REQASE

CONFIDENTIAL

Security & Compliance

Architecture, Data Handling & Operational Security

Document Version: 1.0

Last Updated: February 2026

Classification: External - Enterprise Distribution

Last Updated: February 2026

> This Security & Compliance document is designed for enterprise risk, compliance, and information security teams evaluating Reqase as a secure AI-assisted testing solution. It outlines our architecture, data handling model, compliance posture, and operational security controls.

1. Company Overview

Reqase is an AI-assisted testing application built on the Atlassian Forge platform. We enable QA and software teams to generate test cases and documentation using enterprise AI services.

Reqase:

- Does not operate independent hosting infrastructure
- Does not maintain standalone production servers
- Runs entirely within Atlassian Cloud
- Relies on Atlassian-managed infrastructure for storage and compute

All primary customer data remains within Atlassian Cloud.

2. Architecture & Data Flow

2.1 Hosting Model

Reqase is deployed exclusively within Atlassian Forge and operates without independent infrastructure:

- No self-hosted servers
- No independently managed databases
- No direct inbound internet endpoints
- No external infrastructure operated by Reqase

All compute and storage resources are provisioned and managed by Atlassian.

2.2 High-Level Data Flow

[1] User submits content within Jira or Confluence

[2] Forge invokes the Reqase function inside Atlassian infrastructure

[3] Task-relevant text is transmitted securely to an enterprise AI provider

[4] AI-generated output is returned to Atlassian

[5] Output may be temporarily cached within Forge-managed storage (see Section 3)

2.3 Infrastructure Responsibility Model

Layer	Responsible Party
Cloud Infrastructure	Atlassian
Data Storage	Atlassian
Identity & Access	Atlassian
Application Logic	Reqase
AI Processing	Enterprise AI Provider

3. Data Handling & Storage

3.1 Application-Layer Processing Model

Reqase operates as an application-layer service within Atlassian Cloud.

> **We do NOT:**
>
> - Operate independent storage systems
> - Maintain external backups
> - Export data outside Atlassian-managed infrastructure
> - Use customer data for analytics or profiling
> - Use data for AI model training

3.2 Limited Application-Level Caching

To improve user experience and reduce repeated AI requests, Reqase may temporarily cache certain AI-generated outputs within Atlassian Forge managed storage.

This caching:

- Occurs only within Atlassian-managed infrastructure
- Is not accessible outside the application context
- Is not used for analytics, profiling, or secondary purposes
- Is subject to automated retention limits
- Is periodically deleted

Reqase does not independently control the underlying database infrastructure and cannot directly access Atlassian's storage layer outside application-level permissions.

3.3 Data Location

All customer data, including any temporary cached data, remains within Atlassian cloud regions as configured by the customer's Atlassian tenant. Reqase does not independently determine or control data residency.

4. Third-Party AI Providers

Reqase integrates with enterprise-grade AI providers including:

- OpenAI
- Google (Gemini)
- Anthropic (Claude)

AI Data Handling

- Encrypted in transit (TLS)
- Limited to task-scoped input only
- No independent persistent storage by Reqase
- No AI model training under enterprise configurations

Note: If customers use their own AI subscription (e.g., Azure OpenAI), processing occurs under the customer's direct contractual relationship with the AI vendor.

5. Compliance & Certifications

5.1 Atlassian Compliance

Reqase inherits security controls from Atlassian's cloud environment, which maintains certifications including:

- SOC 2 Type II
- ISO 27001
- GDPR-aligned controls

As Reqase does not operate independent infrastructure, we do not maintain a separate SOC 2 report.

6. Security Controls

Although Reqase does not operate its own infrastructure, we maintain internal security practices including:

- Secure development lifecycle (SDLC)

- Code reviews
- Least-privilege access controls
- Role-based access to development systems
- Encrypted API communications
- Controlled access to application configuration

7. Data Protection Role

Under applicable data protection laws:

- Customer organization = Data Controller
- Atlassian = Primary Data Processor
- Reqase = Sub-processor operating within Atlassian infrastructure

Reqase does not independently determine:

- Data residency
- Core retention policies for Atlassian-stored data
- Infrastructure-level storage configuration

Application-level temporary caching is purpose-limited and automatically managed.

8. Incident Management

In the unlikely event of a security incident:

- We follow documented internal response procedures
- We coordinate with Atlassian where infrastructure-related
- Affected customers will be notified without undue delay where legally required

9. Vendor Risk Positioning

Compared to traditional SaaS providers, Reqase operates with a reduced operational risk profile:

- No independently hosted infrastructure
- No standalone production servers
- No external database systems
- No background behavioral analytics
- Limited and purpose-scoped temporary caching only

10. Changes to This Policy

We may revise this Security & Compliance document periodically to reflect evolving best practices or regulatory requirements. Updates will be posted on our official website and Marketplace listing. Continued use of the application after updates constitutes acceptance of the revised version.

11. Contact

If you have any security or compliance-related questions, please contact:

Email: contact@reqase.com