Security Overview

A procurement-ready summary of how Quantesic handles customer data, who can access it, and what your security team should expect during diligence.

> Quantesic supports post-sale teams working with ticket history, account context, sentiment, and renewals — data that has been handed to you in trust by your customers. Our framework is built around least-privilege access, workspace-scoped tenant separation, encrypted transport, controlled administrative access, and human-reviewed AI output. This document is the public-safe summary; a deeper Trust Pack is available on request for customer-specific review.

01 **Identity-first access**

Access is tied to authenticated users, workspace membership, and role-aware permissions.

02 **Tenant separation**

Customer workspaces are separated so account and support data stays scoped to the right organization.

03 **Controlled operations**

Sensitive actions are mediated by application controls instead of direct browser-side writes.

04 **Privacy-aware AI**

AI output is scoped to the customer's own workflows, visible to users, and designed for human review.

AT A GLANCE

Tenancy model	Workspace-scoped; cross-workspace access is not part of the standard user model.
Data in transit	Encrypted transport (TLS) across all customer-facing surfaces.
Data at rest	Managed cloud infrastructure; customer-facing mutation paths are limited to controlled application workflows.
Model training	Customer data is not used to train cross-tenant or foundation models.
Sub-processors	Disclosed and reviewed during procurement; provider routing reviewable for enterprise teams.
Customer review	Security questionnaires, architecture review, and DPA discussion handled directly with the buyer.
Contact	admin@quantesic.com

SECURITY FRAMEWORK

The control model customers should expect.

Identify who is acting, limit what they can access, protect customer data in transit and at rest, review sensitive AI output, and keep an auditable operational practice.

IDENTITY & ACCESS

Access starts with verified identity.

Users sign in through supported authentication flows and receive access based on workspace membership and role. Administrative functions are reserved for authorized customer administrators and Quantesic support operators.

AUTHORIZATION

Permissions are checked before protected actions.

Quantesic separates ordinary workflow access from sensitive operations — user administration, integration setup, AI provider configuration, and customer support activity.

TENANT SEPARATION

Customer workspaces stay logically separated.

Customer records are organized by workspace and protected by application-layer and data-layer controls that prevent users from accessing another customer's account context.

DATA PROTECTION

Customer data is protected across the service.

Traffic uses encrypted transport. Stored data is handled through managed infrastructure, with direct customer-facing data mutation paths limited to controlled application workflows.

OPERATIONAL ACCESS

Support access is limited and purposeful.

Quantesic administrative access is reserved for support, implementation, tenant management, and reliability work. Customer-visible support and trust activity is designed to make intervention easier to review.

CHANGE DISCIPLINE

Changes go through build and deployment checks.

Product changes go through source control and production build validation. Customer-impacting controls are documented and reviewed as the platform evolves.

PRIVACY & AI

Customer data should not become someone else's shortcut.

Quantesic uses customer context to support that customer's workflows. The product is not designed to turn one customer's private history into another customer's advantage.

No cross-tenant model training.

Customer data is not used to train cross-tenant models or foundation models. Tenant-specific feedback is intended to improve that tenant's own experience.

Human review before action.

AI output is inspectable and editable before it affects customers, escalations, renewal narratives, or account communication.

Source-aware context.

Synthesized context is presented with attention to the underlying customer systems and records that informed the work.

Provider choice for enterprise teams.

For customers with stricter AI governance needs, provider routing, usage visibility, and enterprise key-management options are reviewable during implementation.

TECHNICAL OVERVIEW

Architecture at a safe level of detail.

ACCESS	Every protected workflow requires an authenticated user session. The application evaluates workspace membership and role before returning customer data or allowing a sensitive operation.
DATA	Accounts, tickets, activity, customer notes, AI usage, and integration status are handled as workspace-scoped data.
INTEGRATIONS	Connected systems use provider authorization flows and expose readiness state — configured, connected, operationally active.
SECRETS	Provider keys, OAuth secrets, and enterprise AI keys are treated as sensitive configuration; raw secrets are not displayed back after setup.
AUDIT	Administrative changes, support actions, AI review flows, and customer-impacting events are designed to leave reviewable operational evidence.
RESILIENCE	Availability expectations, support contacts, escalation paths, and customer-specific requirements are agreed during rollout planning.

CUSTOMER REVIEW PACKAGE

What we cover during diligence.

Architecture review

Workspace boundaries, data flow, integration model, administrative access model, and AI processing posture.

Access & role review

Supported sign-in patterns, administrator responsibilities, user roles, support access, and offboarding expectations.

Data handling review

Customer data categories, retention, deletion and export handling, integration scopes, and privacy obligations.

AI governance review

Human review model, provider routing, tenant-level usage visibility, customer key options, and customer-specific AI restrictions.

Contractual review

Security exhibits, confidentiality commitments, data processing terms, sub-processor review, and incident notification terms.

Implementation review

Source systems, rollout scope, admin setup, least-privilege configuration, and mutually agreed support procedures.

PROCUREMENT FAQ

High-signal answers without exposing the blueprint.

Can Quantesic complete a security questionnaire?

Yes. We complete customer security questionnaires and provide supporting explanations appropriate to the review stage and confidentiality requirements.

How is customer data separated?

Data is scoped by workspace and governed by authenticated access, role-aware permissions, and application controls that prevent standard users from crossing workspace boundaries.

Does Quantesic train shared models on customer data?

No. Customer data is not used to train cross-tenant or foundation models.

Can customers request export or deletion?

Yes. Export, deletion, and retention expectations are handled through the customer agreement and implementation process.

Do you publish detailed infrastructure diagrams publicly?

No. Appropriate architecture detail is shared with customers during a controlled security review.

> NEXT STEP
>
> To request the detailed **Trust Pack** — sub-processor list, data flow diagram, retention table, and customer-specific control mapping — email admin@quantesic.com. For a live security walkthrough, book at quantesic.com/book.