Pactone.ai Security Whitepaper

Version 1.0 | June 2026

Executive Summary

Pactone.ai is an agentic Contract Lifecycle Management (CLM) platform built for regulated enterprises. This whitepaper describes how Pactone secures contract data, enforces access at every API boundary, and maintains an immutable audit trail across AI and human actions.

Pactone is operated by **Vectro Consulting Services** and powered by **EXtream.AI**. The platform combines six specialist AI agents, eleven RBAC roles, OPA-enforced policy, and Postgres row-level security to provide a defensible, audit-ready contract workspace.

Why CLM Security is Different

Unlike generic SaaS applications, CLM platforms process an organization's most sensitive assets: contracts containing financial terms, liability caps, intellectual property, trade secrets, and personally identifiable information. A breach or misconfiguration isn't just a compliance violation—it can result in financial loss, litigation, and erosion of customer trust. Pactone was architected from day one with this reality in mind.

What This Whitepaper Covers

Section	Description
Trust Pillars	Foundational security principles underpinning the platform
Architecture Overview	System design and threat model
Identity, Access, and Policy	Authentication, authorization, and tenant isolation
Auditability and the Revalidation Loop	Immutable logging and AI governance

Data Handling	Encryption, residency, and AI training policy
Compliance Posture	Certifications, frameworks, and regulatory alignment
Third-Party Integrations	Sub processors and integration security
Vulnerability Management	Disclosure program and incident response
Security Operations	Monitoring, testing, and continuous improvement

Trust Pillars

Pactone's security model rests on seven foundational pillars:

Pillar	Description
Tenant Isolation	Postgres row-level security on every table ensures data from one customer is never visible to another
OPA-Enforced Authorization	Policy decisions are evaluated at every API boundary—the client is never trusted
Immutable Audit & Risk Logs	Every AI recommendation and human decision is recorded in an append-only trail
AI Revalidation Loop	Every human edit is sanity-checked by AI to detect newly introduced risk
Data Residency & Privacy	Customer contract data is never used to train shared AI models; data stays in-region
Defense in Depth	Multiple layers of security controls across network, application, and data layers
Zero Trust Architecture	Never trust, always verify—applied to every access request

1. Architecture Overview

1.1 System Architecture

Pactone is built as a multi-tenant SaaS platform with a zero-trust security model. The architecture follows industry best practices for cloud-native enterprise applications, with security controls embedded at every layer.

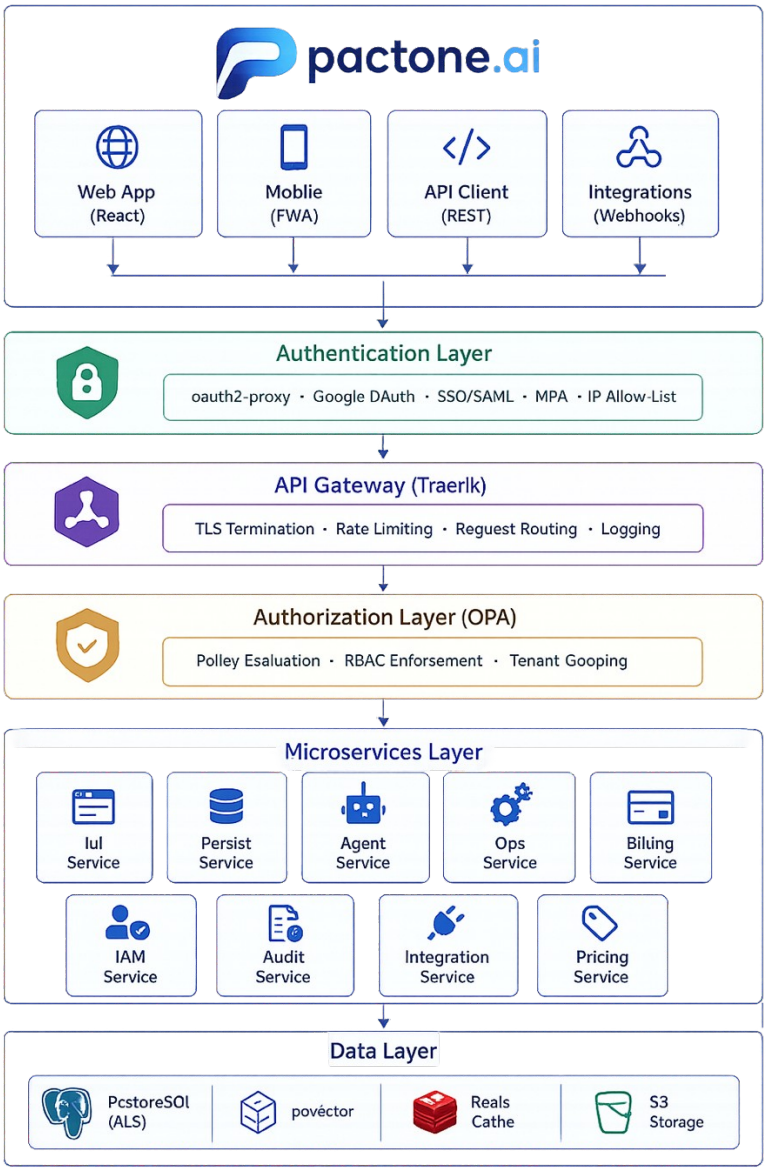

1.2 Threat Model

Pactone maintains a living threat model that is reviewed and updated quarterly. Key threat actors and scenarios include:

Threat Actor	Primary Threat	Mitigation
External Attacker	Data exfiltration via API exploitation	OPA policies, rate limiting, WAF

Malicious Insider	Privilege abuse, data theft	RBAC, RLS, immutable audit logs
AI Prompt Injection	Manipulation of AI outputs	Sanitization, HITL governance, audit
Cross-Tenant Access	Data leakage between customers	RLS, tenant middleware, OPA
Credential Theft	Unauthorized access	MFA, short-lived JWTs, IP allow-listing
Model Poisoning	Corrupted AI recommendations	Revalidation loop, human oversight

1.3 Security Design Principles

Pactone's architecture is guided by the following principles, aligned with NIST Cybersecurity Framework 2.0 and Zero Trust architecture:

Principle	Implementation
Least Privilege	Every user and service receives minimum permissions needed
Defense in Depth	Multiple independent layers of security controls
Zero Trust	Never trust, always verify—applied to every request
Secure by Default	Security is enabled out of the box, not opt-in
Assume Breach	Systems are designed assuming compromise may occur
Immutable Infrastructure	Infrastructure is rebuilt from code, not patched

2. Identity, Access, and Policy

2.1 RBAC — 11 Roles

Pactone ships with a comprehensive role set designed to mirror real-world organizational structures. Roles map to a least-privilege capability matrix and are assignable per workspace.

Role	Capabilities	Access Scope
Admin	Full system control: tenant configuration, user provisioning, OPA policy management, audit log access	All resources
Legal	Draft contracts, redline, analyze risk, manage clause library, view obligations—**cannot approve final execution**	Tenant-scoped

Finance	Manage billing, invoices, payment obligations, and financial contract terms	Tenant-scoped
Procurement	Initiate vendor contracts, manage RFQs, view legal edits—**cannot redline legal clauses**	Tenant-scoped
Ops	Manage workflows, renewals, SLA tracking, and task assignments	Tenant-scoped
Approver	Read contracts awaiting approval, execute (approve/reject) final versions	Tenant-scoped
Reviewer	Draft contracts, suggest edits, comment—**cannot approve or execute**	Tenant-scoped
Auditor	Read-only access to all contracts, obligations, and audit logs—ideal for compliance reviews	Tenant-scoped
Customer	Read-only access to own executed contracts, obligations, and invoices	Self-scoped
Vendor	Read-only access to own SOWs, NDAs, and payment status	Self-scoped
Integrator	API access scoped to programmatic operations; no UI access	Tenant-scoped

Each role is enforced at the database, application, and presentation layers. UI visibility is a UX concern—authorization is always verified server-side.

2.2 OPA-Enforced Policy

Every server function and HTTP route invokes an Open Policy Agent (OPA) policy check before executing business logic.

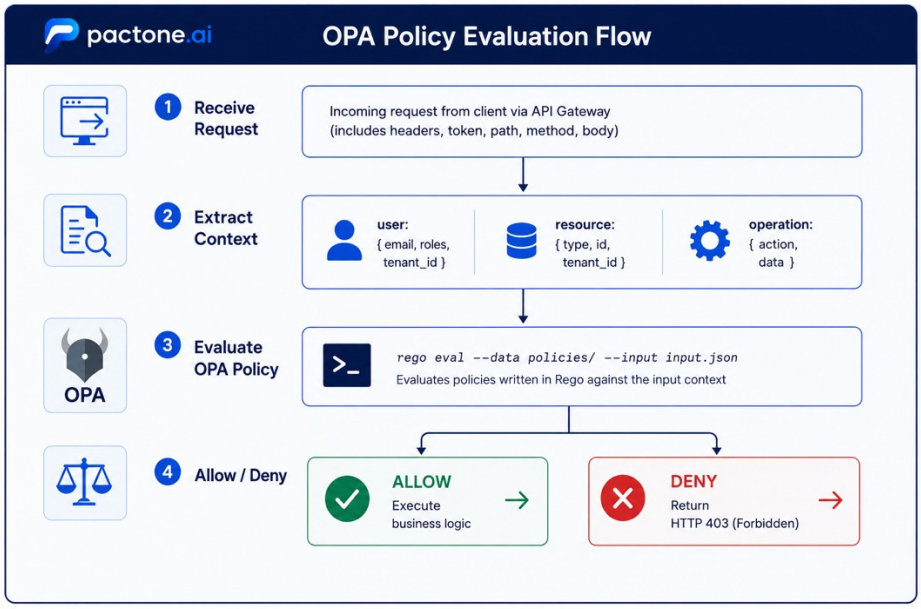

Policy Characteristics:

Attribute	Implementation
Language	Rego (Open Policy Agent)
Versioning	Git-managed with commit history
Review	Mandatory peer review before merging
Testing	Unit-tested with coverage requirements
Hot-Reload	Policy updates applied without service restart

Sample OPA Policy (Simplified):

```
package clm.auth

default allow = false

# Fetch user roles and tenant
user_roles = data.clm.users[input.user_email].roles
user_tenant = data.clm.users[input.user_email].tenant_id

# Admin can do everything
allow {
    "admin" in user_roles
```

```
}

# Legal can redline contracts
allow {
    "legal" in user_roles
    input.resource_type == "contract"
    input.operation in ["read", "update", "redline"]
}

# Approvers can only read and execute
allow {
    "approver" in user_roles
    input.resource_type == "contract"
    input.operation in ["read", "execute"]
}

# Finance can manage billing
allow {
    "finance" in user_roles
    input.resource_type == "billing"
    input.operation in ["read", "update", "pay"]
}

# Customers can only read their own contracts
allow {
    "customer" in user_roles
    input.resource_type == "contract"
    input.operation == "read"
    contract_owner = data.clm.contracts[input.resource_id].tenant_id
    contract_owner == user_tenant
}

# OPA data is published by IAM service
# Policies are version-controlled and audited
```

Key Principle: The client is never trusted as the source of authorization decisions. UI hiding is a UX concern, not a security control.

2.3 Authentication

Pactone supports multiple authentication methods with enterprise-grade security:

Authentication Method	Availability	Description
Email/Password	All tiers	Strong password requirements (12+ characters, complexity enforced)
Google OAuth	All tiers	OAuth 2.0 with Google Workspace accounts only
SSO/SAML	Enterprise	SAML 2.0 with any IdP (Okta, Azure AD, Auth0)
MFA	Enterprise	Optional multi-factor authentication via TOTP
IP Allow-Listing	Enterprise	Restrict access to corporate IP ranges
Session Timeout	Configurable	Per-workspace configurable session expiry

Password Policy:

- Minimum 12 characters
- At least one uppercase and one lowercase letter
- At least one number
- At least one special character
- Password history (last 5 passwords cannot be reused)
- Account lockout after 5 failed attempts

2.4 Tenant Isolation

Every table in the database is scoped by a `tenant_id` column with Postgres Row-Level Security (RLS) policies.

```sql
-- Simplified RLS Example
CREATE POLICY tenant_isolation ON contracts
    USING (tenant_id = current_setting('app.current_tenant')::UUID);

CREATE POLICY tenant_isolation ON obligations
    USING (tenant_id = current_setting('app.current_tenant')::UUID);
```

```
CREATE POLICY tenant_isolation ON users
    USING (tenant_id = current_setting('app.current_tenant')::UUID);

-- All tables are tenant-scoped
-- RLS is enabled by default and cannot be bypassed
```

How It Works:

1. **Middleware:** The Pactone service extracts `tenant_id` from the user's JWT or subdomain
2. **Context Setting:** `current_tenant` is set for the database session
3. **RLS Enforcement:** PostgreSQL automatically filters all queries
4. **No Bypass:** RLS cannot be disabled by application code

Result: Cross-tenant access is structurally impossible at the database layer—even if an API token is compromised.

2.5 API Security

Control	Implementation
Rate Limiting	Per-tenant and per-IP rate limits (configurable)
Request Validation	All inputs validated against JSON schemas
CORS	Strictly configured for approved origins only
JWT	Short-lived tokens (15 minutes) with refresh rotation
API Keys	Scoped to specific operations for Integrator role
Audit Logging	Every API request logged with user, IP, and timestamp

---

3. Auditability and the Revalidation Loop

3.1 Immutable Audit Trail

Every action in Pactone is recorded to an append-only audit log with complete context.

Logged Events Include:

- All AI agent invocations (Create, Redline, Risk, Obligation, Search, Renewal, Revalidation)
- All human decisions (Accept, Reject, Edit)
- All workflow state transitions (DRAFT → REVIEW → APPROVED → EXECUTED)
- All RBAC policy evaluations (Allow/Deny)
- All authentication attempts (success/failure)
- All configuration changes (tenant settings, role assignments)

Audit Record Structure:

json


```
{
  "id": "audit-123e4567-e89b-12d3-a456-426614174000",
  "tenant_id": "TENANT-001",
  "user_id": "user-123e4567",
  "user_email": "legal@acme.com",
  "user_role": "legal",
  "action": "AI_RECOMMENDATION_APPLIED",
  "resource_type": "contract",
  "resource_id": "contract-456",
  "before_state": {
    "clause_7_2": "Supplier's liability shall not exceed fees paid"
  },
  "after_state": {
    "clause_7_2": "Supplier's liability shall not exceed 3x fees paid"
  },
  "ai_agent": "redline_agent",
  "confidence_score": 0.92,
  "reasoning": "Liability cap is below playbook minimum (3x).",
  "ip_address": "192.168.1.100",
  "user_agent": "Chrome/120.0",
  "policy_decision": "ALLOW",
  "policy_evaluation_id": "policy-789",
  "created_at": "2026-06-19T14:32:18Z"
}
```

Key Properties:

Property	Implementation
Append-Only	Records are inserted; never updated or deleted
Immutable	Tamper-evident via cryptographic hash chaining
Replayable	Audit trail can be exported for compliance review
Retention	Configurable per tenant (default: 7 years)
Export	CSV, JSON, and Excel formats
Search	Full-text search across audit records

3.2 The Revalidation Loop

Pactone's **Human-in-the-Loop (HITL)** governance model ensures AI assists but never replaces human judgment.

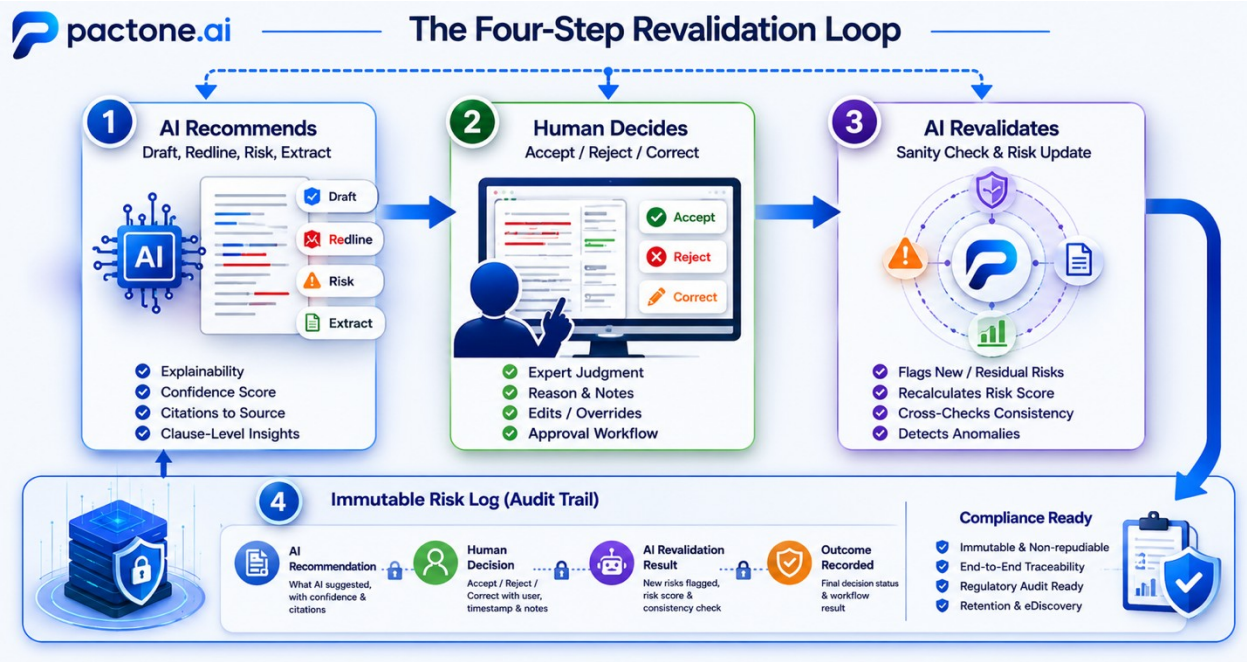

Why This Matters:

Feature	Security Benefit
AI Transparency	Every AI action includes reasoning and citations—no black-box decisions

Human Oversight	AI recommends; humans decide—critical for regulated industries
Revalidation	Human edits are sanity-checked to prevent introduction of new risks
Immutable Record	Complete chain of custody for every clause and decision

AI Governance Controls:

Control	Implementation
Explainability	Every AI output includes reasoning and confidence score
Auditability	All AI decisions logged with before/after state
Human Review	Mandatory human review before final contract execution
Revalidation	Automated sanity-check of human edits
Bias Mitigation	Regular review of AI recommendations for bias

4. Data Handling

4.1 Encryption

Layer	Method	Standard
Data at Rest	AES-256	FIPS 140-2 compliant
Data in Transit	TLS 1.2+	Perfect Forward Secrecy (PFS)
Database	Transparent Data Encryption (TDE)	AES-256
Backup	Encrypted with customer-managed keys	AES-256
AI Gateway	TLS 1.2+	PFS
API Endpoints	TLS 1.2+	PFS

4.2 Data Residency

Pactone supports data residency in three regions:

Region	Location	Availability	Latency
US	AWS US-East-1 (N. Virginia)	All tiers	< 50ms

EU	AWS EU-Central-1 (Frankfurt)	Enterprise	< 50ms
APAC	AWS AP-Southeast-2 (Sydney)	Enterprise	< 80ms

Key Commitments:

- All data and backups remain in the selected region
- No automatic cross-region replication
- Backup retention: 30 days minimum, configurable up to 365 days
- Data subject to local data protection laws (GDPR, CCPA, etc.)

4.3 AI Training Policy

Customer contract data is never used to train shared AI models.

Data Flow for AI Processing:

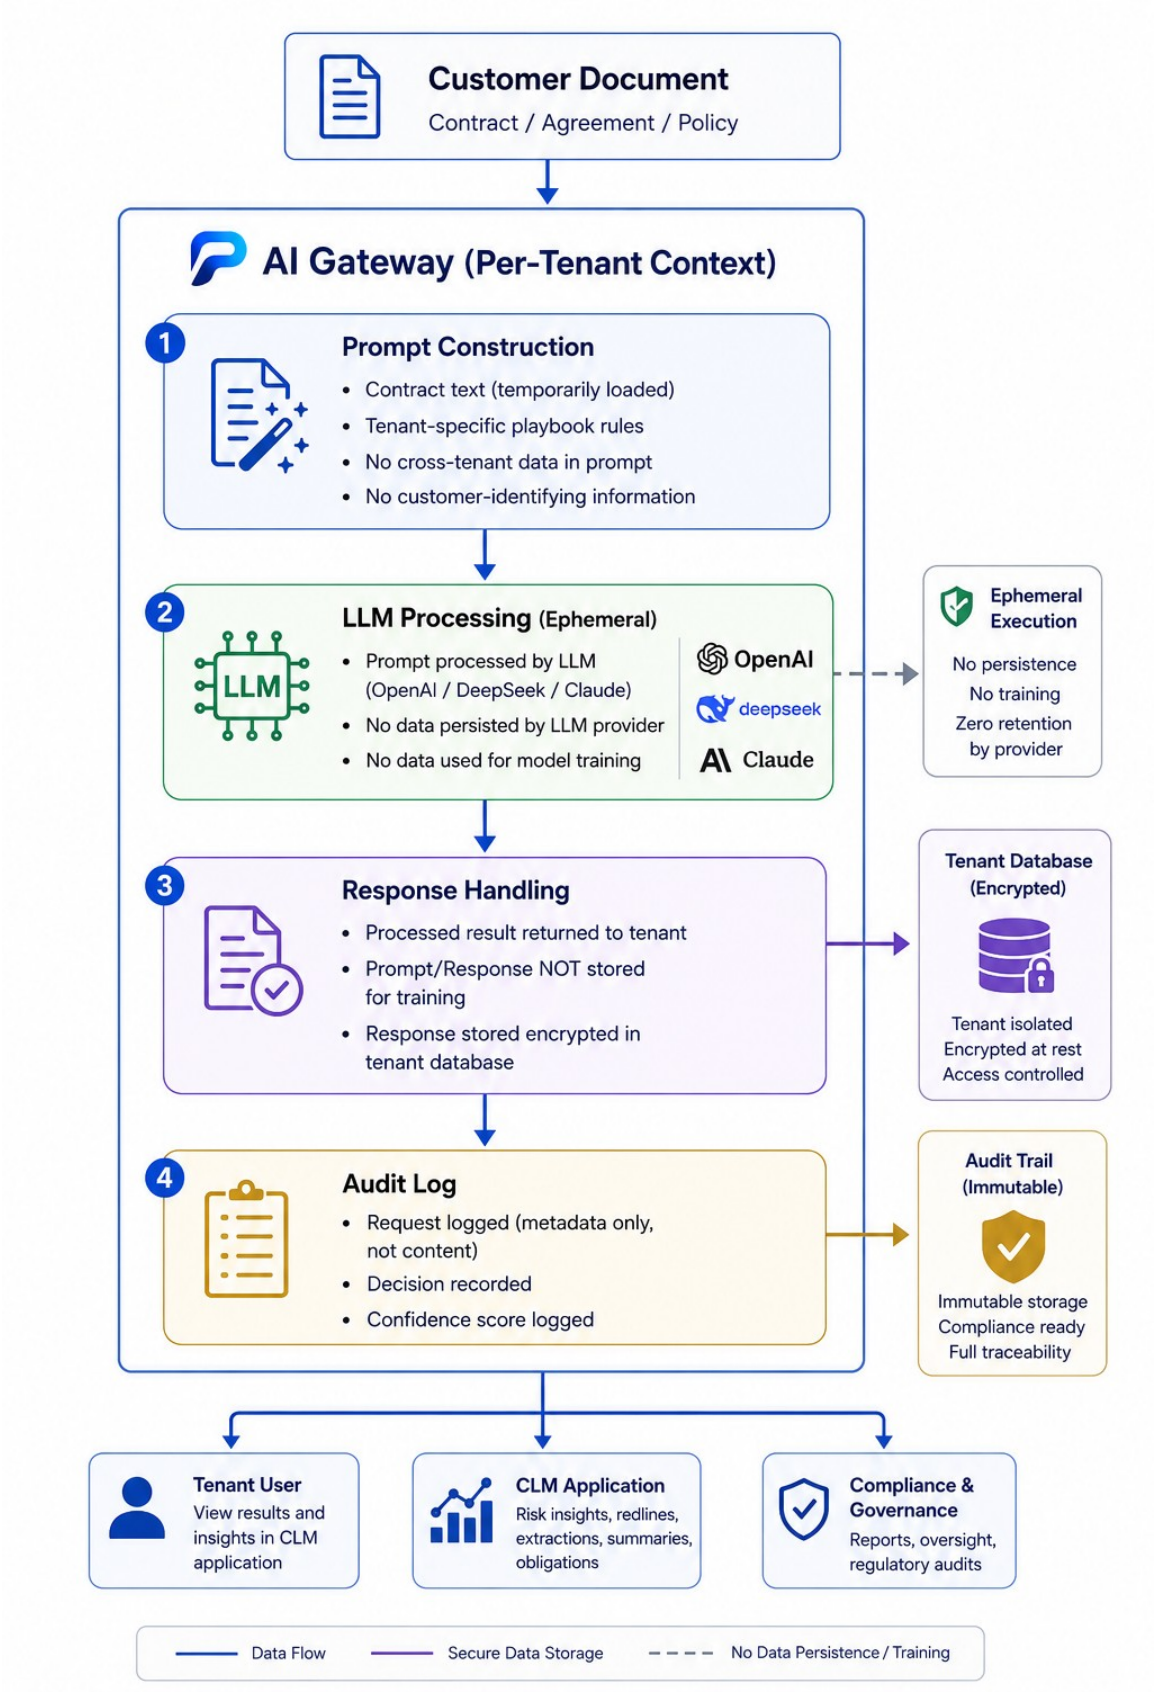
Customer Document
Contract / Agreement / Policy
AI Gateway (Per-Tenant Context)
1
Prompt Construction
Contract text (temporarily loaded)
Tenant-specific playbook rules
No cross-tenant data in prompt
No customer-identifying information
2
LLM Processing (Ephemeral)
LLM
Prompt processed by LLM (OpenAI / DeepSeek / Claude)
No data persisted by LLM provider
No data used for model training
OpenAI
deepseek
Claude
Ephemeral Execution
No persistence
No training
Zero retention by provider
3
Response Handling
Processed result returned to tenant
Prompt/Response NOT stored for training
Response stored encrypted in tenant database
Tenant Database (Encrypted)
Tenant isolated
Encrypted at rest
Access controlled
4
Audit Log
Request logged (metadata only, not content)
Decision recorded
Confidence score logged
Audit Trail (Immutable)
Immutable storage
Compliance ready
Full traceability
Tenant User
View results and insights in CLM application
CLM Application
Risk insights, redlines, extractions, summaries, obligations
Compliance & Governance
Reports, oversight, regulatory audits
Data Flow
Secure Data Storage
No Data Persistence / Training

Key Guarantees:

Guarantee	Implementation
No Cross-Tenant Leakage	Per-tenant prompt context; no cross-tenant data in prompts
No Model Training	Customer data never used to train shared models
Ephemeral Processing	AI interactions are processed, returned, and discarded
Audit-Only Logging	Audit logs capture decisions, not source text
LLM Provider Agreements	Data processing agreements with all LLM providers

4.4 Document Processing

Document Type	Processing Location	Duration	Retention	Encryption
Uploaded Contracts	In-region AI Gateway	Real-time	Stored encrypted	AES-256
Extracted Obligations	In-region database	Real-time	Stored encrypted	AES-256
Redlined Versions	In-region application	Real-time	Stored encrypted	AES-256
Audit Logs	In-region database	Real-time	7+ years	AES-256
Temporary Files	In-region ephemeral storage	Deleted after processing	Not stored	N/A

4.5 Data Classification

Pactone implements a data classification policy aligned with SOC 2 and ISO 27001 requirements:

Classification	Description	Examples	Controls
Critical	Data requiring highest protection	Contract content, financial terms, PII	AES-256 encryption, RLS, audit
Sensitive	Data requiring strong protection	User emails, obligation details	AES-256 encryption, RLS
Internal	Data requiring standard protection	Tenant configuration, audit logs	Standard encryption, RLS
Public	Data requiring minimal	Pricing plans, public	Standard web controls

	protection	documentation	

5. Compliance Posture

5.1 Certifications

Certification	Status	Tier	Description
GDPR-ready	□ Designed for compliance	All tiers	EU data protection requirements met
SOC 2 Type II	□ In progress	Enterprise	AICPA Trust Services Criteria
HIPAA	□ In progress	Enterprise	Health Insurance Portability and Accountability Act
ISO 27001	□ Planned	Enterprise	Information Security Management System
NIST CSF	□ Aligned	All tiers	NIST Cybersecurity Framework 2.0

SOC 2 Type II Details:

SOC 2 Type II is the baseline standard every enterprise CLM vendor should hold. It evaluates controls over an extended period (typically 6-12 months) against five Trust Services Criteria: Security, Availability, Processing Integrity, Confidentiality, and Privacy. Unlike vendor self-attestation, SOC 2 Type II provides independent verification from a CPA firm.

HIPAA Alignment:

Pactone's AI-native CLM platform complies with HIPAA's three core requirements:

- **Privacy Rule:** Protecting PHI confidentiality
- **Security Rule:** Implementing technical safeguards
- **Breach Notification Rule:** Reporting data incidents

5.2 Regulatory Alignment

Framework	Alignment	Description
GDPR	Full	Data subject rights, DPA, cross-border data transfer
CCPA/CPRA	Full	California consumer privacy rights
NIST CSF 2.0	Full	Govern, Identify, Protect, Detect, Respond, Recover functions
ISO 27001:2022	In progress	Annex A controls
PCI DSS	N/A	Not applicable (no credit card storage)

5.3 Compliance Features

Feature	Implementation
Data Subject Rights (GDPR)	Automated data export and deletion workflows
Data Processing Agreement (DPA)	Available on request
Subprocessor List	Available under NDA
Penetration Testing	Annual third-party pentests; summary available on request
Incident Response	Documented and tested; 24/7 security team
Vulnerability Disclosure	security@pactone.ai monitored 24/7
Audit Readiness	SOC 2 and ISO 27001 audit-ready documentation

5.4 Security Assurance

Threat Modeling:

Pactone undergoes regular threat modeling sessions covering:

- Data exfiltration via AI prompts
- Cross-tenant data leakage
- Privilege escalation through role manipulation
- Session hijacking and token replay
- AI prompt injection and model poisoning
- Supply chain attacks on dependencies

Mitigations:

Threat	Mitigation

Prompt Injection	Sanitization layer before LLM calls
Data Exfiltration	RLS + OPA double enforcement
Token Theft	Short-lived JWTs + refresh rotation
AI Hallucination	HITL governance + revalidation loop
Supply Chain	Regular dependency scanning and updates
Insider Threat	Least privilege + immutable audit logs

6. Third-Party Integrations

6.1 Subprocessors

Service	Purpose	Data Processing	Region	Certification
AWS	Infrastructure, database, storage	Encrypted	In-region	SOC 2, ISO 27001
Stripe	Payment processing	Transaction data	US/EU	PCI DSS Level 1
Google OAuth	Authentication	Email, profile	US	SOC 2
OpenAI API	AI processing (optional)	Contract text (not stored)	US/EU	SOC 2
LangChain	AI orchestration	Contract text (not stored)	US/EU	SOC 2
DeepSeek	AI processing (optional)	Contract text (not stored)	CN/EU	SOC 2

Subprocessor Commitments:

- All subprocessors undergo security due diligence
- Data processing agreements in place
- Subprocessor list updated quarterly
- Available under NDA on request

6.2 Integration Security

Integration	Auth Method	Data Shared	Risk Level
DocuSign/Adobe Sign	OAuth 2.0	Contract metadata, signer info	Medium
Salesforce	OAuth 2.0	Deal data, opportunity status	Medium
ServiceNow	API Key + OAuth	SLA metrics, incident records	Medium
Slack	Webhook	Notifications (no contract data)	Low
Workato	API Key	Workflow orchestration metadata	Low
Tableau	API Key	Aggregated analytics data	Low

7. Vulnerability Management

7.1 Disclosure Program

Security researchers and customers are encouraged to report vulnerabilities:

Step	Action
1	Email: security@pactone.ai
2	PGP Key: Available on request
3	Expected Response: Initial acknowledgment within 24 hours
4	Update Frequency: Every 72 hours during investigation
5	Resolution Timeline: Critical: 24 hours; High: 72 hours; Medium: 5 days; Low: 10 days

7.2 Bug Bounty

Pactone operates a responsible disclosure program:

Severity	Description	Reward Range
Critical	Remote code execution, data breach, privilege escalation	$1,000 – $5,000
High	Authentication bypass, unauthorized data access	$500 – $1,000
Medium	CSRF, XSS, information disclosure	$100 – $500
Low	Security misconfiguration, best practice violations	Recognition only

Out of Scope: Phishing, social engineering, physical attacks, third-party apps, DOS attacks.

7.3 Incident Response

Incident Response Process:

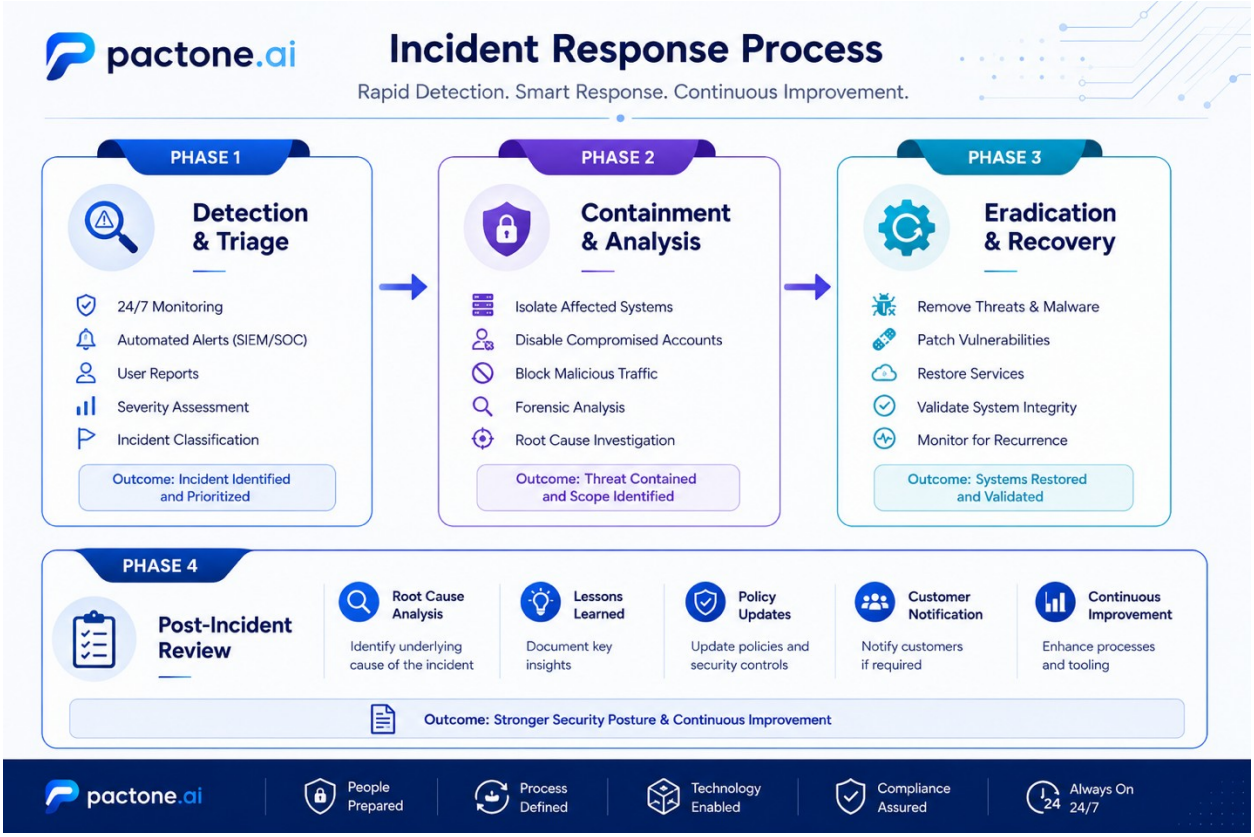

Response Commitments:

Severity	Response Time	Resolution Time	Customer Notification
Critical	< 1 hour	< 24 hours	Within 24 hours
High	< 4 hours	< 72 hours	Within 72 hours
Medium	< 24 hours	< 5 days	Not required
Low	< 48 hours	< 10 days	Not required

8. Security Operations

8.1 Continuous Monitoring

Capability	Implementation
Infrastructure Monitoring	AWS CloudWatch, Prometheus, Grafana
Application Performance	New Relic/DataDog (planned)
Security Logging	Centralized logging with alerting
Threat Detection	AWS GuardDuty (planned)
Vulnerability Scanning	Weekly automated scans
Dependency Scanning	Daily SCA scans (Dependabot, Snyk)

8.2 Testing Regimen

Test Type	Frequency	Scope
Unit Tests	Every commit	All services
Integration Tests	Every PR	API endpoints, OPA policies
Security Scans	Daily	Dependencies, containers
Penetration Tests	Annual	Full application + infrastructure
Red Team Exercises	Quarterly	Simulated attacks on staging
Policy Tests	Every PR	OPA policy coverage

8.3 Security Training

All Pactone engineers and staff undergo:

Training	Frequency	Description
Security Awareness	Annual	Phishing, social engineering, password hygiene
Secure Coding	Annual	OWASP Top 10, secure coding practices
Incident Response	Semi-annual	Tabletop exercises, playbook reviews
Data Privacy	Annual	GDPR, CCPA, data handling procedures

9. Summary of Security Posture

Capability	Implementation Status
Tenant Isolation	✅ Postgres RLS on every table
OPA Policy	✅ Versioned, tested, hot-reloadable
RBAC (11 Roles)	✅ Comprehensive role matrix
MFA	✅ Enterprise tier
SSO/SAML	✅ Enterprise tier
IP Allow-Listing	✅ Enterprise tier
Immutable Audit	✅ Append-only with replay capability
AI Revalidation	✅ Full HITL loop
Encryption at Rest	✅ AES-256
Encryption in Transit	✅ TLS 1.2+
Data Residency	✅ US, EU, APAC
GDPR-ready	✅ Designed for compliance
SOC 2 Type II	✅ In progress
HIPAA	✅ In progress
ISO 27001	✅ Planned
Customer Data Training	✅ Never used for model training
Subprocessor Transparency	✅ Available under NDA
Pentests	✅ Annual with summary available
Incident Response	✅ Documented and tested
Zero Trust Architecture	✅ Applied at all layers

Contact

For security questions, vulnerability reports, or to request the full compliance pack:

Purpose	Contact
Security & Vulnerability Reports	security@pactone.ai
Sales & Demos	ask@pactone.ai
Support	support@pactone.ai
Admin & Billing	admin@pactone.ai

Procurement	procure@pactone.ai
DPA Requests	dpa@pactone.ai

Response Commitment: All security inquiries are acknowledged within **24 hours** and addressed with appropriate urgency.

Document Control

Attribute	Value
Version	1.0
Last Updated	June 2026
Next Review	December 2026
Classification	Public
Owner	Security Team, Pactone.ai

Built for regulated industries. Security by design. Audit by default.