Jump

SEPTEMBER 2026

Security at Jump.

How we protect your data and keep your clients' information secure — purpose-built for financial advisors.

- SOC 2 Type II
- GDPR · CCPA
- EU AI Act

DOCUMENT
Security Program Whitepaper

CLASSIFICATION
Confidential

TRUST CENTER
security.jump.ai

DESIGNED FOR FINANCIAL SERVICES

The AI meeting platform built for financial professionals.

Jump is the AI meeting platform built specifically for financial professionals. We help teams capture, summarize, and act on every client conversation — while meeting the rigorous security and compliance expectations of the financial services industry.

This document describes the security program, controls, and practices Jump maintains to protect customer data. Our independent audit and penetration test reports, and all other available documentation, are published in our Trust Center at security.jump.ai.

CONTENTS

01	**Overview** Program governance, certifications, and insurance	03
02	**Data handling** Encryption, segregation, retention, redaction, ownership	03
03	**AI governance** Model providers, training opt-out, human-in-the-loop, AI security testing	04
04	**Access control** SSO, MFA, role-based access, recertification	05
05	**Cloud & infrastructure security** Network and platform monitoring	05
06	**Application security** Secure development, SAST/DAST, pen testing, bug bounty	05
07	**Incident response** Tabletop testing and SIEM-based monitoring	06
08	**Business continuity & disaster recovery** Redundancy, backups, corporate office	06
09	**Personnel security** Hiring, training, and workstation hardening	06
10	**Risk management** Risk assessments and third-party risk management	07

© 2026 JUMP

01

SECTION ONE

Overview

Jump's security program is managed by in-house, dedicated personnel, led by the Head of Security. The Head of Security reports directly to executive leadership and the board regularly, establishing a tone from the top. Security policies and procedures are documented and reviewed by management annually.

Jump maintains SOC 2 Type II compliance, with an audit period spanning a 12-month observation window year over year, with no gap in between. In addition to SOC 2 compliance, Jump aligns with applicable privacy and AI law and regulations, such as GDPR, CCPA, and the EU AI Act.

Cybersecurity insurance is maintained and renewed annually. Our certificate of insurance is available in our Trust Center for review.

02

SECTION TWO

Data handling

Data encryption

Customer data is encrypted using current industry-standard algorithms and methods — AES-256 at rest and TLS 1.2+ in transit. Jump relies on GCP for key management, such that no Jump employee has access to the underlying encryption keys.

Data segregation

Jump provides a multi-tenant, cloud-based SaaS solution. Customer data is logically segregated to prevent commingling, with isolation enforced at every function level.

Storage & retention

Jump supports US, Canada, and EU hosting regions. Customer data, including backups, is stored in the region designated.

By default, customer data is retained for the duration of the agreement. However, customers may define granular retention periods, including zero data retentino, within the application per meeting asset (e.g., audio/video recordings, meeting transcriptions, and AI-generated summaries).

Data redaction

Jump natively supports redaction of Non-Public Information (NPI) and Personally Identifiable Information (PII) from transcripts and any downstream content. Customers are strongly encouraged to enable this feature.

DATA OWNERSHIP

Customer data is owned by the customer. Jump does not train on or monetize customer data, or use it for advertising. Customer meeting content — recordings, transcriptions, summaries — is never used for analytics.

© 2026 JUMP

03

SECTION THREE

AI governance

AI model providers

Jump uses enterprise API deployments of third-party Large Language Models (LLMs) to power some features of the application (e.g., AI-generated meeting notes). Jump identifies our third-party model providers in our list of sub-processors in our Trust Center.

Training & opt out

Jump does not use any customer data to train, fine-tune, or improve any AI model. Jump has executed Data Processing Agreements (DPAs) with all AI model providers, and enforces zero-data-retention and 'no training'. All customers are opted out of model training by default.

HUMAN-IN-THE-LOOP

Jump operates a human-in-the-loop model. Jump may draft or propose actions (e.g., CRM updates and emails), but autonomous action is never taken on behalf of the user. Jump does not send or write to any external system until an end user has explicitly reviewed and approved the action. Automated actions that write to external systems are always gated by human approval.

AI performance & security testing

Jump monitors AI model performance using structured evaluation sets to benchmark precision, recall, and F1 scores across model versions. These metrics are tracked to detect performance drift, regression, or safety degradation before and after production deployments. Before any model is promoted to production, it is assessed for disparate performance across relevant demographic, use-case, and input-type dimensions.

Jump conducts adversarial security testing specific to AI systems, including:

- Automated red-teaming to test for direct and indirect prompt injection across key AI workflows.
- Annual third-party AI penetration testing focused on prompt injection and data exfiltration risks.

AI SAFEGUARD	HOW IT WORKS
No customer-specific RAG	Jump does not maintain a customer-specific knowledge base or fine-tune models on customer data. All outputs are generated from workflow context provided at runtime.
Model versioning & rollback	Models are version-pinned. Jump can roll back to a prior stable version if a newly deployed model exhibits unexpected behavior or performance degradation.
Bias review	Bias mitigation is a required checkpoint in evaluation and deployment. Findings are documented, reviewed by AI governance, and factored into model selection and prompt design.

© 2026 JUMP

04

SECTION FOUR

Access control

Jump enforces the use of a strong password and MFA for all internal and production systems access, where technically feasible. Access reviews are performed regularly to recertify production systems access, with privileged access reviewed under additional scrutiny.

Upon employee termination, Jump automatically deactivates the user account in Jump's IdP, propagating account deactivation down to integrated applications.

05

SECTION FIVE

Cloud & infrastructure security

Network security

Jump applies a zero-trust approach, treating all connections as hostile even within the cloud environment. All access not explicitly permitted is denied by default.

Network connections between application, data, and workstation segments are controlled by firewall rules, and firewall rule sets are reviewed at least annually. Firewall rule changes are managed via code, requiring an independent review and approval.

Platform monitoring

Jump implements an IDS/IPS, in addition to leveraging a Cloud Security Posture Management (CSPM) solution to perform continuous cloud and infrastructure scanning for vulnerabilities, misconfigurations, and threats. Alerts are sent to the Security team for review and triage.

Infrastructure performance and system health is monitored as well, with autoscaling enabled to meet capacity and resource needs.

06

SECTION SIX

Application security

Secure development

In addition to requiring independent peer review and approval on production application code changes, all changes must complete QA testing, linting, and SAST prior to merge. DAST is also performed.

Security testing

Jump engages an independent third party to perform annual penetration testing, scoped to the Jump application and cloud infrastructure. Jump also conducts AI-specific security assessments focused on prompt injection and data exfiltration risks.

Jump hosts a bug bounty program with an ongoing responsible disclosure channel for the security research community.

© 2026 JUMP

07

SECTION SEVEN

Incident response

Jump maintains a documented Incident Response Policy and Plan that is tested at least annually. IR tabletop exercises involve stakeholder involvement and are scoped to scenarios relevant to Jump's threat and risk landscape.

Application and system event logs are piped to a SIEM. Logs are immutable, such that modifications cannot be made. Alerts are configured and sent to on-call personnel for review and triage.

08

SECTION EIGHT

Business continuity & disaster recovery

Redundancy

Jump is hosted across multiple GCP availability zones, with immediate, automated failover in the event of zone failure. Autoscaling is enabled to meet infrastructure and capacity needs.

Backups

Backups are performed daily, retained for 14 days, and encrypted at rest. Jump maintains a formal BCP/DR plan that is reviewed annually and tested through annual tabletop exercises and database restoration exercises.

Corporate office

Jump does not host or store production data at its corporate office. The corporate network is physically and logically segregated from the production environment and access to the network does not provide any elevated access to the production environment. The majority of Jump's workforce is remote, and all personnel are able to work remotely, with no impact or interruption to business operations.

09

SECTION NINE

Personnel security

We ensure that our employees are qualified and aware of their security responsibilities by requiring:

- **Background checks.** Pre-employment screening is conducted for all employees, as legally permissible.
- **Confidentiality agreements.** Employees must sign a confidentiality agreement prior to their first day.
- **Security policies.** Employees must review and sign their acknowledgement of Jump's security policies upon hire and annually thereafter. Policy violations are subject to disciplinary measures, up to and including termination.
- **Security awareness training.** Employees must complete security awareness training within 30 days of hire and at least annually thereafter. Developers complete additional training scoped to secure coding practices.

Employee workstations

Jump does not permit BYOD. Employees use Jump-provisioned laptops that are centrally managed with an MDM. Employee laptops are installed with anti-malware, an EDR, and an enterprise DLP solution; and are enabled with host-based firewall, automatic OS updates, removable-media restrictions, password protection, and inactivity timeout.

© 2026 JUMP

10

SECTION TEN

Risk management

Risk assessment

A documented Risk Management Policy governs risk identification and documentation, risk rating, and treatment. Jump conducts formal risk assessments at least annually, with identified risks documented in the risk register for tracking until close.

Risks are regularly reporting to senior and executive management, with material risks escalated to the board as appropriate.

Third-party risk management

Jump maintains a formal vendor management program that includes conducting due diligence and a risk assessment prior to onboarding, and performing ongoing monitoring for all third parties.

Additional scrutiny is provided to sub-processors and third parties providing an AI service. Vendor agreements include appropriate confidentiality, security, and privacy obligations, where applicable.

© 2026 JUMP

FOR MORE INFORMATION

Self-serve access to every report we publish.

Our SOC 2 Type II report, penetration test report, and all other security documentation are available from Jump's Trust Center. Customers may also subscribe to receive email notifications for changes, such as when documentation is updated or when new sub-processors are added.

TRUST CENTER

security.jump.ai

SECURITY CONTACT

security@jumpapp.com

LEGAL RESOURCES

jump.ai/msa | jump.ai/terms-of-service

© 2026 JUMP