AI Assistant in Jamf Pro - Architecture & Security

Today's IT administrators face an increasingly complex landscape with limited resources at their disposal. Jamf's AI Assistant for Jamf Pro addresses these challenges by bringing together artificial intelligence capabilities with the established Jamf ecosystem, transforming these obstacles into opportunities for enhanced operational excellence.

This document explores the comprehensive security architecture that forms the foundation of this offering. Our design of AI Assistant maintains data integrity within established organizational boundaries while providing an intuitive interface for managing administrative complexity.

Current Capabilities

AI Assistant for Jamf Pro currently operates exclusively in a **read-only** capacity. This fundamental design choice strengthens security while providing valuable functionality through three core skills:

- **Reference skill**: Retrieves and consolidates information from Jamf documentation and support history to answer your questions with verifiable sources.
- **Explain skill**: Generates detailed explanations of specific Jamf Pro objects (policies, scripts, smart groups, extension attributes and more to come) by accessing their current configuration data through the Jamf Pro API.
- **Search skill**: Translates natural language questions into database queries, allowing administrators to quickly gather insights about their managed devices.

This read-only implementation means the AI Assistant can help you understand your environment, retrieve information, and generate insights—but it cannot modify configurations, create or delete objects, or make changes to your Jamf Pro instance.

Future Capabilities

While the current implementation is entirely read-only, Jamf plans to introduce write-based skills in the future. These write-based capabilities will be available on an opt-in basis and will be accompanied by additional technical documentation. These write-based skills will incorporate Jamf's same commitment to security and organizational control that underpins AI Assistant for Jamf Pro.

Security Driven Architecture

AI Assistant's Jamf Pro skills operate within a carefully orchestrated ecosystem where organizational boundaries are fundamental to every operation. This architectural approach delivers key advantages:

- **Organization-Centric Authentication**: The system maintains continuous awareness of organizational context, ensuring data never crosses organizational boundaries through persistent context validation and isolation.
- **AWS-Native Security Integration**: By operating within the same AWS environment as most Jamf Pro tenant data, we eliminate unnecessary data transmission risks and leverage AWS's robust security features. (AI Assistant is available in Jamf Pro Azure deployments and data is transmitted securely over AWS PrivateLink, though responses may entail relatively higher latency compared to AWS hosted instances)

Disabled by Default

In line with industry best practices and drawing inspiration from proven security approaches, the AI Assistant for Jamf Pro is designed to be 'off by default.'

Nothing is enabled until an administrator explicitly opts in. Opt-ins are organized by capability **tier**, not by every individual "skill," so teams can align features with their security posture:

Capability Tier	What it Enables	Default State
Retrieval-Augmented Generation (RAG)	Assistant accesses only public Jamf knowledge sources to answer questions. No tenant data is touched.	Off
Read-Only Jamf Pro Tools	Assistant can call read-only APIs to pull configuration data, inventory records, etc., but can never alter Jamf Pro.	Off

Every activation or deactivation of AI functionality is carefully logged, providing a comprehensive audit trail for compliance and internal reviews. This logging includes:

- WHEN: The time and date of the change
- WHO: The administrator who made the modification
- WHAT: The specific features affected and the organizational context of the change

Data Residency and Processing

Our commitment to data security extends beyond just access controls. AI Assistant uses Anthropic LLMs served by Amazon Bedrock. By leveraging AWS Bedrock within the same secure environment as Jamf Pro tenant data, information never leaves your protected AWS ecosystem. All AI Assistant processing takes place inside Jamf's secured AWS environment, ensuring your data remains under Jamf's control end-to-end. **Your prompts and the outputs you receive are processed only in encrypted, transient memory, are not logged or retained by AWS Bedrock or Anthropic and are never used to train or fine-tune the underlying foundation models.**

Critically, our architecture ensures that any data processed for AI inference is never repurposed to train or enhance models. This strict data boundary extends to all aspects of the system:

- Secure, private AI model inference within the AWS environment
- No data is used for model training
- Complete audit trails track all AI operations

Note (Beta): During the AI Assistant beta program, all AI Assistant workloads are deployed in U.S. AWS regions. The service is accessible globally, but data will remain within the United States during the 30-day retention window. **By General Availability (GA), Jamf will introduce region-specific deployments so customer data can stay within designated geographic boundaries.**

Data Retention and Deletion

All customer-specific data handled by the AI Assistant—including conversation history, Explain Skill object payloads, Inventory Search result sets, generated diagrams, and system logs—is stored in organization-scoped, encrypted DynamoDB tables for **30 calendar days**. After 30 days the data is automatically and irreversibly purged via AWS Lifecycle policies. No data is ever used to retrain models or shared across organizations.

Data Usage by Skill

- **Explain skill** – To answer admin questions, the configuration details of each Jamf Pro object in question may be fetched, rendered into prose, and written to encrypted logs. These entries follow the same **30-day** retention window.
- **Search skill** – To answer inventory-related questions, the Assistant may load a working set of up to 5,000 device records into memory per query. Only the generated SQL and the resulting insights/responses are retained, and these are automatically purged after

30 days. Based on the query, fetched information may include inventory fields that contain personally identifiable information, including:
  - *device name, last IP address, last reported IP, user username, user realname, user email, user position, user phone, department id, building id, room, application names and paths, group membership, custom extension attributes*
- **Reference skill** – Uses only publicly available Jamf documentation; **no customer data are retained**.

All retention and purge operations are logged and auditable.

(See **“Disabled by Default”** for how organizations opt in to these features.)

Security Implementation: A Defense-in-Depth Strategy

Our comprehensive security strategy implements controls across multiple dimensions, ensuring protection at every layer of the system. This multi-faceted approach includes data protection, authentication lifecycle management, and organizational isolation.

Authentication Gateway

Tyk API Gateway validates incoming requests before proceeding to the authentication flow. By integrating Tyk, we ensure that:

- **Validated Requests:** API calls are checked for requisite security headers.
- **Proactive Threat Management:** Potentially malicious requests are identified and filtered out at the network edge.
- **Continuous Monitoring:** Real-time detection mechanisms are in place to swiftly counter any emerging threats.

Data Protection

For data in motion, we implement:

- TLS 1.2/1.3 encryption for all communications
- JWT token validation at service boundaries
- Robust parameter validation to prevent injection attacks

For data at rest, we ensure:

- DynamoDB encryption with org-specific access keys

- AWS KMS integration for cryptographic key management
- Organization-isolated database schema design

Authentication Lifecycle

Authentication in our system operates as a continuous process throughout the request lifecycle:

1. Initial authentication leverages existing Jamf identity infrastructure through Auth0
2. Token validation includes proper signature verification and claims extraction
3. Context propagation embeds organization identifiers in all internal requests
4. Each component independently validates authentication context before processing

Preserving Organizational Boundaries

Our multi-level isolation strategy ensures complete organizational separation through both storage layer and functional isolation. At the storage layer, DynamoDB implements sophisticated composite keys that incorporate organizational identifiers.

This key structure ensures proper data isolation while maintaining efficient access patterns for different operational needs. The functional layer builds upon this foundation by ensuring that each specialized component operates with explicit organizational context.

Future Horizons: Strategic Evolution

As AI Assistant continues to evolve, we're focusing on several strategic enhancements:

- **Expanded Function Registry**: We're building a growing ecosystem of specialized capabilities that maintain our security-first approach while extending functionality.
- **Advanced Tenant Controls**: New features will provide granular control over AI functionality, allowing organizations to customize their security posture on a per tenant basis.
- **Seamless Integration**: Additional Jamf product connections will expand the assistant's capability while maintaining our security standards.

Conclusion

Jamf's AI Assistant for Jamf Pro represents our vision for simplifying administrative complexity while maintaining the highest security standards. Our architecture ensures that the AI Assistant maintains data security and organizational isolation throughout all interactions while providing powerful AI capabilities to enhance the Jamf Pro experience.