Kiteworks

SOLUTION GUIDE

Shadow AI in the Kingdom of Saudi Arabia

How to Identify, Contain, and Govern the Kingdom's Most Consequential AI Security Risk



98%
of Saudi public servants use AI at work — the highest rate globally
Public Sector AI Adoption Index 2026

91%
of Saudi organizations experienced an AI-related security incident in the past year
Cisco Cybersecurity Readiness Index 2025

$670K
average Shadow AI uplift added to each data breach
IBM Cost of a Data Breach Report 2025

Saudi Arabia is the world's most AI-enthusiastic nation — and by direct consequence, one of the most exposed to Shadow AI: the unauthorized, ungoverned use of AI tools that creates regulatory violations, data breaches, and audit failures faster than any organization can manually track. This guide explains the risk, maps it to the Kingdom's specific regulatory obligations, and lays out the architecture required to govern it.

Executive Summary

Saudi Arabia ranks first globally in public sector AI adoption — 98% of public servants use AI at work, two-thirds every day.[1] The Kingdom declared 2026 the Year of Artificial Intelligence, has committed over $40 billion in AI investment through its Public Investment Fund, and has embedded AI as a structural requirement across 66 Vision 2030 goals.[2] That ambition is real, coordinated, and accelerating.

The governance infrastructure has not kept pace. Only 25% of Saudi organizations have achieved mature or progressive cybersecurity readiness.[3] 91% experienced an AI-related security incident in the prior year.[4] In agencies that have not provided approved AI tools, 64% of public servants use personal account logins at work and 70% use AI without their manager's knowledge.[5] The result is a governance gap that is simultaneously a data security risk, a regulatory compliance failure, and a financial liability.

At the center of that gap sits Shadow AI — the unauthorized, ungoverned use of AI tools by employees operating outside official IT and compliance frameworks. Shadow AI is not a peripheral risk in Saudi Arabia. It is the primary unresolved vulnerability running through an ecosystem shaped by some of the world's most prescriptive data protection regulations.

Key Findings

- Shadow AI is responsible for 20% of AI-related security breaches globally[6] and adds an average of $670,000 to each breach cost.[7]
- 97% of organizations that experienced AI-related breaches lacked adequate AI access controls.[8]
- 83% of organizations lack automated controls to prevent employees from uploading sensitive data to AI tools.[9]
- A single Shadow AI incident in Saudi Arabia simultaneously implicates seven overlapping regulatory frameworks: PDPL, SAMA CSF, NCA ECC, NCA DCC, NCA CCC, SDAIA AI Ethics Principles, and SDAIA GenAI Guidelines — creating multi-regulator exposure with no close parallel in any other major AI-adopting economy.
- SDAIA issued 48 enforcement decisions against organizations in 2025, signaling that PDPL enforcement is operational, not aspirational.[10]
- Third-party involvement in data breaches doubled year-over-year, from 15% to 30% of confirmed incidents, according to Verizon's 2025 DBIR.[11] Consumer AI tools are, by definition, unvetted third parties.
- The average annual cost of insider threats — with Shadow AI now the top driver — has reached $20 million.[12]

What Organizations Must Do

- Run a Shadow AI discovery audit immediately. Most organizations find dozens of unauthorized AI applications.[13]
- Enforce governance at the data layer — not the model layer. System prompts are not audit-defensible controls under NCA DCC or SAMA CSF.
- Map every AI data flow against PDPL cross-border transfer restrictions before deployment. Most consumer AI tools process on U.S. and European servers.
- Provide approved AI tools with proper governance. Restricting access without an alternative creates more Shadow AI, not less.[14]
- Govern AI agents as you would govern employees — authenticated per session, authorized per operation, logged with full attribution.

The Business Case

Organizations with AI-powered defenses save approximately $2 million per breach and resolve incidents 80 days faster.[15] 63% of organizations globally associate compliance investment with improved security posture.[16] The organizations that build AI governance infrastructure before regulators compel it will deploy AI faster, close deals more credibly, and respond to enforcement inquiries in hours rather than weeks.

The Problem: Adoption Without Governance

Saudi Arabia declared 2026 the Year of Artificial Intelligence. Vision 2030 ties 66 of its national transformation goals directly to AI and big data.[17] The Public Investment Fund has committed over $40 billion to AI ventures.[18] AI companies in the Kingdom attracted $9 billion in funding in 2025 alone. MENA IT spending is projected to reach $169 billion in 2026.[19]

The adoption results match the ambition. According to the Public Sector AI Adoption Index 2026 — a survey of 3,335 public servants across 10 countries — Saudi Arabia ranked first globally with a score of 66 out of 100.[20] India and Singapore, in second and third, sit eight points behind. 98% of Saudi public servants use AI at work. Two-thirds use it every day.[21]

But adoption statistics and governance statistics tell opposite stories. Only 25% of Saudi organizations have achieved mature or progressive cybersecurity readiness, according to Cisco's 2025 Cybersecurity Readiness Index.[22] 91% experienced an AI-related security incident in the past year.[23] In agencies that have not provided approved tools, 64% of public servants use personal account logins at work, and 70% use AI without their manager's knowledge.[24]

That gap between adoption velocity and governance maturity is where Shadow AI lives. And in a country with the world's most prescriptive multilayered data protection regime — the PDPL, SAMA's Cybersecurity Framework, NCA's ECC/DCC/CCC controls, and SDAIA's AI Ethics Principles — it is also where regulatory exposure accumulates silently, interaction by interaction, every day.[25]

The Core Risk in One Sentence

Every time a Saudi employee pastes work data into an unauthorized AI tool, they potentially create a PDPL cross-border transfer violation, a SAMA CSF control failure, an NCA DCC audit trail gap, and an SDAIA AI Ethics breach — simultaneously, in a single interaction, with no alert fired and no log entry created.

What Shadow AI Is — and Why It's Harder to Govern Than Shadow IT

Shadow AI is the use of AI tools — generative AI chatbots, code assistants, large language models, AI analytics platforms — without the knowledge, approval, or oversight of IT, security, or compliance functions.[26] It is related to Shadow IT but structurally more dangerous.

Factor	Shadow IT	Shadow AI
Visibility	Detectable through network logs	Blends invisibly into normal user workflows
Primary risk	Unauthorized apps	Data exfiltration, AI outputs, decision contamination
Adoption speed	Gradual, identifiable	Rapid, viral — invisible to IT
Governance approach	IT-driven	Requires legal, compliance, and security alignment
Regulatory impact	Moderate	High — PDPL, SAMA, NCA exposure simultaneously
Data exposure	Application access	Training model ingestion — potentially permanent

The most dangerous dimension is what happens to data after it leaves the organization. Information uploaded to public AI tools can become permanently embedded in third-party training models.[27] Unlike a shadow application that can be blocked or deprovisioned, data already ingested into an external AI system may be irrecoverable. You cannot fulfill a PDPL deletion obligation, respond to a subject access request, or reconstruct a breach notification timeline for data in a system you did not know existed.

The Five Shadow AI Behaviors Most Common in KSA Organizations

1. Employees pasting confidential work data into consumer AI interfaces (ChatGPT, Gemini, Copilot) via personal account credentials — bypassing all enterprise monitoring
2. Using personal-subscription AI tools for work tasks, with no IT visibility or policy enforcement
3. Deploying AI-enabled browser extensions or plugins without security assessment
4. Connecting approved work systems to unapproved AI tools via unofficial integrations
5. Using AI notetakers that retain full meeting recordings — including sensitive personnel, financial, or operational discussions — in third-party storage with no data governance

Why Training Alone Won't Solve It

A 2026 UpGuard finding is uncomfortable for any organization relying on awareness training as its primary Shadow AI control: employees who understand AI security requirements are more likely to use unauthorized tools, not less.[28] As AI knowledge increases, so does confidence in personal risk judgment — and willingness to override policy when productivity seems to justify it.

60% of employees globally accept security risks specifically to work faster with unsanctioned AI tools.[29] In Saudi Arabia, where 84% of public servants received employer-provided AI training but 50% describe it as box-ticking,[30] that dynamic plays out at national scale. Training without governance architecture is not a control. It is a documented liability.

The Scale of the Exposure

What the Data Shows Globally

- 83% of organizations lack automated controls to prevent employees from uploading sensitive data to AI tools[31]
- 86% of organizations are blind to their AI data flows; the average enterprise has 1,200 unofficial applications[32]
- Shadow AI was responsible for 20% of AI-related security breaches in 2025[33]
- 97% of organizations that experienced AI-related breaches lacked adequate AI access controls[34]
- 68% of organizations experienced data leaks directly linked to AI tool usage
- End-users are responsible for 68% of errors leading to breaches[35] — in a workforce where AI is used daily and training is superficial, that is an operating baseline, not a benchmark

Three Reasons KSA Exposure Is Worse Than the Global Average

1. Highest daily AI use rate in the world[36] — highest frequency of potential exposure events per day, per organization
2. Strictest data localization requirements of any major AI-adopting economy — PDPL requires personal data on Saudi citizens to remain within the Kingdom;[37] most consumer AI tools process on U.S. and European servers, making every unauthorized use a potential automatic violation
3. Only 25% of Saudi organizations have achieved mature or progressive cybersecurity readiness[38] — three-quarters of the market lacks controls to detect, contain, or respond to Shadow AI incidents

The Verizon 2025 Data Breach Investigations Report — 22,052 security incidents, 12,195 confirmed breaches, the largest dataset in a single DBIR edition[39] — specifically flagged GenAI-related data leakage as a new internal data egress channel, driven by employees accessing AI via non-corporate identities without integrated authentication. Third-party involvement in breaches doubled year-over-year, from 15% to 30% of confirmed incidents.[40] Consumer AI tools are unvetted third parties. The connection is direct.

The Financial Exposure

$670K
Shadow AI average breach cost uplift per incident
IBM Cost of a Data Breach Report 2025

$20M
Average annual insider threat cost; shadow AI is the top driver
DTEX / Ponemon 2026

73 days
Median third-party breach disclosure lag — while the PDPL notification clock runs
Black Kite 2026

Organizations with AI-powered defenses saved approximately $2 million per breach and resolved incidents 80 days faster, according to the same IBM/Ponemon research.[41] The gap between governed and ungoverned AI is not only regulatory. It is directly financial.

What Data Is Actually at Risk

A Kiteworks AI Data Security Report found that 27% of organizations admit more than 30% of the information sent to AI tools contains private data — national ID equivalents, medical records, financial data, trade secrets.[42] In the KSA context, that maps to: citizen national ID and PII, SAMA-regulated banking and credit records, government procurement and contract data, healthcare records under PDPL's high-risk processing provisions, and energy sector operational data tied to critical national infrastructure.

Regulatory Exposure: Seven Frameworks, One Incident

Most jurisdictions have one primary privacy law. Saudi Arabia has seven overlapping regulatory frameworks governing the same data simultaneously. A single Shadow AI incident triggers multi-regulator exposure that has no close parallel in any other major AI-adopting economy.

Framework	Shadow AI Trigger	Controls Violated	Enforcer
PDPL (SDAIA)	Cross-border data transfer; processing without consent	Consent, localization, DPIAs, breach notification	SDAIA
SAMA CSF	Financial data in unauthorized AI tool	Encryption, access control, third-party oversight	SAMA
NCA ECC-1:2023	Ungoverned third-party AI with CNI access	Third-party security; cloud computing controls	NCA
NCA DCC-1:2022	Data lifecycle violations; no audit trail	Data classification, protection, audit logging	NCA
NCA CCC-2:2024	Non-NCA-approved cloud AI service	Data localization; CST/CSP approval required	NCA / CST
SDAIA AI Ethics Principles	No transparency, accountability, privacy controls	Principles 3 (Transparency), 4 (Accountability), 5 (Privacy)	SDAIA
SDAIA GenAI Guidelines 2025	Classified data in consumer GenAI tool	Internal policy; privacy impact assessments	SDAIA

SDAIA and the PDPL: Enforcement Is No Longer Theoretical

SDAIA issued 48 enforcement decisions against organizations in 2025 — the first substantive wave since the PDPL became enforceable in September 2023.[43] Common violations: processing personal data without lawful basis, collecting data beyond stated purposes, insufficient technical controls. Every unauthorized employee AI interaction potentially checks all three boxes.

SDAIA also achieved ISO 42001 certification in July 2024 — the international standard for AI management systems.[44] ISO 42001 compliance is becoming the baseline expectation for procurement and regulatory processes across the Kingdom. Organizations without it will face increasing friction from both regulators and enterprise customers.

SAMA: Financial Sector Carries Double Exposure

SAMA-regulated institutions — banks, insurance companies, fintechs, and their third-party vendors — face the most prescriptive Shadow AI exposure in the Kingdom. SAMA's dedicated AI regulatory framework (2023) requires explainable AI-driven financial decisions, bias audits, documented risk assessments, and human accountability for AI outcomes.[45] Shadow AI tools satisfy none of these requirements by definition.

The dual-regulator structure — SAMA for banking and payments, the Capital Market Authority for securities — means a single Shadow AI incident can trigger parallel investigations from two agencies with different standards, timelines, and penalty frameworks.[46] Globally, 90% of breaches affecting financial institutions carry financial motives.[47]

The Incoming Complexity: The Draft Global AI Hub Law

The CST released its draft Global AI Hub Law in April 2025 — the first G20 proposal allowing foreign governments to operate on Saudi soil under their own data protection law, through three hub models including a "data embassy" concept with no direct precedent.[48] Once enacted, Shadow AI that routes data through foreign AI services outside Hub agreements creates new jurisdictional exposure categories that current governance frameworks are not designed to catch.

Where Shadow AI Hits Hardest by Sector



Government and Public Sector

With 98% AI adoption and two-thirds using AI daily,[49] Saudi government is the highest-volume Shadow AI risk environment in the world. The data moving through unauthorized tools includes citizen PII, tax records, and law enforcement files — categories subject to PDPL's most restrictive protections. SDAIA's AI Ethics Principles require government entities to designate AI Ethics Officers and subject AI systems to periodic verification.[50] Shadow AI deployments are never reviewed under those mandates. Cumulative, unaudited exposure accrues daily.



Financial Services

AI is embedded in fraud detection, credit scoring, and trading across Saudi financial institutions. The attack surface from unauthorized tools is acute. Globally, financial data breaches average $6 million per incident.[51] Dual-regulator exposure — SAMA plus CMA — means a single Shadow AI incident can trigger parallel investigations from two agencies with different standards and timelines.[52]



Energy and Critical Infrastructure

Vision 2030's energy diversification has accelerated OT/IT convergence across Saudi Arabia's energy sector. The Dragos 2026 OT/ICS Cybersecurity Report documented 96 ransomware incidents against Middle East OT/ICS environments.[53] Shadow AI agents operating in IT environments adjacent to operational systems create unmonitored access paths that neither OT security nor traditional IT controls are configured to detect.



Healthcare

Health data is among the PDPL's most sensitive categories — high-risk processing that explicitly requires a Data Protection Impact Assessment before any deployment.[54] Shadow AI processing patient records is, by definition, high-risk processing with no DPIA conducted, no consent obtained, no localization assessed, and no audit trail. Vision 2030's health sector digitization is driving AI adoption at pace; governance has not kept up.

The Next Wave: Shadow AI Agents

The Shadow AI problem described so far — employees pasting data into consumer AI tools — is serious. In 2026, it is also escalating structurally into something harder to govern.

The shift is from Shadow AI tools to Shadow AI agents: systems that set objectives, execute multi-step task sequences, call APIs, access databases, and interact with external services continuously — without human prompting for each action. An AI agent is not a tool a person uses. It is a digital worker operating with accumulated access and persistent autonomy.

The DTEX/Ponemon 2026 Insider Threat Report found that 19% of organizations have deployed AI agents into daily workflows — while only 19% classify AI agents as equivalent to human insiders for governance purposes. 44% believe malicious use of AI agents will significantly increase data theft risk. 73% worry that unauthorized AI use is creating invisible data loss pathways.[55] The technology has outrun the policy framework designed to govern it.

The Saudi Threat Context

CrowdStrike 2026: 89% year-over-year increase in AI-enabled adversary attacks.

IBM X-Force 2026: 44% increase in attacks beginning with exploitation of public-facing applications — including AI interfaces and agent endpoints.

Average eCrime breakout time: 29 minutes from initial access to lateral movement.

For Saudi organizations already operating unmonitored Shadow AI agents, an attacker who compromises one inherits whatever access it has accumulated. 82% of 2026 detections were malware-free — attackers using legitimate tools and legitimate credentials. A Shadow AI agent in the wild is indistinguishable from normal activity.

Standard DLP was designed for humans sending files. It cannot evaluate AI agent API calls, MCP tool invocations, or multi-step data workflows. It cannot authenticate AI agent identity, enforce minimum-necessary access at the operation level, or produce the delegation-chain audit trail that regulators require.[56] Shadow AI agents operating without governance are indistinguishable from legitimate activity in most security stacks. There is no alert to fire.

The Solution: Five Things Organizations Must Do

The regulatory architecture exists. SDAIA, SAMA, and NCA have constructed frameworks that, when operationally implemented, eliminate most Shadow AI compliance exposure. The missing element is the operational data governance infrastructure to enforce those requirements at machine speed — across millions of daily AI interactions — without depending on individual employee judgment.

1

Run a Shadow AI Discovery Audit — Now

You cannot govern what you cannot see. A Shadow AI audit using network monitoring, SaaS discovery, and DSPM tooling will typically surface dozens of unauthorized AI applications in a mid-sized organization.[57] The average enterprise has 1,200 unofficial applications.[58] Most organizations are surprised by what their audit finds.

This is not a one-time exercise. Shadow AI discovery must be continuous — the tool landscape changes weekly, and employee behavior adapts to whatever restrictions are imposed. Discovery is the prerequisite for every other action on this list.

2

Enforce Governance at the Data Layer — Not the Model Layer

System prompts and model-level guardrails are not audit-defensible controls under NCA DCC, SAMA CSF, or SDAIA's enforcement framework. A SDAIA investigator will not accept a system prompt as evidence that personal data access was authorized. Governance must be enforced at the data layer: attribute-based access control, FIPS 140-3 validated encryption, and tamper-evident audit logs for every AI interaction.[59] This is the only architecture that satisfies all seven KSA frameworks simultaneously — because all seven ultimately require provable control over data, not policy statements about AI behavior.

3

Map Every AI Data Flow Against PDPL Localization Requirements

Before any AI tool deployment, three questions must be answered and documented: Does this tool process personal data of Saudi citizens or residents? Where does processing occur geographically? Is the processing environment within an NCA CCC-certified cloud or an on-premises deployment within KSA?[60,61]

Most consumer AI tools process on U.S. and European servers. The PDPL's cross-border transfer restrictions make that a potential automatic violation with every employee interaction.[62] Data flow mapping determines which tools are permissible to deploy at all — and the 94-day median remediation time for exposed credentials means delays in assessment directly affect the organization's ability to meet notification obligations.

4

Provide Approved Tools — or Accept More Shadow AI

Restricting AI access without providing approved alternatives does not reduce Shadow AI. It drives usage underground. The data proves this directly: in agencies without approved tools, 64% of public servants use personal logins and 70% use AI without manager knowledge.[63] Secure enablement — governed AI tools with proper data controls, audit logging, and compliance guardrails — is the only approach that simultaneously reduces Shadow AI risk and retains productivity gains.

5

Govern AI Agents as You Would Govern Employees

Every AI agent deployed without governance is a non-human identity with persistent data access and no audit trail.[64] AI agents must be authenticated per session, authorized per operation, and logged with full attribution — what data, which agent, which human authorizer, when, and for what purpose. That is the application of zero-trust principles extended to non-human identities — the only governance model that produces audit-defensible evidence for SDAIA, SAMA, and NCA.

The Governance Architecture: What It Must Include

The table below maps each KSA regulatory requirement to the specific technical capability required to satisfy it. An organization that implements all seven elements has a defensible governance position across the full regulatory stack.

Architecture Requirement	What It Does	Regulatory Frameworks Satisfied
Zero-trust AI data access	Every AI request authenticated and authorized before data is provided — no implicit access once connected	PDPL cross-border transfer; SAMA CSF access control; NCA DCC audit trail
ABAC policy enforcement at the data layer	AI inherits user permissions and cannot exceed authorization boundaries; enforced per operation, not per session	SAMA CSF; NCA DCC; SDAIA AI Ethics Principle 5
FIPS 140-3 validated encryption in-jurisdiction	Data encrypted to federal audit standard; encryption keys retained within KSA	PDPL data localization; SAMA CSF encryption; NCA CCC
Tamper-evident, immutable audit logs	Every AI interaction logged with full attribution — data, agent, user, time, authorization	NCA DCC audit trail; SDAIA AI Ethics Principle 4; SAMA CSF
Real-time SIEM integration	AI access events fed to security operations immediately — no delayed or batched logging	NCA ECC threat detection; SAMA CSF monitoring
ISO 42001-aligned AI governance structure	AI Ethics Officers designated; AI systems subject to periodic verification; aligned to standard SDAIA has itself adopted	SDAIA AI Ethics Principles; SDAIA GenAI Guidelines 2025
On-premises or private cloud deployment	Processing within KSA-authorized cloud environments; no data routed through unapproved foreign AI providers	PDPL localization; NCA CCC; SAMA CSF third-party oversight

What Kiteworks Delivers

Kiteworks provides the governed data layer that enterprise AI requires in regulated environments. Its Private Data Network delivers:

- Zero-trust AI data access via the Kiteworks AI Data Gateway — every AI request authenticated, authorized, and logged before data is provided
- Compliant AI (MCP Server + AI Data Gateway) — zero-trust governance extended to AI agents and RAG pipelines with OAuth 2.0 authentication and ABAC policy enforcement at the data layer
- FIPS 140-3 validated encryption with in-jurisdiction key custody — satisfying PDPL localization and NCA CCC requirements
- Tamper-evident, immutable audit logs for every AI interaction, feeding real-time to SIEM — satisfying NCA DCC, SAMA CSF, and SDAIA AI Ethics Principle 4
- Flexible deployment: on-premises, private cloud, hybrid — data stays within KSA-authorized environments
- Automated compliance reporting with exportable evidence packages for SDAIA, SAMA, and NCA audit requirements

Kiteworks protects over 100 million end-users and thousands of global enterprises and government agencies.

kiteworks.com/platform/compliance/compliant-ai

The Business Case for Acting Now

IBM/Ponemon data shows that organizations with AI-powered defenses save $2 million per breach and resolve incidents 80 days faster.[65] 63% of organizations globally associate compliance investment with improved security posture. 52% report enhanced customer trust as a direct outcome.[66] Those are business outcomes, not regulatory checkboxes.

Saudi Arabia's $9 billion AI investment in 2025[67] is building AI capability across every sector of the economy. The organizations that govern that capability first will deploy AI faster — because compliance review is automated and continuous rather than manual and periodic. They will close deals more credibly — because they can produce evidence packages on demand. And when a SDAIA or SAMA inquiry arrives, they will answer it in hours, not weeks.

48

SDAIA enforcement decisions issued in 2025 — enforcement is operational, not aspirational

IAPP, 2025

30%

of breaches now involve third-party access — up from 15% the prior year

Verizon DBIR 2025

$2 M

average breach cost savings for organizations with AI-powered defenses

IBM Cost of a Data Breach 2025

Two Paths Forward

There are two ways a Saudi organization discovers its Shadow AI exposure.

The first: a SDAIA enforcement investigator contacts you citing unauthorized personal data processing. Your legal team scrambles to reconstruct what happened. IT discovers multiple teams have been using consumer AI tools for months. The data trail is incomplete. The PDPL breach notification clock may already be running — you don't know yet, because you don't know when the first unauthorized transfer occurred. The SAMA inquiry arrives the following week.

The second: your security team runs a Shadow AI audit, surfaces and closes the exposure, and builds a governed AI program with approved tools, data-layer policy enforcement, and tamper-evident audit logs. A SDAIA inquiry arrives the following year. Your team produces the evidence package in four hours.

Those paths diverge at exactly one decision point: whether to build governance infrastructure before the regulator makes it a legal obligation.

Saudi Arabia has proven it can achieve AI adoption at national scale.[68] The governance challenge is categorically different — it requires protection that operates at machine speed, across millions of daily interactions, without creating friction that drives AI use underground. The regulatory scaffolding is in place. The missing element is the operational infrastructure to enforce it.

Organizations that build that infrastructure now will not merely achieve regulatory compliance. They will have the trusted AI foundation that Vision 2030's knowledge economy requires.

Shadow AI Governance Checklist for KSA Organizations

Use this checklist to assess your current Shadow AI governance posture across all relevant regulatory frameworks.

Discovery & Visibility	
☐	Comprehensive AI tool audit completed across all departments in the past 90 days
☐	Continuous SaaS discovery monitoring identifies unauthorized AI applications
☐	AI-enabled browser extensions and plugins reviewed on managed devices
☐	AI notetakers and meeting-summary tools inventoried and assessed for data retention
PDPL Compliance	
☐	Every AI tool assessed for whether it processes personal data of Saudi citizens or residents
☐	Geographic location of processing documented for each AI tool
☐	Cross-border transfer adequacy assessments completed before any deployment outside KSA
☐	DPIAs conducted for AI tools categorized as high-risk processors
☐	Breach notification process can meet PDPL timelines
SAMA CSF (Financial Sector)	
☐	Third-party AI tools in financial operations subject to SAMA CSF vendor risk assessments
☐	FIPS 140-3 encryption enforced for all AI data interactions
☐	AI-driven financial decisions subject to explainability and human oversight requirements
NCA Controls	
☐	AI tools verified as operating within NCA CCC-approved cloud environments
☐	Tamper-evident audit trail exists for all AI data access satisfying NCA DCC requirements
☐	Third-party AI vendors assessed under the same process as other third-party providers (NCA ECC)

SDAIA AI Ethics & GenAI Guidelines	
☐	AI Ethics Officer designated as required by SDAIA AI Ethics Principles
☐	Internal policy explicitly prohibits classified data entry into GenAI systems
☐	AI systems subject to periodic verification and validation
☐	AI-generated content is labeled to notify end-users
Governance Architecture	
☐	AI access governed at the data layer — not via system prompts or model-level guardrails
☐	Zero-trust authentication enforced for AI agent data access
☐	AI governance structures ISO 42001-aligned
☐	Real-time SIEM integration in place for AI access logs

Glossary of Key Regulatory Acronyms

Acronym	Definition
ABAC	Attribute-Based Access Control — grants access based on attributes of users, data, environment, and operations
CCC	Cloud Cybersecurity Controls — NCA framework for cloud security and data localization requirements
CST	Communications, Space and Technology Commission — KSA regulator for telecommunications, IT, and the draft Global AI Hub Law
DCC	Data Cybersecurity Controls — NCA framework for data lifecycle protection and audit trail requirements
DPIA	Data Protection Impact Assessment — mandatory for high-risk AI processing under the PDPL
ECC	Essential Cybersecurity Controls — NCA foundational framework including third-party and cloud controls
FIPS 140-3	Federal Information Processing Standard for cryptographic module validation — required for audit-defensible encryption
ISO 42001	International standard for AI management systems — adopted by SDAIA July 2024
MCP	Model Context Protocol — open standard for AI system integration with enterprise data repositories
NCA	National Cybersecurity Authority — KSA's apex cybersecurity regulator
PDPL	Personal Data Protection Law — KSA's foundational data privacy law, enforceable from September 2023
PIF	Public Investment Fund — Saudi Arabia's sovereign wealth fund; committed over $40B to AI ventures
SAMA	Saudi Arabian Monetary Authority — central bank and financial sector cybersecurity regulator
SDAIA	Saudi Data and Artificial Intelligence Authority — apex authority for AI governance and PDPL enforcement

Endnotes

All statistics and regulatory citations in this guide are sourced from the documents listed below. Kiteworks proprietary research is identified as such; all other statistics are attributed to the originating third-party publisher.

1. Kiteworks. "Saudi Arabia's AI Leadership: Securing Public Sector Innovation." kiteworks.com (citing Public Sector AI Adoption Index 2026, PublicFirst).
2. Grant Thornton Saudi Arabia. "Saudi Arabia's Year of Artificial Intelligence 2026." grantthornton.sa.
3. Cisco. "KSA Findings from its Cybersecurity Readiness Index 2025." tradingview.com/news.
4. Cisco. "KSA Findings from its Cybersecurity Readiness Index 2025." tradingview.com/news.
5. Kiteworks. "Shadow AI Crisis: Solving Government's Hidden Security Risk." kiteworks.com (citing Public Sector AI Adoption Index 2026).
6. HelpNetSecurity. "Shadow AI, Deepfakes, and Supply Chain Compromise – Financial Sector Cyber Threats Report." helpnetsecurity.com, April 2026.
7. IBM / Ponemon Institute. Cost of a Data Breach Report 2025. ibm.com/reports/data-breach.
8. HelpNetSecurity. "Shadow AI, Deepfakes, and Supply Chain Compromise – Financial Sector Cyber Threats Report." helpnetsecurity.com, April 2026.
9. Kiteworks. "AI Data Security Report: 83% of Organizations Lack Basic Controls." kiteworks.com/video-ai-data-security-report.
10. IAPP. "Saudi Arabia's Data Protection Authority Steps Up Enforcement." iapp.org.
11. Verizon. 2025 Data Breach Investigations Report. verizon.com/business/resources/reports/dbir.
12. DTEX Systems / Ponemon Institute. 2026 Insider Threat Report. dtex.com.
13. Kiteworks. "AI Data Security Report: 83% of Organizations Lack Basic Controls." kiteworks.com/video-ai-data-security-report.
14. Kiteworks. "Shadow AI Crisis: Solving Government's Hidden Security Risk." kiteworks.com (citing Public Sector AI Adoption Index 2026).
15. IBM / Ponemon Institute. Cost of a Data Breach Report 2025. ibm.com/reports/data-breach.
16. Kiteworks. 2026 Data Security, Compliance & Risk: Data Sovereignty Report. kiteworks.com.
17. Grant Thornton Saudi Arabia. "Saudi Arabia's Year of Artificial Intelligence 2026." grantthornton.sa.
18. Fast Company Middle East. "Saudi Arabia data center sector surges under Vision 2030 with $4.26 billion investment." fastcompanyme.com.
19. Gartner / YouTube. "Middle East to Pump $169 Billion into AI Compute Boom." youtube.com.
20. PublicFirst. Public Sector AI Adoption Index 2026. indexaiglobalpublicservices.publicfirst.co.
21. Kiteworks. "Saudi Arabia's AI Leadership: Securing Public Sector Innovation." kiteworks.com (citing Public Sector AI Adoption Index 2026, PublicFirst).
22. Cisco. "KSA Findings from its Cybersecurity Readiness Index 2025." tradingview.com/news.
23. Cisco. "KSA Findings from its Cybersecurity Readiness Index 2025." tradingview.com/news.
24. Kiteworks. "Shadow AI Crisis: Solving Government's Hidden Security Risk." kiteworks.com (citing Public Sector AI Adoption Index 2026).
25. V-Comply. "Saudi Data & AI Authority (SDAIA) Compliance." v-comply.com.
26. IBM. "What Is Shadow AI?" IBM Think Topics. ibm.com/think/topics/shadow-ai.
27. Kiteworks. "AI Data Security Report: 83% of Organizations Lack Basic Controls." kiteworks.com/video-ai-data-security-report.
28. UpGuard. "Shadow AI Use Widespread, Poses Security Risks." linkedin.com/posts/thomaslafayette.
29. BlackFog. "Shadow AI Threat Grows Inside Enterprises." blackfog.com.
30. PublicFirst / Kiteworks. KSA Country Factsheet – Public Sector AI Adoption Index 2026. kiteworks.com.
31. Kiteworks. "AI Data Security Report: 83% of Organizations Lack Basic Controls." kiteworks.com/video-ai-data-security-report.
32. Kiteworks. "AI Data Security Report: 83% of Organizations Lack Basic Controls." kiteworks.com/video-ai-data-security-report.
33. HelpNetSecurity. "Shadow AI, Deepfakes, and Supply Chain Compromise – Financial Sector Cyber Threats Report." helpnetsecurity.com, April 2026.
34. HelpNetSecurity. "Shadow AI, Deepfakes, and Supply Chain Compromise – Financial Sector Cyber Threats Report." helpnetsecurity.com, April 2026.
35. Verizon. 2025 Data Breach Investigations Report. verizon.com/business/resources/reports/dbir.
36. PublicFirst. Public Sector AI Adoption Index 2026. indexaiglobalpublicservices.publicfirst.co.
37. Sahl. "Saudi PDPL Compliance Using AI: Roadmap for 2025." getsahl.io.
38. Cisco. "KSA Findings from its Cybersecurity Readiness Index 2025." tradingview.com/news.
39. Verizon. 2025 Data Breach Investigations Report. verizon.com/business/resources/reports/dbir.
40. Verizon. 2025 Data Breach Investigations Report. verizon.com/business/resources/reports/dbir.
41. IBM / Ponemon Institute. Cost of a Data Breach Report 2025. ibm.com/reports/data-breach.
42. Kiteworks. "AI Data Security Report: 83% of Organizations Lack Basic Controls." kiteworks.com/video-ai-data-security-report.
43. IAPP. "Saudi Arabia's Data Protection Authority Steps Up Enforcement." iapp.org.
44. Modulos AI. "Navigating AI Governance in Saudi Arabia." modulos.ai.
45. ITButler. "SAMA Compliance for Artificial Intelligence (AI)." itbutler.sa.

Endnotes (cont.)

46. Legal 500. "Country Comparative Guides 2026 – Saudi Arabia Fintech." legal500.com.

47. HelpNetSecurity. "Shadow AI, Deepfakes, and Supply Chain Compromise – Financial Sector Cyber Threats Report." helpnetsecurity.com, April 2026.

48. JD Supra. "Saudi Arabia's Global AI Hub Law." jdsupra.com.

49. Kiteworks. "Saudi Arabia's AI Leadership: Securing Public Sector Innovation." kiteworks.com (citing Public Sector AI Adoption Index 2026, PublicFirst).

50. Morrison Foerster. "AI Library Saudi Arabia – SDAIA AI Ethics Principles." mofo.com.

51. IBM / Ponemon Institute. Cost of a Data Breach Report 2025. ibm.com/reports/data-breach.

52. Legal 500. "Country Comparative Guides 2026 – Saudi Arabia Fintech." legal500.com.

53. Dragos. 2026 OT/ICS Cybersecurity Report. dragos.com.

54. Sahl. "Saudi PDPL Compliance Using AI: Roadmap for 2025." getsahl.io.

55. DTEX Systems / Ponemon Institute. 2026 Insider Threat Report. dtex.com.

56. Kiteworks. "Compliant AI: Secure Data Governance for AI Agents." kiteworks.com/platform/compliance/compliant-ai.

57. LinkedIn / Sumardi. "7 Ways Kiteworks Boosts Your DSPM Investment." linkedin.com.

58. Kiteworks. "AI Data Security Report: 83% of Organizations Lack Basic Controls." kiteworks.com/video-ai-data-security-report.

59. Kiteworks. "Compliant AI: Secure Data Governance for AI Agents." kiteworks.com/platform/compliance/compliant-ai.

60. Sahl. "Saudi PDPL Compliance Using AI: Roadmap for 2025." getsahl.io.

61. NCA. "Cloud Cybersecurity Controls CCC-2:2024." nca.gov.sa.

62. Crowell. "The Middle East's Big Bet on AI and Data Security." crowell.com.

63. Kiteworks. "Shadow AI Crisis: Solving Government's Hidden Security Risk." kiteworks.com (citing Public Sector AI Adoption Index 2026).

64. DTEX Systems / Ponemon Institute. 2026 Insider Threat Report. dtex.com.

65. IBM / Ponemon Institute. Cost of a Data Breach Report 2025. ibm.com/reports/data-breach.

66. Kiteworks. 2026 Data Security, Compliance & Risk: Data Sovereignty Report. kiteworks.com.

67. Fast Company Middle East. "Saudi Arabia data center sector surges under Vision 2030 with $4.26 billion investment." fastcompanyme.com.

68. Kiteworks. "Saudi Arabia's AI Leadership: Securing Public Sector Innovation." kiteworks.com (citing Public Sector AI Adoption Index 2026, PublicFirst).

Copyright © 2026 Kiteworks. Kiteworks' mission is to empower organizations to effectively manage risk in every send, share, receive, and use of private data. The Kiteworks platform provides customers with a secure data exchange that delivers data governance, compliance, and protection in a unified control plane. Kiteworks unifies, tracks, controls, and secures sensitive data moving within, into, and out of their organization, significantly improving risk management and ensuring regulatory compliance on all private data exchanges. Headquartered in Silicon Valley, Kiteworks protects over 100 million end-users and thousands of global enterprises and government agencies.

www.kiteworks.com

June 2026