Kubernetes network policies enforce zero-trust architecture between services.
- Ingress is restricted to authorized endpoints with Google-managed TLS certificates.
- Access to the GCE metadata server (169.254.169.254) is explicitly blocked to prevent credential theft.
- Prometheus metrics endpoints are restricted to internal network ranges (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16).
- Public API endpoints enforce rate limiting by operation category (metadata, data, search, upload, compute).

7. Audit Logging

- All significant actions are logged to an immutable ClickHouse-backed audit trail.
- Audit events capture: actor ID, action, resource type, resource ID, changes, status, and timestamp.
- Audit logs are retained for 365 days.
- API request metadata is logged with sampling for performance analysis.

8. Vulnerability Management

- Python code is scanned for security vulnerabilities using Bandit (integrated into pre-commit hooks).
- Dependency vulnerabilities are monitored via GitHub Dependabot for Python, JavaScript, and GitHub Actions.
- Container images are scanned before deployment.
- Pre-commit hooks enforce architecture boundary checking to prevent unauthorized cross-module imports.

9. Incident Response

Refer to the Incident Response Plan (separate document) for detailed procedures. Key points:

- Security incidents must be reported immediately to the security owner.
- Incidents involving PHI must follow HIPAA Breach Notification procedures.
- Post-incident reviews are conducted within 72 hours.

10. Vendor Management

- All third-party vendors that process customer data must have a signed contract or terms of service.
- Vendors processing PHI must have a signed Business Associate Agreement (BAA).
- Vendor security posture is evaluated before onboarding and reviewed annually.
- A current vendor inventory is maintained in the compliance directory.

11. Physical Security

Mixpeek operates entirely on cloud infrastructure (Google Cloud Platform, AWS). Physical security is delegated to cloud providers:

- **GCP**: SOC 2 Type II, ISO 27001, HIPAA BAA available
- **AWS**: SOC 2 Type II, ISO 27001, HIPAA BAA available

12. Employee Security

- All employees receive security awareness training upon hire and annually thereafter.
- Access is provisioned on hire and revoked within 24 hours of departure.
- Employees must use company-approved devices with full-disk encryption.

13. Business Continuity

Refer to the Business Continuity Plan for detailed procedures. Key points:

- Cross-region database backups (MongoDB: us-east1 primary, us-west1 DR).
- Backup retention: 30 days.
- Quarterly disaster recovery drills with documented results.
- PodDisruptionBudgets ensure service availability during maintenance.

14. Policy Review

This policy is reviewed and updated at least annually, or when significant changes occur to Mixpeek's infrastructure, operations, or regulatory requirements.

This document is maintained by Mixpeek Engineering and is subject to change. Questions should be directed to info@mixpeek.com.