PERIDOT

Secure AI Application Platform

CISO SECURITY BRIEF

AI adoption has outpaced control. Peridot closes the gap—by design, not policy.

Your data never leaves your environment. Your users build AI apps. Your team stays in control.

THE PROBLEM

Employees use ChatGPT, Claude, and other AI tools—copying sensitive data into prompts, building apps outside IT, with zero visibility for security teams. The result: data leakage, compliance exposure, and shadow AI you can't audit.

ARCHITECTURE

- Runs inside your VPC or on-prem infrastructure
- Split-plane: control and data plane fully separated
- No external data processing or storage
- Stateless execution — zero data persistence
- Bring Your Own Key (BYOK) for all LLM access

COMPLIANCE IMPACT

Keeps all data processing within your controlled environment — materially reducing third-party exposure and simplifying SOC 2, HIPAA, and GDPR compliance. No new data processors. No amended DPAs. No model training on your data.

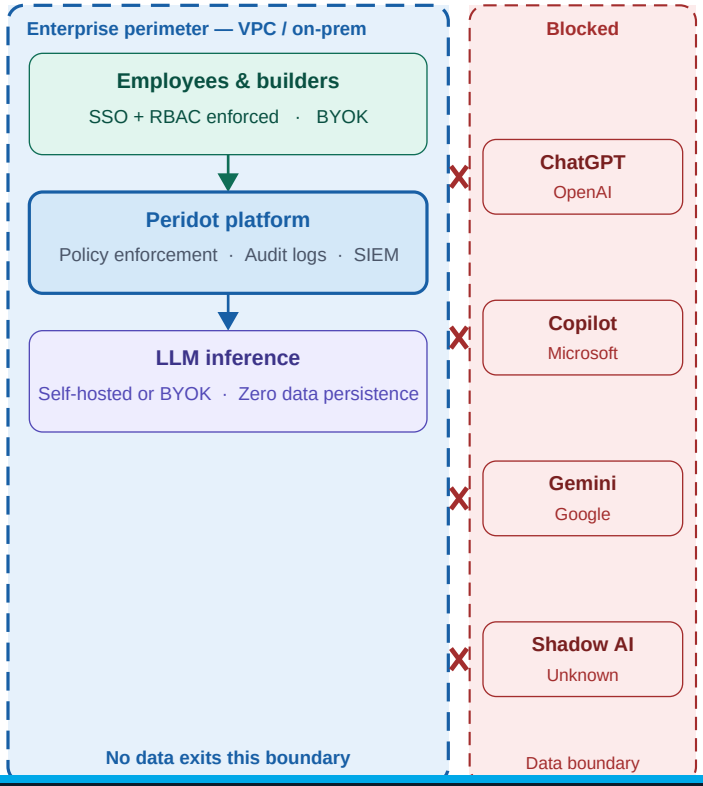

SECURITY MODEL

- Zero data egress — nothing leaves your environment
- No model proxying or hidden external routing
- BYOK: your keys, your control over LLM access
- Policy enforcement runs before execution
- No cross-tenant data access by design

GOVERNANCE & VISIBILITY

- Full audit logs for all actions and data access
- SIEM integration: Splunk, Datadog, and others
- Identity-based access: SSO + RBAC enforced
- No black-box model behavior — full observability

ISOLATION & BOUNDARIES

- Tenant-level isolation between all workloads
- Credential isolation — no shared secrets
- Application-level boundaries enforced at runtime
- No cross-system data leakage paths

RISK PROFILE VS. TYPICAL AI TOOLS

Risk vector	Typical AI tool	Peridot
Data leaves environment	Yes	**Never**
Model training on your data	Possible	**Impossible**
Audit trail	Limited	**Complete**
Compliance burden	Increases	**Reduces**

Ready to eliminate AI data risk in your environment?

Request a security architecture review: **peridot.company/audit**