AI Assistant
Security Technical Paper

Published: April 2026 | Distribution: Public

Executive Summary

AI Assistant is a conversational interface integrated into Jamf Pro, Jamf Account, and Jamf Protect — powered by Claude (Anthropic) via AWS Bedrock. It provides production tools for inventory queries, configuration analysis, compliance checking, and knowledge retrieval.

This document describes the security architecture, data handling practices, and privacy controls that govern AI Assistant's operation across all Jamf Cloud regions (US, EU, APAC).

AI Assistant is built on four security principles, each enforced at the architectural level rather than solely at the policy or prompt layer: disabled by default, least-privilege access control, read-only enforcement at the API layer, and transparent/attributable responses. AWS Bedrock Guardrails provide additional content monitoring and prompt injection detection across all environments.

Architecture Overview

Infrastructure

AI Assistant is deployed across all Jamf Cloud regions. Customer data is processed in the region where their Jamf environment is hosted and is not transferred across regional boundaries.

Region	AWS Bedrock Region	Status
United States	us-east-1	Production
European Union	eu-central-1	Production
Asia-Pacific	ap-northeast-1	Production

Model

AI Assistant uses Claude (Anthropic) accessed through AWS Bedrock. Bedrock provides the inference layer between Jamf's infrastructure and Anthropic's model. For the current model version, see the Jamf Learning Hub.

Anthropic sub-processor relationship: Anthropic does not receive or have access to customer data. Model inference occurs within AWS Bedrock infrastructure, which operates within Jamf's AWS environment. Customer queries, tool results, and conversation context are processed within Bedrock and are not transmitted to Anthropic. For details on how AWS Bedrock handles data privacy and security, see the AWS Bedrock Data Protection documentation.

AWS Bedrock security properties:

- Customer data is not used to train or fine-tune Anthropic's models
- Data is processed in-region and does not leave the AWS region where the customer's environment is hosted
- SOC 2 Type II compliant
- AWS enterprise security controls apply to all inference requests

Model updates: Jamf manages model versioning through AWS Bedrock. The current model version is maintained in the Jamf Learning Hub — organizations with change management requirements should monitor the Learning Hub for model version changes.

Tool Architecture

AI Assistant uses a tool-calling architecture: when a user submits a query, the model determines which tools to invoke, executes them against specific Jamf APIs using the user's existing permissions, and synthesizes results into a response.

All tools are read-only. AI Assistant tools span five categories: knowledge retrieval (Jamf documentation and knowledge base), configuration access (policies, profiles, scripts, blueprints, etc.), inventory queries (device data — Mac and mobile), compliance checking (benchmarks against CIS, NIST, and DoD STIG, etc.), and security intelligence (mobile app risk assessments). Jamf Protect tools (alert analysis, malware lookup) are available in limited beta. For the current tool catalog including availability and product requirements, see the Jamf Learning Hub.

Third-party data flows: Three tools query external services outside Jamf infrastructure. These integrations are documented for transparency.

- **Apple OS Lookup** queries Apple's Global Device Management Framework (GDMF) API (gdmf.apple.com), a public Apple endpoint. No customer data is transmitted — the tool retrieves publicly available Apple OS release information only.
- **App Lookup** queries the iTunes Search API (itunes.apple.com) as a fallback data source for application version and patch information. No customer data is transmitted — the tool retrieves publicly available app metadata only.
- **Mobile App Risk** queries NowSecure's MARI (Mobile Application Risk Intelligence) database to retrieve security assessments. The only data transmitted is the application's store identifier (e.g. iOS bundle ID) and platform (iOS or Android). No device data, user identities, or organizational information is transmitted.

Security Design Principles

Disabled by default. AI Assistant is disabled for all organizations until explicitly enabled by an administrator in Jamf Account. Individual tool groups require separate activation — enabling AI Assistant Core does not automatically enable Jamf Pro tools or any future product integrations. No AI capability is available to users until their organization has made an affirmative choice to enable it, and administrators can disable any tool group at any time.

Least-privilege access control. All tool queries execute under the permissions of the authenticated user, inheriting existing Jamf Pro RBAC controls without modification. AI Assistant does not elevate privileges or access data the user cannot already access directly. A user without permission to view a policy cannot retrieve it through AI Assistant.

Read-only enforcement at the API layer. AI Assistant calls Jamf Pro APIs using the authenticated user's session token — there is no separate service account with elevated credentials. All API calls issued by AI Assistant tools are GET requests. No tool in the system issues a POST, PUT, PATCH, or DELETE request against Jamf Pro. This is an architectural constraint enforced at the implementation layer — not a prompt-layer instruction or a policy that could be overridden through clever prompting. Regardless of how a query is structured, AI Assistant cannot modify device configurations, deploy policies, remove applications, or change enrollment states.

Transparent, attributable responses. Every response surfaces its sources, allowing administrators to verify answers against authoritative documentation. Tool results returned to the model are structured data rather than free-form text, making every response traceable to its source.

Bedrock Guardrails. AWS Bedrock Guardrails are deployed across all AI Assistant environments. The guardrail configuration includes content monitoring for multiple harm categories (violence, sexual content, hate speech, insults, misconduct) and prompt injection detection at high sensitivity. All guardrail events are traced and logged, providing a full audit record of flagged inputs and outputs.

Data Handling

Data Flow

When a user submits a query, the following sequence occurs:

1. **Query processing:** The user's natural language query is received by the AI Assistant backend
2. **Tool execution:** Relevant tools query Jamf APIs using the authenticated user's permissions
3. **Context assembly:** The user's query, relevant tool results, and the current conversation thread are prepared for inference
4. **Model inference:** The inference request is processed by AWS Bedrock and a response is generated
5. **Response delivery:** The generated response is returned to the user in the Jamf interface

What Data is Processed During Inference

Data Type	Processed by Inference Layer	Notes
User query	Yes	The natural language question as submitted
Tool results	Yes	Inventory data, configuration details relevant to the query
Conversation history	Yes	Full thread history for the current conversation, loaded from persistent storage; retained for 30 days
User credentials or tokens	No	Never included in model context
Full database contents	No	Only query-relevant results are included

Data Residency

AI Assistant respects Jamf's regional data boundaries. Inference requests are routed to the AWS Bedrock deployment in the same region as the customer's Jamf environment:

- **US customers:** Data processed in AWS us-east-1
- **EU customers:** Data processed in AWS eu-central-1
- **APAC customers:** Data processed in AWS ap-northeast-1

Device inventory, configuration data, and other customer-specific data are not transferred between regions.

Note on knowledge retrieval: Knowledge retrieval queries Jamf's documentation corpus only — it does not access device inventory, configuration details, or other customer-specific data. All AI Assistant queries, including knowledge retrieval, are processed in the customer's assigned region.

Session Isolation

Each AI Assistant conversation is scoped to the authenticated user and their organization. Conversation context is not shared between users or between customer organizations. A query from one organization cannot surface inventory data or configuration details from another organization.

Conversations are retained for 30 days, then automatically and permanently deleted. Retention is enforced at the storage layer via DynamoDB TTL — not a scheduled cleanup task that could be deferred or skipped. Each conversation is accessible only to the user and organization that created it. Conversation data is stored exclusively in Jamf's infrastructure — queries and responses are not logged or retained by AWS Bedrock or Anthropic beyond the inference request itself.

Retention and Audit Logging

Conversation content is retained for 30 days. After 30 days, conversation data is deleted and cannot be recovered.

Audit logs are maintained in Jamf Account under Activity History → AI Assistant. The audit log captures all administrative changes to AI Assistant configuration, including:

- AI Assistant enabled or disabled
- Tool groups added, removed, or updated
- The identity (name and email) of the administrator who made each change
- The date and time of each change

Audit log entries are accessible to Organization Administrator and Administrator roles in Jamf Account. The audit log provides a complete record of configuration changes.

Data Type	Retention Period	Notes
Conversation content	30 days	Automatically deleted; cannot be recovered
Audit log (configuration changes)	Jamf Account standard retention	Accessible in Jamf Account Activity History
Model inference context	Not retained beyond session	Discarded when session ends
Model training	Not applicable	Anthropic does not train on AWS Bedrock customer data

Access Control

Authentication

AI Assistant inherits the authenticated user's Jamf session. No separate login, API key, or credential is required. Users who are not authenticated to their Jamf environment cannot access AI Assistant.

Enabling AI Assistant requires the Administrator or Organization Administrator role in Jamf Account. Standard users and read-only roles cannot enable, disable, or modify AI Assistant tool group settings. All changes made by administrators are recorded in the Activity History audit log.

Authorization

All tool queries execute under the permissions of the authenticated user. AI Assistant does not elevate privileges or bypass existing Jamf Pro role-based access controls:

- Inventory queries return only devices the user has permission to view
- Configuration explain results respect existing object-level access controls
- Compliance data access follows standard Jamf Pro RBAC
- A user without access to a policy cannot retrieve that policy's details through AI Assistant

Enabling and Disabling AI Assistant

AI Assistant is disabled by default for all organizations. Administrators enable it explicitly in Jamf Account under Organization → AI Assistant.

Disabling AI Assistant is immediate and reversible. An administrator unchecks the Enable AI Assistant checkbox in Jamf Account. This disables all AI Assistant capabilities for all users across the organization instantly. Individual tool groups (Jamf Pro read-only tools) can also be disabled independently without disabling AI Assistant Core.

Environment-level scoping provides an additional control layer for organizations that want a more cautious rollout. When enabling Jamf Pro read-only tools, administrators can restrict access to specific environments and tenants rather than enabling across all environments. This allows organizations to pilot AI Assistant in a sandbox or staging environment before enabling it in production, without any change to the production environment.

Tool Availability by Product

Current tool availability, product requirements, and beta status are maintained in the Jamf Learning Hub.

Compliance

AI Assistant operates within Jamf's existing compliance program. For the current list of Jamf certifications, see the Jamf Trust Center or contact your account team.

AWS Bedrock compliance (applicable to inference layer):

- SOC 2 Type II
- ISO 27001

FedRAMP and StateRAMP: AI Assistant is not available in StateRAMP or FedRAMP-authorized environments. Contact your Jamf account team for roadmap details on future FedRAMP and StateRAMP availability.

Penetration testing: AI Assistant has undergone penetration testing as part of Jamf's security review program. Results are available to customers under NDA upon request through your Jamf account team.

Security Controls Summary

Control	Implementation
Encryption in transit	TLS 1.2+ for all communications
Encryption at rest	AWS KMS encryption
Authentication	Inherits Jamf Pro session — no separate credential required
Administrator roles required	Administrator or Organization Administrator in Jamf Account
Authorization	Jamf Pro RBAC enforced on all tool queries — no privilege escalation
Data residency	Regional processing — US/EU/APAC, no cross-region transfer
Anthropic data access	Anthropic does not receive customer data — inference stays within AWS Bedrock
Model training	Customer data not used for model training (AWS Bedrock)
Third-party processors	NowSecure (app risk intelligence) — app identifiers only; Apple GDMF (OS versions) — public data only
Audit logging	Configuration changes logged in Jamf Account Activity History — by user, action, and timestamp
Conversation retention	30 days
Read-only operation	Enforced at implementation layer — all Jamf Pro API calls are GET requests; no write methods exist in the tool code

Session isolation	Conversations scoped to authenticated user and organization — not accessible to other users or organizations
Disabled by default	Disabled for all organizations until explicitly enabled by an administrator
Disable capability	Immediate — uncheck Enable AI Assistant in Jamf Account; reversible at any time
Environment scoping	Pro tools can be scoped to specific environments/tenants for controlled rollout
Web Application Firewall (WAF)	Applied at API Gateway layer in production and stage environments
Bedrock Guardrails	Content monitoring across harm categories and prompt injection detection at high sensitivity; all events traced and logged
Penetration testing	Conducted prior to GA; results available under NDA

Document Information

Published	April 2026
Distribution	Public
Location	jamf.it/aiassistant

For More Information

- **Jamf Trust Center** — Current Jamf certifications, compliance documentation, and security posture: https://www.jamf.com/trust-center/
- **Jamf Learning Hub** — Current AI Assistant tool catalog, product requirements, and model version: https://learn.jamf.com/home
- **This document** — The most current version of this paper is available at: jamf.it/aiassistant
- **Questions?** Contact your Jamf account team.