



Pritect Security Architecture Whitepaper

Pritect Security Architecture Whitepaper

Version 4.1 · April 2026

Document title	Pritect Security Architecture Whitepaper
Version	4.1
Date	April 2026
Classification	Confidential — distribute to prospects and customers under written agreement
Document owner	Cybersecurity Partner, White Label Consultancy AS
Review cycle	Quarterly; next scheduled review July 2026
Contact	security@pritect.ai
Authoritative source	Live status, certifications, MSA, DPA and sub-processor list at pritect.ai/trust

Revision history

Version	Date	Summary of change
1.0	November 2025	First customer release.
2.0	January 2026	Added detail on enterprise features and expanded security narrative.
3.0	March 2026	ISO/IEC 27001 alignment, CAIQ v4 structure, enterprise commitments.
4.0	April 2026	New AI & LLM section, sub-processor change-notification commitment, refreshed compliance roadmap.
4.1	April 2026	Clarification of breach notification roles (Controller vs Processor).

This whitepaper is a summary written for security, privacy and procurement reviewers evaluating the Pritect platform. For binding terms, live certification status, and the authoritative sub-processor register, please refer to the Pritect Trust Center at pritect.ai/trust. A supporting spreadsheet with detailed CAIQ v4 and common questionnaire responses is available on request.

1. Executive summary

Pritect is a multi-tenant SaaS platform for privacy, data protection, cybersecurity, AI governance and enterprise risk management, operated by White Label Consultancy Group of Companies. Its security architecture is designed so that tenant isolation, authentication, authorisation and auditability are enforced independently at multiple layers, with customer data held in the European Union and governed under a published Data Processing Agreement.

Four ideas summarise how we think about security:

- Defence in depth. Authorisation is enforced at the database, API and user-interface layers; no single control protects customer data on its own.
- Security as code. Security policies are expressed in version-controlled configuration, enforced by automated pipeline gates, and applied identically to every environment.
- Data stays in the EU. Customer data is hosted in Ireland on managed infrastructure; transfers to any sub-processor outside the EU rely on the transfer mechanisms set out in our Data Processing Agreement.
- Transparent posture. Certifications, legal documents and the sub-processor list are published at pritect.ai/trust. Where a control is on our roadmap rather than in place today, we say so.

This version of the whitepaper has been reduced in length relative to earlier versions. Detailed operational evidence — CAIQ v4 answers, control mappings and supporting artefacts — is provided in a separate spreadsheet available to prospective and current customers under a confidentiality agreement.

2. Quick reference card

Ten questions that come up in nearly every security questionnaire. Each item is expanded in the sections that follow or, where more detail is useful, in the supporting CAIQ spreadsheet.

#	Question	Pritect answer
1	Is data encrypted in transit and at rest?	Yes. TLS 1.3 for data in transit; AES-256 for data at rest on managed EU PostgreSQL. Sensitive fields can be additionally encrypted at the application layer.
2	Where is customer data stored?	In the European Union. Primary database in Ireland; edge delivery via a global CDN with EU-primary routing. No customer-identifying data is stored outside the EU.
3	Do you support MFA and SSO?	Yes. TOTP-based MFA with per-tenant policies (optional, admin-required or required). SSO via SAML 2.0 and OIDC with domain verification and auto-provisioning.
4	How is tenant data isolated?	Tenant isolation is enforced at three independent layers: session and JWT context, API-level permission checks, and PostgreSQL Row-Level Security on every customer table.
5	What is your uptime target?	99.9% monthly availability as a non-binding operational target (MSA §2.6). A 99.5% binding SLA is available on the Enterprise Custom tier.
6	Are you ISO 27001 or SOC 2 certified?	Not yet. ISO/IEC 27001:2022 is targeted in 2026 and SOC 2 Type II in 2027. Our controls are already aligned to ISO 27001:2022 Annex A, see the Trust Center for current status.
7	What is your breach notification timeline?	Supervisory authority (Pritect Controller) and affected customers (Pritect Processor) are notified without undue delay and no later than 72 hours after we become aware of a personal data breach (GDPR Arts. 33–34; DPA §6.1).
8	How do you handle data subject requests?	We do not expect receiving a DSR on controller's behalf due to nature of our platform. However, should we receive one we will notify the controller immediately and redirect the request. For each tenant on the Pritect platform, the platform provides a DSR portal.
9	What are your RTO and RPO targets?	RTO 4 hours / RPO 1 hour, supported by managed point-in-time recovery. Targets are documented in the BCDR plan and shared under NDA on request.
10	Do you use customer data to train AI models?	No. Customer data is not used to train generic LLMs. We explain our AI and Mistral usage in Section 7 of this whitepaper.

3. Platform scope and data residency

Pritect is delivered exclusively as a cloud-native multi-tenant SaaS service. If you require a different setup than our standard solution we might be able to help you through a custom solution. There are no on-premises components, no customer-managed infrastructure and no hybrid deployments. This means that every customer benefits from the same security controls, the same patch cadence and the same configuration baseline, there is no configuration drift between tenants.

Customer data: all personal data, assessment content, inventory records, risk registers, audit entries and uploaded files are stored in managed PostgreSQL hosted in Ireland (EU) by Supabase. Frontend content is delivered through a globally distributed Anycast CDN configured to route EU traffic to EU edge nodes first. Backups are held by the managed provider within the same regional boundary as the live database.

3.1 Sub-processors and transfer mechanisms

The authoritative sub-processor register is maintained at pritect.ai/trust/subprocessors. The table below lists the current sub-processors at the time of publication. Where a sub-processor operates outside the EU, the transfer is covered by a data transfer mechanism.

Sub-processor	Service	Location
Supabase, Inc.	Managed PostgreSQL, authentication, storage, edge functions	Ireland (EU)
Mistral AI SAS	Large-language-model provider for opt-in AI features	France (EU)
Friendly Captcha GmbH	Bot protection on public forms	Germany (EU)
Vercel, Inc.	Frontend CDN, edge network and performance monitoring	United States (SCCs + DPF)
Resend, Inc.	Transactional email	United States (SCCs + DPF)
Mapbox, Inc.	Maps and geocoding for in-product location features	United States (SCCs + DPF)
Stripe, Inc.	Payment processing (billing data only)	United States (SCCs + DPF)

We commit to providing at least 30 days' prior written notice of a new or replaced sub-processor that processes personal data. Customers may object to a new sub-processor in accordance with the DPA.

4. Architecture overview

The architecture separates concerns into four layers: an edge layer that delivers the frontend through a CDN with a web application firewall and DDoS protection; an application layer of serverless functions that verify authentication and authorisation before executing business logic; a data layer of managed PostgreSQL with Row-Level Security enforced on every customer table; and an operational layer that governs deployment, monitoring and incident response. Each layer implements independent security controls so that the failure of a single control does not expose customer data.

4.1 Authentication and authorisation

Users authenticate through a managed identity provider that issues short-lived, cryptographically signed JWT access tokens. Multi-factor authentication is available in all environments, enforced per tenant, and configurable for different role populations. Enterprise customers can federate identity through SAML 2.0 or OIDC with domain verification and just-in-time provisioning.

Authorisation is enforced at three independent layers: the user interface layer checks route guards and feature availability; the API layer validates role-based permissions and tenant membership before business logic runs; and the database layer evaluates Row-Level Security policies that restrict every query to the rows belonging to the requesting tenant. The three layers are independent by design so that a gap in one cannot be used to bypass the others.

4.2 Tenant isolation

Tenant isolation is implemented primarily at the database layer. Every customer-owned table carries a tenant identifier, and Row-Level Security policies filter reads and writes to that identifier on every query. Because the policy is enforced inside the database engine, application code cannot disable it, and a compromised application component cannot return another tenant's rows. Audit-log tables are additionally protected by database constraints that prevent accidental deletion.

4.3 Secure development and change control

All changes follow a mandatory path from development to pre-production and only then to production. Automated gates run unit tests, integration tests and accessibility checks before a change is eligible for promotion. Destructive database operations require out-of-band approval and are blocked from being introduced through ordinary feature branches. Production credentials are verified on every deployment, and deployments emit summaries used for later audit.

4.4 Network and edge protection

Edge traffic is terminated on a managed CDN with TLS 1.3, HSTS, a strict Content Security Policy, X-Frame-Options set to deny, a Permissions-Policy that disables device APIs, and a Referrer-Policy that limits referrer leakage. A web application firewall filters common HTTP attack patterns, and DDoS absorption is inherited from the CDN. Webhooks that leave the platform are protected against server-side request forgery by hostname and DNS-resolution restrictions.

5. Cryptography

The table below summarises where cryptography is used, which algorithms we rely on and how keys are managed. Today, symmetric keys for data-at-rest encryption are managed by our infrastructure providers. Customer-managed encryption keys (CMEK / BYOK) are on the 2027 roadmap; if this is a procurement-blocking requirement for your organisation please let us know so we can prioritise appropriately.

Usage	Algorithm	Notes
Data in transit	TLS 1.3	Enforced end-to-end, including WebSocket connections and outbound webhooks.
Data at rest	AES-256	Applied by the managed database and storage providers. Sensitive fields may additionally be encrypted at the application layer.
API key storage	SHA-256	API keys are never stored in plaintext. Issued values are shown once at creation.
JWT signing	RS256	Signing-key rotation is managed by the identity provider.
Session tokens	Managed provider	Short-lived access tokens; idle timeout and hard sign-out supported.

6. Identity, access and administrative controls

6.1 Role-based access control

Permissions are expressed as granular permission keys assigned to roles. Roles are distinct at the platform level and at the tenant level: a platform administrator cannot access tenant data without separate tenant-level membership, and tenant administrators cannot alter platform configuration. This separation is enforced at the API layer and, where data access is involved, at the database layer.

6.2 Session and credential management

Sessions are time-bound with idle-timeout enforcement and a warning before expiry. A hard sign-out invalidates refresh tokens and requires full reauthentication. Password policy requires high-entropy passwords and integrates with breached-password screening; customers who prefer to eliminate password authentication may enforce SSO at the tenant level.

6.3 Privileged access to customer data by Pritect staff

Access to production systems and to customer data by Pritect staff is limited to a small number of named individuals and governed by the following controls: MFA on all administrative accounts; role-based, least-privilege assignments reviewed quarterly; documented justification for any access to tenant content; separation between engineering delivery accounts and any account capable of reading tenant data; and audit logging of administrative actions that is retained for at least twelve months.

6.4 API keys and machine-to-machine access

API keys are scoped to specific permissions, rate-limited per key, and fully auditable. Keys are hashed at rest and are revocable with immediate effect; every revocation is logged. Webhook endpoints require authenticated delivery with signature validation.

7. AI features and use of Mistral

Pritect uses a large-language-model (LLM) provider, Mistral AI, to power optional AI-assisted features inside the product, for example, drafting narratives in a risk register, suggesting ROPA entries, or summarising long policy text. Given that AI governance is one of the domains Pritect is sold to support, we hold our own AI usage to the same standards we help customers enforce.

The table below summarises how we use Mistral and the controls that apply. The authoritative version of these commitments is the Data Processing Agreement and the sub-processor register at pritect.ai/trust.

LLM provider	Mistral AI SAS, headquartered in Paris, France. Models are served from EU infrastructure.
Where Mistral is used	Only in opt-in AI features of the product, such as AI-assisted drafting of ROPA entries, risk narratives and questionnaire responses. AI features are clearly labelled in the UI, the system can be used without them.
Data minimisation	Only the content the user explicitly submits to an AI feature is sent to Mistral, no broader customer dataset is forwarded.
Prompt-injection defence	User-supplied content is sanitised to strip LLM-specific control markers before it reaches the model; structured system prompts separate trusted instructions from user content.
Training on customer data	Customer content is not used to train generic Mistral models. Our commercial terms with Mistral prohibit the use of our inputs or outputs for model training. This flows through to customers through our DPA.
Retention at the model provider	Prompts and completions are retained by Mistral only for the short operational window needed to deliver the response and to protect against abuse. Retention terms are mirrored in the sub-processor register.
Transfer mechanism	Processing takes place in the EU. No personal data is transferred outside the EU for AI processing under current configuration.
Logging and auditability	Every AI-feature invocation produces an audit-log entry with the tenant, user, feature, timestamp and a hash of the prompt. Customers can export this log for their own AI-governance records.
Output handling	AI-generated content is presented as a suggestion and must be accepted by a human before it becomes part of a customer record. Outputs are treated as user-generated content and subject to the same access controls.

7.1 Customer controls

- Per-user visibility. Users can see exactly when an output was produced by an AI feature; AI-generated suggestions are clearly labelled.

If you are evaluating Pritect against an internal AI-governance policy, we are happy to respond to a dedicated AI questionnaire (ISO/IEC 42001, NIST AI RMF, or your own) — contact security@pritect.ai.

8. Business continuity and disaster recovery

8.1 Backups and point-in-time recovery

The primary database is backed up by the managed provider with point-in-time recovery available within a defined retention window. Backups inherit the EU residency of the live database and are encrypted at rest. Backup integrity is verified continuously by the provider; restoration drills are performed on a defined cadence as part of the BCDR plan, and any deficiencies are tracked as corrective actions.

8.2 Recovery objectives

The platform targets a Recovery Time Objective of 4 hours and a Recovery Point Objective of 1 hour. These targets are documented, reviewed annually, and shared with customers under NDA on request. Customers with stricter operational requirements can discuss enhanced targets under the Enterprise Custom tier.

8.3 Availability

The platform targets 99.9% monthly availability as a non-binding operational commitment defined in the Master Service Agreement. Planned maintenance is announced on the status page no less than 72 hours in advance; unplanned downtime triggers a root-cause analysis initiated within 48 hours of resolution.

8.4 Regional architecture

The database currently operates in a single EU region. A multi-region active-passive architecture is being evaluated in Q4 2026 alongside the ISO 27001 programme; any change will be communicated through the Trust Center and this whitepaper.

9. Data lifecycle: retention, export and deletion

Customer data is handled throughout its lifecycle in a way that supports GDPR rights, regulatory recordkeeping, and operational recovery. The retention periods below apply unless a contract, legal hold or applicable regulation requires a different period.

Data category	Retention	Basis
Audit logs (authentication, authorisation, admin)	12 months; extendable to 24 months on request	ISO 27001 A.5.33; GDPR Art. 5(1)(e); NIS2 Art. 23
API request logs	30 days	Operational monitoring
Security event logs	12 months	Incident investigation and forensic analysis
Customer data during contract	Duration of contract + 30 days	GDPR Art. 6(1)(b); contractual obligation
Customer data post-termination	Deleted within 30 days of termination	GDPR Art. 17; MSA
Backup snapshots	Per managed-provider policy (typically 7–30 days)	Business continuity
Training completion records	Employment duration + 12 months	Compliance evidence

9.1 Export and portability

Customers can export their data in standard, machine-readable formats (CSV) through an in-product self-service flow. Additional formats are available on request. Exports are authenticated, logged and rate-limited.

9.2 Deletion and soft delete

Records that users deactivate are soft-deleted: marked as deactivated via a timestamp, so that the platform can support forensic review, audit holds, and legitimate reversal. On contract termination, customer data is deleted in accordance with the DPA, usually within 90 days, with a written deletion attestation available on request.

10. Incident response and breach notification

Detection and triage	Automated alerts trigger initial triage promptly during business hours with an on-call rotation for out-of-hours events.
Containment	Containment actions are initiated within four hours of a confirmed incident classification.
Supervisory authority	For personal-data breaches, notification to the relevant supervisory authority within 72 hours of becoming aware (GDPR Art. 33) for any personal data for which Pritect is the Controller.
Customer notification	Affected customers are notified without undue delay and no later than 72 hours after we become aware of the breach (GDPR Art. 34; DPA §6.1). Notification is sent to the customer's security and privacy contacts of record.
Root-cause analysis	A post-incident report is completed within 10 business days covering root cause, timeline, impact and remediation, and shared with affected customers on request.
Severity classification	Four-tier scale (P1–P4). P1 incidents activate the full response team with continuous status updates until resolution.
Annual testing	Incident-response procedures are tested annually as part of the ISO 27001 management-review cycle.

10.1 Vulnerability disclosure

Security researchers and customers may report vulnerabilities to security@pritect.ai. Reports are acknowledged within one business day and triaged within two business days. We commit to a good-faith safe-harbour position for researchers who act within our disclosure policy: we will not pursue legal action for research that stays within the scope of the policy and avoids privacy-invasive or service-disruptive testing. Confirmed vulnerabilities are triaged with CVSS v3.1 and remediated against the following targets: Critical 24–48 hours, High 7 days, Medium 30 days, Low 90 days. A formal bug-bounty programme is planned in 2027.

10.2 Third-party security testing

Annual external penetration testing by an independent firm is planned from 2026, aligned with the ISO 27001 programme. Summary letters of attestation will be available to enterprise customers under NDA after each test.

11. Compliance alignment and certifications

We are explicit about what is in place today versus what is on our roadmap. The authoritative status of every item below is published and kept current at pritect.ai/trust.

Framework	Status	Target / notes
GDPR (EU 2016/679)	Active	Supported by the DPA, DPIA process and EU-centric processing of personal data
ISO/IEC 27001:2022	Programme underway	Certification audit planned 2026; controls aligned to Annex A
SOC 2 Type II	Planned	Audit planned 2027 (audit firm to be selected)
ISO/IEC 27701 (privacy)	Roadmap	2027–2028 once 27001 is operating
ISO/IEC 42001 (AI management)	Roadmap	Under evaluation for 2027, aligned to our AI-governance module

In addition to our own certifications, our core infrastructure providers hold independent third-party certifications that cover the layers we build on. Supabase, Vercel and GitHub each maintain SOC 2 Type II; Vercel and GitHub additionally hold ISO 27001. Customers may request these provider reports as part of vendor due diligence.

12. Shared responsibility

Security is a shared responsibility. Pritect provides a platform and operates it securely; the customer governs how the platform is used within their organisation. The table below summarises each side's obligations.

Pritect (White Label Consultancy)	Customer
Secure the platform: infrastructure, application, pipeline and shared services.	Govern user access: joiner/mover/leaver, access reviews, segregation of duties.
Operate the service: patching, monitoring, incident response, continuity.	Configure roles and permissions: align product configuration with your policy.
Enforce tenant isolation: database, API, and UI-layer authorisation.	Classify data: decide what content is placed in the platform and at what sensitivity.
Maintain audit logs: including admin actions and security events.	Integrate identity: configure SSO with your identity provider if federation is used.
Publish transparency: Trust Center, DPA, sub-processor register, notices.	Respond to data subjects: use the DSR tools and follow your own GDPR procedures.

13. Contact and next steps

- Security questions: security@pritect.ai
- Trust Center (certifications, MSA, DPA, sub-processor list, notices): pritect.ai/trust
- Status and incidents: status.pritect.ai

We are happy to complete your standard vendor security questionnaire, walk through the platform with your security team, and provide a sandbox environment. For a customer-specific pen-test attestation, AI governance questionnaire, or DORA control-mapping, email security@pritect.ai.

This document is superseded quarterly. Please use the version at pritect.ai/trust as the authoritative source.

