TOTALLY PRIVATE AI

AI Adoption for Regulated Organisations

US Edition

A positioning brief for digital transformation advisors, enterprise architects, and technology consultants serving US federal agencies, state governments, and regulated enterprises.

February 2026 · Version 1.0 · Confidential

1. The AI Adoption Imperative in the United States

The United States is in the midst of a rapid, policy-driven acceleration of AI adoption across both government and the private sector. Executive Orders on artificial intelligence have established national priorities for safe, secure, and trustworthy AI deployment. The Office of Management and Budget's guidance on federal AI use (M-24-10) requires agencies to implement governance frameworks, designate Chief AI Officers, and manage AI risk, while simultaneously expanding AI capabilities.

In the private sector, AI adoption in regulated industries is moving from experimentation to production. Financial institutions, healthcare systems, and legal organizations are under competitive pressure to deploy AI for operational efficiency, client service, and regulatory compliance. The market is not waiting.

For enterprise architects and transformation consultants advising these organizations, the mandate is clear: deliver AI capabilities into production environments that are subject to some of the world's most demanding regulatory and litigation frameworks. Yet across sectors, adoption is stalling at the same point. The firms that solve this first will set the terms for the next wave of enterprise AI. The rest will be responding to RFPs that someone else shaped.

2. The Compliance Bottleneck

US regulated organizations face structural tension. AI adoption is a strategic and competitive imperative, but the regulatory environment creates significant barriers to deploying public AI platforms.

The compliance environment is dense and sector-specific:

- **Healthcare:** HIPAA's Privacy and Security Rules govern protected health information (PHI) with strict requirements for access controls, audit trails, encryption, and business associate agreements. Any AI platform processing PHI must satisfy these requirements by design, not by policy overlay.
- **Financial services:** GLBA and SOX impose data protection and internal controls requirements. The SEC's cybersecurity disclosure rules (2023) require public companies to report material cyber incidents. NYDFS 23 NYCRR 500 mandates encryption, access controls, audit trails, and third-party service provider oversight for financial institutions operating in New York.
- **State privacy laws:** CCPA/CPRA (California), VCDPA (Virginia), CPA (Colorado), CTDPA (Connecticut), and a growing patchwork of state-level privacy legislation create overlapping data handling, consent, and consumer rights obligations. There is no single federal privacy law. Organizations must work across a fragmented compliance environment.
- **Federal government:** NIST SP 800-53 provides the control framework for federal information systems. FedRAMP governs cloud service authorization for government use. CMMC requirements apply to defense contractors handling Controlled Unclassified Information (CUI). FISMA establishes baseline security standards across federal agencies.

When a hospital system, bank, or federal agency evaluates a public AI platform (ChatGPT, Copilot, Gemini), the conversation typically ends at one question: where does the data go? If the answer involves data leaving the organization's controlled environment, processing on vendor infrastructure with limited visibility, or insufficient audit controls, the engagement stops.

The organizations that need AI capabilities the most (federal agencies, banks, hospital systems, law firms) are the least able to adopt it under current market conditions.

For consultants advising these organizations, this is the moment that defines the engagement. You have built the AI business case, secured stakeholder alignment, and mapped the

deployment timeline. Then compliance asks: where does the data go? If you do not have a defensible answer, the engagement stalls, the budget review gets pushed, and the client starts looking for an advisor who does.

3. How TPAI Removes the Bottleneck

Totally Private AI (TPAI) is an air-gapped AI platform purpose-built for organizations that cannot accept the data exposure risk of public AI services. It provides access to leading foundation models, including Claude (Anthropic), Llama (Meta), and Mistral, within a security architecture designed around the requirements of regulated industries.

The practical proposition for regulated programs: maintain sovereign control over customer data while enabling production AI use in weeks, not quarters.

What "Air-Gapped" Means in Practice

- **No internet egress:** Application subnets have no public internet gateway or NAT. All AWS service access occurs via VPC private endpoints. This is a verifiable network configuration, not a marketing claim. For clarity, 'air-gapped' here means no public internet egress from application subnets and private service connectivity via approved endpoints; it does not imply physically disconnected hardware.
- **Per-user encryption:** Envelope encryption with per-user data encryption keys (DEKs) protected by AWS KMS. No cross-tenant data access by design.
- **Full audit trail:** Per-request logging with model, user, timestamp, and token metadata. SIEM-exportable for integration with existing security operations centers.
- **US data residency:** Deployed on AWS US infrastructure across multiple availability zones (us-east-1, us-east-2, us-west-1, us-west-2). US-region deployment supports data residency objectives; jurisdictional obligations remain subject to customer contract and applicable law. Region availability last reviewed: February 2026.
- **Minimal data exposure by design:** For Pro deployments, customer data is encrypted using AWS Customer Managed Keys (CMKs); TPAI does not access these keys in the normal course of operations. For Enterprise deployments, data is encrypted per user with tenant-controlled keys and administrator key copies for business continuity. In both models, TPAI holds no PII in its own systems. In normal operations, there is no readable customer data for TPAI to produce in a legal discovery scenario; actual data recoverability depends on deployment model, customer key governance, and applicable legal process.

Deployment Flexibility

TPAI is available as a managed SaaS service or deployed directly into the client's own AWS account. The client-deployed model is particularly relevant for organizations with the strictest data sovereignty requirements: data never leaves the client's cloud environment, and TPAI has no ongoing access to the deployment. The client controls the encryption keys, network boundaries, and data lifecycle.

How TPAI Compares

Approach	Time to Deploy	Data Sovereignty	Indicative Cost
Public AI platforms (ChatGPT, Copilot, Gemini)	Immediate	Data processed on vendor infrastructure. Location and handling typically outside customer control.	Low per-seat
Cloud AI private endpoints (Azure Private Link, etc.)	Days to weeks	Private connectivity improves network boundary control; data processing remains on provider-managed infrastructure unless client-hosted alternatives are used.	Moderate
Build your own AI stack	12-18 months	Full control, but significant engineering, maintenance, and ongoing operational burden.	$1.5M+ plus ongoing
TPAI	**Weeks**	**Air-gapped, per-user encryption, US-resident. Client-deployed option removes ongoing vendor access to customer data.**	**Pro: $299/seat/month** **Enterprise: from $130K + $30K/month**

TPAI is currently in beta. ISO 27001 certification targeted Q3 2026; SOC 2 Type II targeted Q4 2026. Security evidence packs are available under NDA for solution diligence, covering US, ME, EU, and APAC jurisdictions.

4. Use Cases by Sector

The following examples illustrate where transformation consultants and enterprise architects are most likely to encounter AI adoption demand constrained by compliance requirements, and where TPAI can unlock delivery.

Financial Services

Banks, insurance companies, and investment firms subject to GLBA, SOX, SEC disclosure rules, and NYDFS 23 NYCRR 500 face stringent third-party risk and data handling requirements. TPAI's architecture addresses core concerns in vendor due diligence assessments. Use cases include **research and analysis, regulatory document review, client communication drafting, and compliance monitoring**, within a security posture that supports institution assessments against applicable federal and state financial regulations.

Healthcare

Healthcare organizations processing PHI face HIPAA Privacy and Security Rule requirements that demand encryption, access controls, audit trails, and formal business associate agreements for any third-party platform. TPAI's zero-PII-retention architecture and per-user encryption make it viable for **clinical decision support, medical literature review, administrative automation, and training** where patient data confidentiality is non-negotiable. The architecture supports privacy risk committee review, data classification workflows, and audit evidence requirements that health-sector compliance teams expect from any AI vendor.

Legal

Law firms and in-house legal teams handle privileged communications, M&A documentation, and litigation strategy. In the US litigation environment, the discoverability of AI-processed data is an active and evolving area of concern. TPAI ensures client-matter data remains encrypted per user with no cross-tenant access, supporting matter-level confidentiality and privileged workflow boundaries. Use cases include **contract review, legal research, document summarization, and due diligence support**.

Public Sector & Government Contractors

Federal and state agencies are under direct mandate to expand AI capabilities while implementing governance safeguards. NIST SP 800-53 control baselines and FedRAMP authorization requirements create a compliance floor that generic public AI offerings typically cannot meet without significant additional controls, boundary design, and accreditation planning.

For systems integrators and government contractors delivering AI program, TPAI provides a sovereign deployment foundation: air-gapped, per-user encrypted, and deployable into the agency's own AWS account. Use cases for pilot and innovation program include **internal knowledge management, document analysis, policy drafting, and constituent-service automation**. The prime contractor manages accreditation and authority-to-operate processes; TPAI provides the underlying infrastructure and security evidence to support that process. Formal certifications (ISO 27001, SOC 2 Type II) are in progress.

5. What This Means for Your Practice

You have been in this room before. The AI strategy is compelling. The business case is solid. The executive sponsor is engaged. And then the question lands: *How do we keep our data sovereign?* The room shifts. Legal leans forward. The CISO raises a hand. And if you do not have a credible, specific answer, the engagement loses momentum in a way that is very difficult to recover.

Until now, the available options were to recommend that clients accept data exposure risk (untenable for regulated entities), build custom AI infrastructure (high cost, long timelines, ongoing maintenance), or defer AI adoption entirely (unacceptable given strategic mandates).

TPAI provides another option: a purpose-built platform deployable in weeks, operating within regional AWS infrastructure, designed around the compliance frameworks your clients are already subject to.

For advisory practices, this translates into shorter discovery-to-solution cycles, fewer architecture exception debates in steering committees, and a faster path from strategy deck to pilot. When the data sovereignty question has a credible answer, the entire engagement accelerates.

The advisor who can answer the sovereignty question in the room, with evidence, closes the engagement. The one who cannot watches it go to committee and quietly die.

Practically, this means:

- **Shorter sales cycles:** The compliance objection is addressed at the architecture level, not debated project by project.
- **Procurement-ready documentation:** Regional evidence packs covering KSA PDPL, UAE PDPL, NCA ECC-2, SAMA CSF, CBUAE, and DFSA requirements are available on request.
- **Flexible engagement models:** TPAI works with consultants as channel partners with referral, co-sell, and white-label options.
- **Credible client recommendation:** You can recommend a platform that satisfies both the transformation mandate and the compliance requirement, without hedging.

When TPAI Is Not the Right Fit

Not every AI use case requires air-gapped infrastructure. If your client has no data sovereignty requirement, processes no sensitive or regulated data, and operates small teams with non-confidential workloads, a standard public AI platform may be entirely adequate. TPAI is designed for organizations where the compliance, privacy, and litigation risk profile makes public AI untenable, and where the alternative of building from scratch is too slow and too expensive.

The question is not whether your clients will adopt AI. It is whether you will be the advisor who showed them how to do it safely.

6. How to Engage

TPAI is working with a select number of advisory partners during the beta program. If you advise regulated organizations on AI strategy or digital transformation, we would welcome a conversation.

- **All enquiries:** rcastillo@totallyprivate.ai
- **Evidence packs:** Available for US, EU, Middle East, and APAC jurisdictions
- **Web:** https://totallyprivate.ai

Your Data. Your AI. Totally Private.