DevAI Suite — Security White Paper

Version 1.0 — April 2026 Classification: Public · safe to share with prospects, customers, and auditors.

Executive Summary

DevAI Suite is the first Manufacturing-as-a-Service (MaaS) platform — a unified APQP programme management, supplier excellence, and smart factory intelligence environment for automotive and aerospace manufacturers. The platform handles **regulated manufacturing data** (PPAP submissions, FMEA workbooks, supplier qualifications, plant telemetry) for customers operating under **IATF 16949** and **AS9100 / AS9145**.

This white paper documents how DevAI Suite is designed and operated to protect that data. It is structured around the **NIST Cybersecurity Framework 2.0** (Govern, Identify, Protect, Detect, Respond, Recover) and maps each function to specific controls, technologies, and operational practices. Where certifications are pending, we say so honestly with a defensible audit timeline rather than overclaim.

The intended audience is enterprise security review teams, vendor risk-management functions, and procurement officers performing due diligence ahead of contract signature.

Posture	Status
Multi-tenant isolation	PostgreSQL Row-Level Security on every customer table
Encryption in transit	TLS 1.2+, HSTS with `includeSubDomains`, AEAD ciphers only
Encryption at rest	Cloud volume-level + envelope encryption for sensitive credentials
Authentication	PBKDF2-HMAC-SHA256 passwords (bcrypt verified for legacy accounts), Google + Microsoft Entra ID SSO, TOTP MFA
Audit logging	Append-only per-org log of security-relevant events
GDPR / CCPA	Compliant; DPA available; DSR pipeline live
SOC 2 Type II	In progress (observation H2 2026, audit Q4 2026)
ISO/IEC 27001:2022	Planned (certification audit H1 2027)
Disaster recovery	Documented runbook, 30-min RTO, 5-day point-in-time backups, quarterly rehearsals

1. Govern — Security Program & Accountability

1.1 Security ownership

A named security lead reports to the founder/CEO and owns:

- Security policy authoring and review (annual cadence; ad-hoc on incident).
- Vendor risk reviews and subprocessor onboarding decisions.

- Incident response coordination.
- Audit and certification programme management.
- Customer security questionnaire responses.

The platform-admin role is bound to a specific user account via the `is_platform_admin` flag (replacing the deprecated email-comparison check); privilege derives from the database row, not from email. MFA is mandatory for the platform-admin role.

1.2 Policies in force

Policy	Scope	Review cadence
Information Security Policy	Organisation-wide	Annual
Acceptable Use Policy	Personnel	Annual
Access Control Policy	Engineering + ops	Annual
Vendor Management Policy	Procurement	Annual
Incident Response Plan	Engineering + leadership	Annual + post-incident
Disaster Recovery Plan	Engineering	Quarterly rehearsal
Secure Development Lifecycle	Engineering	Annual
Data Retention & Deletion Policy	Organisation-wide	Annual
Cryptographic Key Management	Engineering	Annual
Change Management	Engineering	Annual

Policies are kept under version control and reviewed by the security lead before each release cycle. Customer-eligible excerpts are available under NDA.

1.3 Personnel security

- Background checks for personnel handling regulated customer data.
- Confidentiality agreements signed at engagement start.
- Onboarding includes security awareness training; annual refresh.
- Access provisioning follows the principle of least privilege; quarterly access review.
- Offboarding is automated where possible (SSO de-provisioning, secret rotation, repository revocation).

1.4 Compliance & certification roadmap

Standard	Status	Target
GDPR (EU 2016/679)	Compliant	Continuous
CCPA / CPRA	Compliant	Continuous
SOC 2 Type II	Observation period H2 2026; audit Q4 2026	Report by Q1 2027
ISO/IEC 27001:2022	ISMS scoping H2 2026	Certification H1 2027
ISO/IEC 27701	Bundled with 27001	H2 2027
IEC 62443 (industrial)	Roadmap	On enterprise demand

For the audit firms engaged, control narratives, and pre-audit evidence packages, contact security@devaisuite.com.

2. Identify — Asset Management & Risk

2.1 Asset inventory

The platform's asset surface is bounded and inventoried:

- **Application apps:** `devai-api` (FastAPI backend), `devai-web` (Next.js frontend), `devai-admin` (admin/billing), `devai-landing` (marketing site). All deployed to Fly.io (Frankfurt FRA).
- **Data stores:** Fly Postgres (`devai-db`) for all relational data including audit logs; Cloudflare R2 buckets `devai-suite-prod` (global templates, no customer data) and `devai-tenant-docs-prod` (per-org documents).
- **Cache / queues:** Redis (rate limit, ephemeral); Inngest (background-job orchestration).
- **External providers:** see §6 Vendor Management.

A complete subprocessor list with data-residency attribution lives at `/subprocessor-list.html` .

2.2 Data classification

Class	Examples	Handling
Public	Marketing pages, security white paper	No restriction
Internal	Source code, internal runbooks	Access-controlled, not customer-visible
Customer Confidential	APQP projects, FMEA, PPAP, supplier records, uploaded documents, AI prompts/ responses	RLS-isolated; encryption at rest; access-logged
Regulated PII	Customer email, name, address (where collected), auth credentials	Same as above + dedicated PII handling per §3.4
Secrets	API keys, encryption keys, OAuth client secrets	Fly Secrets store; envelope encryption for stored credentials; never in code or logs

2.3 Risk management

Material risks are tracked in an internal risk register reviewed quarterly. Top items as of 2026-Q2:

- **Single-region Postgres** — single point of failure for the database tier; mitigated by 5-day point-in-time backups + documented restore. Migration to multi-region or Fly Managed Postgres is on the roadmap as customer scale justifies.
- **LLM provider dependency** — AI Coordinator depends on third-party model APIs (Together AI as the default provider, with Anthropic and OpenAI on elevated/fallback lanes); fallback chain configured, no single-provider lock-in.
- **Embedding-dimension lock** — pgvector schema is fixed at Voyage AI `voyage-3-large` 1024 dimensions; provider or dimension

rotation requires a coordinated re-embed of the corpus documented in the DR runbook.

2.4 Threat model

The platform is threat-modelled against:

- **External attackers** — credential stuffing, account takeover, injection, SSRF, supply-chain compromise.
- **Malicious or compromised tenants** — attempted cross-tenant data access via API enumeration, RAG poisoning, LLM prompt injection.
- **Insider threat** — over-privileged engineers, departed personnel.
- **Third-party compromise** — subprocessor breach affecting our data.

Mitigations are described in the relevant sections below.

3. Protect — Controls

3.1 Multi-tenancy & data isolation

Every customer's data lives in a single shared PostgreSQL database, isolated by **PostgreSQL Row-Level Security**. Every table that holds customer data has an RLS policy keyed to the request's tenant context. The application layer sets the tenant on every database session through the request authentication middleware:

```
ALTER TABLE deliverables ENABLE ROW LEVEL SECURITY;
CREATE POLICY tenant_isolation ON deliverables
  USING (org_id = current_setting('app.current_org_id', true)::int);
```

A request that leaks tenant context (even via an internal application bug) cannot read or write rows belonging to another tenant — the database refuses. Cross-tenant SELECTs return zero rows by design.

The `app.db.tenant.set_tenant_context` helper is the single seam through which tenant context is applied; it is exercised on every request that resolves a user, before any business logic runs.

3.2 Per-tenant document storage

RAG documents live in Cloudflare R2 with object keys prefixed by `org_id`. Reads are scoped by the same tenant context, with a server-side check before any presigned-URL issuance.

Embeddings (Voyage AI `voyage-3-large`, 1024-dim) are stored in pgvector tables that carry an `org_id` column under the same RLS policy as the source documents.

3.3 Encryption

In transit. TLS 1.2+ on all customer-facing endpoints. Strong cipher suites only (AES-GCM, ChaCha20-Poly1305). HSTS is enforced with `max-age=31536000; includeSubDomains`. The `preload` directive is intentionally omitted until every subdomain meets the eligibility criteria. X-Frame-Options is `DENY` and X-Content-Type-Options is `nosniff` on every response.

At rest. Postgres data is encrypted at rest by Fly's underlying volume layer. Cloudflare R2 encrypts at rest by default (AES-256). Sensitive integration credentials stored in `integration_connectors.credentials_json` are envelope-encrypted with a per-deployment Fernet key (`DEVAI_CREDENTIALS_ENCRYPTION_KEY`); the key is rotated on a documented cadence and never leaves the secret store.

Key management. Application keys (JWT signing, Fernet, OAuth client secrets, third-party API keys) are stored in Fly Secrets, surfaced to the running process at boot, and never persisted to disk. Key rotation procedures are documented per provider in `docs/runbooks/disaster-recovery.md` §3.

3.4 Authentication & access control

Identity options:

- Email/password — passwords are hashed with PBKDF2-HMAC-SHA256 (passlib `CryptContext`), with a per-password random salt and iterated key stretching; bcrypt hashes from legacy/migrated accounts remain verifiable. Password history is not enforced; minimum length is 12 characters with complexity requirements (configurable per-org).
- Google SSO via OAuth 2.0.
- Microsoft Entra ID SSO via OAuth 2.0 / OIDC, defaulting to the `organizations` tenant (work/school accounts) — personal MSA accounts are blocked by the `is_business_domain` gate.
- SSO-only users have `NULL` passwords and the password-login route refuses them rather than offering a leaky fallback.

Multi-factor authentication. TOTP-based, optional per-org-policy, mandatory for the platform-admin role. The MFA enforcement check runs in the auth dependency on every protected request; users with MFA enabled but unverified can only reach the `/auth/mfa/verify` and `/auth/mfa/status` endpoints until they verify.

Role-based authorization. Every protected route calls `ensure_permission(db, user, resource_type, action, org_id)` which checks the user's role-permission grant against the requested action. Decisions are recorded to an append-only `policy_decisions` audit table.

Session tokens. Short-lived JWTs (default 12 hours) with rotation on sensitive actions (password change, MFA enrollment). Stored in `HttpOnly; Secure; SameSite=Lax` cookies AND mirror in `localStorage` for the SPA — both must agree on every request. Token reuse after invalidation is detected by checking against a blocklist updated on logout / revocation.

CAPTCHA. Cloudflare Turnstile gates registration, login, password-reset, MFA enrollment, and the public demo sandbox provisioning route. The

server-side `verify_turnstile_token` helper logs Cloudflare's `error-codes` field on failure for diagnosis.

Rate limiting. A Redis-backed middleware enforces per-endpoint per-IP limits — auth endpoints capped at 10 requests/minute, AI Coordinator feedback at 30/minute, default at 100/minute. Limits fail open when Redis is unreachable (availability over rigid limits) but log the failure for review.

3.5 Application security

- OWASP ASVS Level 2 alignment is the baseline.
- Input validation via Pydantic schemas at every API boundary; type strictness enforced at compile time with mypy.
- ORM-only data access (SQLAlchemy); no string-concatenated SQL.
- HTTP responses carry strict security headers (HSTS, X-Frame-Options=DENY, X-Content-Type-Options=nosniff, Referrer-Policy=strict-origin-when-cross-origin). CSP is on the roadmap and gated behind a thorough testing pass to avoid breaking the AI Coordinator embeds.
- SSRF defense: outbound HTTP from the application is bounded to a known allowlist of provider hostnames; arbitrary outbound is not permitted from request-bound code paths.
- Output encoding: React (web app) and Jinja2 (admin) auto-escape; explicit `dangerouslySetInnerHTML` usage is grep-scrubbed in CI.

3.6 LLM-specific protections

The platform integrates with Together AI (default inference), Anthropic and OpenAI (elevated/fallback inference lanes), Voyage AI (document

embeddings), and RunPod (optional self-hosted serving, off by default). Specific risks:

- **Prompt injection from uploaded documents.** All RAG-retrieved content is delivered to the LLM in a clearly-delimited "context" block; instructions in retrieved documents are scoped against the system prompt and post-filtered for known jailbreak patterns. Tool-calling proposals are surfaced for user review before execution (the AIProposal lifecycle).
- **Data exfiltration via LLM.** No customer data leaves the platform's tenant context except as part of an LLM inference call; the AI provider's data-handling clauses are reviewed at vendor onboarding, and per-org opt-in is required before sending data to providers that retain it for training.
- **Cost exhaustion.** Per-org and per-user quotas (configurable; default conservative). The public demo sandbox enforces a 10-message AI Coordinator quota per session.
- **Tenant-private fine-tuning.** Phase 3 of the platform's AI track introduces per-tenant LoRA adapters trained on the org's own data only. Adapters never cross tenants; the inference router selects the right adapter by `org_id` before calling the model server.

3.7 Secrets management

Surface	Storage	Rotation
Application secrets (JWT, Fernet, API keys)	Fly Secrets	Per provider — see `docs/runbooks/disaster-recovery.md` §3
Customer integration credentials (Salesforce, SAP, etc.)	Envelope-encrypted in `integration_connectors.credentials_json`; the key is in Fly Secrets	Customer-driven
OAuth client secrets	Fly Secrets	Annual or post-incident
Third-party API keys (Together, Anthropic, OpenAI, etc.)	Fly Secrets	Per provider; documented in DR runbook

No secret is ever committed to source control; pre-commit hooks and CI scan for accidental commits.

4. Detect — Monitoring & Audit

4.1 Application logging

Structured JSON logs are emitted for every request and security-relevant event. PII scrubbing is applied at the logging filter (`app.core.logging_filters.PIIScrubFilter`) before logs leave the application;

emails, IP addresses, and other identifiers are redacted in non-debug environments.

4.2 Audit trail

A dedicated `org_security_audits` append-only table records security-relevant events scoped by org:

- Authentication: login, logout, MFA enrollment, MFA challenge result, password reset, SSO bind/unbind.
- Authorization: role grants, role revocations, permission denials.
- Data access: PPAP submissions, FMEA edits, AI proposal accept/reject, integration credential read/write, sandbox toggle.
- Org lifecycle: creation, deletion request, deletion cancellation, conversion from demo.

Customer admins can request export of their org's audit history via the DSR pipeline. Retention is 7 years for regulated artifacts, 1 year for general events, configurable per-customer in the DPA.

4.3 Error tracking

Sentry captures application exceptions with the FastAPI integration (transactions tagged by endpoint), the SQLAlchemy integration (slow-query detection), and the PII scrubbing filter applied at the logging layer. Profiling sample rate is set to 1% of traced transactions to balance cost and signal.

4.4 Infrastructure monitoring

Health probes (`/livez` , `/readyz`) report into Fly's automated health checking. Readiness pings the database (hard dependency) and Redis (soft, fail-open). Status page integration is on the roadmap.

4.5 Anomaly detection

Rate-limit middleware emits standard `X-RateLimit-{Limit,Remaining,Reset}` headers and logs anomalies (sudden spikes from a single IP) for review. Integration with a SIEM is on the roadmap once SOC 2 Type II is signed.

5. Respond — Incident Response

5.1 Incident classification

Severity	Examples	Response
P0	Confirmed data breach; full platform outage > 30 min	Page on-call immediately; status page update within 15 min; war-room
P1	Partial outage; non-data-affecting security event	On-call paged; status update within 30 min
P2	Service degradation; suspicious activity under investigation	Triage during business hours
P3	Reported vulnerability under triage; no active exploit	Acknowledged within 48h

5.2 First-30-minute checklist

A paste-ready operator checklist lives in `docs/runbooks/disaster-recovery.md` §5:

1. Confirm impact via `/readyz`.
2. Probe Fly machine status and recent logs.
3. Triage the failure class (DB / Redis / app bug / Fly infra).
4. Post status page update before fixing — sets customer expectations.
5. Mitigate per the relevant runbook section.
6. Communicate to affected tenants directly.

5.3 Customer notification

Where legally required for personal-data breaches under GDPR Article 33, customer notification is targeted within **72 hours of confirmed awareness**. The notification includes the nature of the breach, categories and approximate number of data subjects, likely consequences, and the measures taken or proposed to address it.

5.4 Post-incident review

A blameless post-mortem is written within 5 business days for every P0 and P1 incident, capturing the timeline, root cause, contributing factors, and concrete follow-ups. Post-mortems are filed internally and shared with affected customers under NDA.

5.5 Vulnerability disclosure

We publish a `security.txt` at `/.well-known/security.txt` with a report contact (security@devaisuite.com) and a 90-day disclosure window. Researchers acting in good faith are not subject to legal action under our safe-harbor language.

6. Recover — Disaster Recovery & Continuity

6.1 Backup posture

PostgreSQL automated daily snapshots are retained for 5 days. RTO target for full-stack restoration is 30 minutes; RPO is up to 24 hours of writes (sub-day point-in-time recovery requires migration to Fly Managed Postgres, on the roadmap).

6.2 Restoration procedure

Documented in detail in `docs/runbooks/disaster-recovery.md` §1. Summarised:

1. Stop the API to prevent split-brain writes.
2. Identify the snapshot ID before the corruption window.
3. Restore as a new Postgres app (do NOT overwrite in place — preserves the original for forensics).
4. Validate against business-metric baselines.
5. Repoint `DATABASE_URL` to the restored cluster.
6. Bring the API back up.
7. Smoke-test via `/readyz` and the per-provider health endpoints.

6.3 R2 (object storage) recovery

Cloudflare R2 is single-region. "Failover" for the global template bucket means regenerating from source code via `publish_global_library`; for the tenant-doc bucket it requires R2 versioning to be enabled (verified at deployment).

6.4 DR rehearsal cadence

Quarterly rehearsals against staging Postgres (`devai-api-stg-pg`):

Quarter	Focus
Q1	Postgres restore (full §1.5 staging rehearsal)
Q2	Secret rotation (rotate one provider end-to-end)
Q3	R2 bucket regenerate from source
Q4	Tabletop incident drill (paper exercise)

A failed rehearsal triggers a P1 to update either the runbook or the underlying configuration before the next rehearsal cadence.

6.5 Business continuity

Critical business functions (billing, customer support, on-call rotation) are documented in playbooks separate from the platform itself, so platform unavailability does not block customer recovery efforts.

7. Vendor & Subprocessor Management

7.1 Onboarding

Each new subprocessor is reviewed for:

- Data residency (where they store our customer data).
- Contractual posture: DPA, SCCs (for non-EEA processors handling EU personal data), liability terms.
- Security posture: published certifications (SOC 2, ISO 27001), public security pages, breach history.
- Operational dependency: criticality to the service, ease of switching.

The decision is recorded in the vendor register; renewal review is annual.

7.2 Current subprocessor list (excerpt)

Provider	Purpose	Region	Key compliance
Fly.io	Compute (app + workers)	EU (Frankfurt)	SOC 2 Type II; ISO 27001-certified datacenters
Neon	Managed PostgreSQL + pgvector (tenant data at rest)	EU (Frankfurt, aws-eu-central-1)	SOC 2 Type II, ISO 27001:2022, ISO 27701
Cloudflare	DNS, R2 storage, Turnstile	EU / global edge	SOC 2 Type II, ISO 27001, ISO 27701, PCI DSS
Together AI	LLM inference (default)	US	SOC 2 Type II, ISO 27001:2022
Voyage AI	Document embeddings (RAG)	US	DPA in place; SOC 2 / ISO 27001 to confirm (MongoDB company)
Anthropic	LLM inference (Claude), elevated lanes	US	SOC 2 Type II, ISO 27001:2022, ISO 42001
OpenAI	LLM inference (GPT), configurable fallback	US	SOC 2 Type II, ISO 27001
Stripe	Billing & payments	US, EU	SOC 1, SOC 2 Type II, PCI DSS Level 1, ISO 27001
Resend	Transactional email	US	SOC 2 Type II
Sentry	Error tracking (PII-scrubbed)	US	SOC 2 Type II, ISO 27001
Inngest	Background job orchestration	US	SOC 2 Type II

Provider	Purpose	Region	Key compliance
RunPod	Self-hosted model serving (optional; off by default)	US / EU	SOC 2 Type II (Secure Cloud)
HuggingFace	Adapter artifact storage	US	SOC 2 Type II

The full, always-current list with version numbers and links to each provider's DPA is published at `/subprocessor-list.html`. Customer notification of subprocessor changes follows the cadence in the customer's DPA (typically 30 days advance notice for material changes).

7.3 Subprocessor change review

Customers are notified of subprocessor changes per DPA cadence and may object within the contractual window. Where a customer's objection cannot be reconciled with the new subprocessor, the contract permits termination per the DPA.

8. Secure Software Development Lifecycle

8.1 Source control & review

- Single canonical repository on GitHub.
- Branch protection on `main`: required PR reviews, required status checks, force-push disabled.
- Pre-commit hooks: lint, type check, secret scan.
- CI on every PR: full test suite (currently ~3,000 tests across smoke and full tiers).

8.2 Testing

Test type	Scope	Run cadence
Unit tests	Service logic, validators	Every PR
Integration tests	Routes + DB + middleware via FastAPI TestClient	Every PR
Smoke tests	Highest-leverage 200-test subset	Every PR (PR-tier CI)
Full sweep	All 3,000 tests	On push to `main` / `staging`
Schema-drift gate	OpenAPI spec generated against committed Python	Every PR
Vitest (frontend)	React components, hooks	Every PR
Type checks	mypy (Python) + tsc (TypeScript)	Every PR
Security regressions	Specific test files for RLS, RBAC, auth, billing	Every PR (smoke tier)

8.3 Dependency management

- GitHub Dependabot enabled; weekly PRs for security advisories.
- Manual review of every dependency update; security-relevant patches land within 7 days for high-severity advisories, 30 days for medium.
- No outdated framework majors: Python 3.11 (CI canonical), Node 20+, FastAPI / Pydantic / SQLAlchemy current.

8.4 Migration discipline

Database schema changes flow through Alembic migrations, reviewed and idempotent. Production runs `alembic upgrade head` as a release-

command pre-flight. Rollbacks are tested; destructive migrations carry a `--confirm` gate.

8.5 Static analysis

- Type checking is the primary static-analysis tool (mypy + tsc strict).
- `flake8` for Python style.
- ESLint for TypeScript.
- A future SAST integration (Semgrep or similar) is on the SOC 2 roadmap.

8.6 Configuration validation at startup

`validate_production_config(settings)` runs in the FastAPI lifespan startup. The application refuses to boot when:

- A required production secret is missing.
- `STRIPE_LIVE=true` is set but the Stripe key is not an `sk_live_` key (live-mode mismatch).
- Other invariants required to avoid unsafe states.

This catches configuration drift before it can serve a single request.

9. Data Handling

9.1 Data lifecycle

Stage	Handling
Collection	Customer-driven via the platform UI, API, or document upload. We do not scrape or buy customer data.
Storage	Encrypted at rest, RLS-isolated, region-pinned to EU (FRA) for the primary database.
Use	Bounded by tenant context on every request. AI inference calls are clearly delimited and the AI provider is named in the org's audit log.
Retention	Per the customer's DPA. Default: artifacts retained for the contract duration + 7 years for regulated data; general data retained for the contract duration + 1 year.
Deletion	Self-service org deletion via the admin UI (Sprint 8 D3 — 30-day grace + restore window, then hard-delete via the daily purge cron). DSR-driven user deletion is processed within 30 days.

9.2 Data residency

The primary database and application tier run in Fly's Frankfurt (FRA) region. Cloudflare R2 buckets are in EU jurisdiction; LLM provider regions vary and are documented in the subprocessor list.

For customers with strict residency requirements, on-premise / dedicated-region deployment is available on Enterprise plans.

9.3 Data subject rights (GDPR / CCPA / CPRA)

The DSR pipeline supports:

- Right to access — full export of the customer's data on request.
- Right to deletion — within 30 days, with confirmation to the requester.

- Right to rectification — handled via the platform UI by org admins.
- Right to portability — JSON export aligned with the same export contract as access.
- Right to object / restrict — handled per the DPA's processing-purposes table.

Requests are logged and tracked through the audit trail.

9.4 Cross-border transfer

Where personal data of EU subjects is processed by US-based subprocessors (LLM providers, Stripe, etc.), the transfer is governed by:

- Standard Contractual Clauses (SCCs) included in the DPA.
- Supplementary measures where required (encryption-in-transit, no plaintext logging at the provider).

The full transfer-impact assessment is internal but available to enterprise customers under NDA.

10. AI-Specific Considerations

10.1 What's sent to LLM providers

Inference calls send:

- The user's prompt (system prompt + user message).
- Retrieved RAG context (user's own documents only).
- Per-tenant adapter pointer (Phase 3 inference router only — never an adapter from another tenant).

What's NOT sent:

- Other customers' data.
- Internal application secrets, API keys, system credentials.

- The customer's authentication tokens.

10.2 Provider data-handling

Each LLM provider's data-handling clauses are reviewed at vendor onboarding. Providers that retain prompts/responses for training are gated behind a per-org opt-in flag (`ai_training_opt_in`); the default is opt-out.

10.3 AI proposal review

The platform's AI Coordinator never executes destructive actions silently. Any LLM-generated change to customer data flows through the AIProposal lifecycle:

1. AI generates a proposal (FMEA row, dashboard tile, document classification).
2. Proposal is stored with `status=draft` and surfaced to the user.
3. User reviews and explicitly accepts; rejection is also recorded with optional correction text for tenant-learning.
4. Only on accept does the platform commit the change to customer data.

This prevents silent data corruption from a hallucinated LLM response and creates an audit trail for every AI-driven change.

10.4 Tenant-learning isolation

Phase 3 of the AI track introduces per-tenant LoRA adapters trained on the org's own corrections to AI proposals. Critical isolation properties:

- Training data for each adapter is sourced exclusively from one org's `tenant_learning_events` rows, scoped by RLS.
- Trained adapters are tagged with `org_id` and stored in HuggingFace Hub under the organisation's namespace.
- The inference router selects the active adapter for the request's `org_id` before dispatching to RunPod / Together; cross-tenant

adapter selection is impossible by design (tested in the Phase 3c wiring suite).

- Demo orgs (`is_demo=True`) and sandbox orgs (`is_sandbox=True`) are excluded from training corpora.

11. Customer Responsibilities (Shared Model)

Security is shared. Customers are responsible for:

- **User management:** adding, removing, and reviewing authorised users in their org. Disable departed users promptly.
- **Credential hygiene:** strong passwords; MFA where the org policy permits it; not sharing credentials.
- **Configuration:** enabling allowed-email-domains, MFA enforcement, and the data-residency settings appropriate for the customer's compliance posture.
- **Data classification:** assessing what data they choose to upload to the platform. The platform is suitable for typical APQP/ manufacturing data; customers with stricter classifications (export-controlled, classified, etc.) should consult before upload.
- **Output validation:** AI-generated outputs must be reviewed before being relied upon for regulated decisions (PPAP submissions, customer deliverables, etc.). The AI Coordinator's confidence indicators are a starting point, not a substitute for SME review.
- **Integration credentials:** rotating their own credentials for connected systems (Salesforce, SAP, ERP, MES) per their internal policy.

A copy of this customer-responsibility list is in the Master Service Agreement and the DPA.

12. Contacts

Channel	Address
Security inquiries	security@devaisuite.com
Vulnerability disclosure	/.well-known/security.txt
Privacy / DPA / DSR	privacy@devaisuite.com
Status page	https://status.devaisuite.com
Sales / contracts	sales@devaisuite.com

PGP public key for security correspondence is available on request.

13. Document control

Field	Value
Version	1.0
Effective date	March 12, 2026
Last updated	April 29, 2026
Owner	Security Lead
Review cadence	Annual + post-incident
Distribution	Public (this version); internal-only versions track in-flight controls

This white paper is reviewed at least annually and on every material change to the platform's security posture. The authoritative source is `docs/security/WHITEPAPER_EN.md` in the platform's source repository.

A Spanish-language version is maintained in parallel at `docs/security/WHITEPAPER_ES.md` and rendered at `/es/security-whitepaper.html`.