

FOR CIOs, CISOs, CTOS, AND DPOs

THE BUSINESS CASE FOR TRUSTED, GOVERNED AI

Eggsplain Delivers The Intelligence. Kiteworks Protects The Information. Together They Provide A Secure Foundation For Enterprise AI: From Cloud To Sovereign To On-Premises.

Sovereign AI

GDPR · NIS2 · AI Act

Hybrid to On-Prem

IN PARTNERSHIP WITH EGGSPLAIN

Purpose & Executive Summary

This playbook helps partners lead executive-level business conversations, not product demonstrations, on AI governance, data sovereignty, and secure digital transformation, moving the discussion from technology to business value. AI is no longer a future initiative; it is already part of daily operations. The challenge is no longer deploying AI; it's controlling it. Today's executive conversations center on governance, compliance, data sovereignty, security, predictable AI economics, business risk, and trust.

The Challenge

Most organizations already have AI; the challenge is no longer deploying it, it's controlling it. Executive teams are asking who can access sensitive data and AI agents, where information is actually processed, whether the organization can demonstrate GDPR and NIS2 compliance, and whether leadership can prove governance and oversight to the board.

The Solution

Eggsplain delivers the intelligence: a sovereign AI platform with orchestration, agents, applications, and private models that runs hybrid, on-premises, or in the cloud. Kiteworks protects the information: a control plane for secure data exchange with a central policy engine, encryption, and a complete audit trail across every channel. Together, the two platforms give organizations trusted, governed, secure AI.

Proven At Enterprise Scale

100M+

Protected End Users

100%

On-Premises, Private AI Infrastructure

99.9%

Uptime SLA

Executive Challenges



AI governance

How do we control what employees and AI agents can access?



Data sovereignty

Where is our information processed, and who controls it?



Compliance

Can we demonstrate compliance with GDPR, NIS2, and future AI regulations?



Risk management

How do we reduce cyber risk while increasing AI adoption?



Cost predictability

Can we scale AI without unpredictable operational costs?



Executive accountability

Can management demonstrate governance and oversight?

Better Together: Eggsplain, Protected By Kiteworks

Eggsplain Provides
A sovereign AI platform for the enterprise.

✓ Sovereign AI platform

✓ AI orchestration and agents

✓ AI applications and private models

✓ AI App Store

✓ Hybrid and on-premises deployment

Kiteworks Provides
The control plane that protects it.

✓ Secure data exchange and managed file transfer

✓ Secure email and API governance

✓ Central policy engine

✓ Complete audit trail

✓ Encryption and compliance reporting

Customer Business Outcomes

Reduced operational risk

Improved regulatory compliance

Better executive oversight

Increased customer trust

Greater control over sensitive information

Faster, more confident AI adoption

Lower governance complexity

Reduced dependency on hyperscalers

Durable As AI Changes

GDPR, NIS2, and the AI Act are the current test, not the last one. Governance sits in the control plane, not in any single model.

NEW MODELS
Added or swapped without changing who may reach which data.

AUTONOMOUS AGENTS
Governed as identities under the same rule set as people.

FUTURE REGULATION
Evidence already captured, so new reporting is a query.

Customer Journey

1

Workshop
Strategy & governance objectives

2

Assessment
AI landscape, data flows, risk

3

Proof of Value
Governance & policy enforcement

4

Deployment
Eggsplain, Kiteworks, secure AI

5

Expansion
New apps, departments, App Store

Typical Customer Challenges

CHALLENGE	WHY THIS MATTERS: EGGSPLAIN + KITeworks
AI governance	Unified AI platform with policy control
GDPR compliance	Secure data handling & audit trails
NIS2 readiness	Central governance & risk management
Data sovereignty	Sovereign & on-prem deployment
AI security / shadow AI	Protected, controlled AI platform
Executive oversight	Centralized visibility & reporting

Discovery Questions

• What is your AI strategy?

• Where does your sensitive data reside today?

• How do you control AI access?

• Are you prepared for NIS2 and GDPR?

"AI adoption isn't limited by technology anymore. It's limited by trust, governance, and control. That's exactly where Eggsplain and Kiteworks create value together."

Eggsplain + Kiteworks · Executive Business Case Playbook

One Request, Whatever The Sector

The same everyday moment in three industries: someone asks for sensitive information, and an AI model answers. No new logins and no second rulebook — the permissions the employee already holds decide what the model is allowed to see.

Identity You Already Maintain Becomes The AI Permission Model

Card or token, role and title, department, unit, and group membership are read from the directory in place today. The control plane applies that same set of rights to data and AI access, so nobody has to maintain a parallel entitlement model for models and agents.

	HEALTHCARE	FINANCE	DEFENCE
01 Who asks, and for what	A nurse asks the AI assistant to summarise a patient history before a consultation.	A relationship manager asks how exposed the bank is to one client across loans, trading, and derivatives.	A design engineer asks for the tolerance data on a part in a live weapons programme.
02 How they are identified	Signed in with her SITHS card; role, title, and department come from the existing catalogue.	Signed in with the bank's SSO; her desk, coverage list, and licence status come from the HR and trading systems.	Signed in with his smart card; clearance, nationality, and programme assignment come from the personnel directory.
03 What decides access	The record entitlements she already holds. A colleague in billing asking the same question gets nothing.	Only the client accounts on her coverage list, and only the exposure figures her mandate covers.	Clearance level and need-to-know for that specific programme, enforced document by document.
04 Policy before processing	Data class and consent checked; identifiers can be minimised before the prompt is built.	Advisory-side deal data is walled off, so the model never sees inside information before it is public.	Export-control marking is checked first: an ITAR-marked drawing is not summarised for a non-national, however senior.
05 Where it runs	A private model inside the hospital perimeter. Patient data does not leave and is not used for training.	A private model in the bank's own data centre. Positions and client names never reach an external API.	The model runs on an air-gapped network inside the facility. Nothing traverses the internet.
06 Evidence out	Who asked, which record, which model, what was returned — available to the data protection officer.	Question, accounts touched, and answer are kept together, so compliance can reconstruct who saw what before a trade.	Every access is logged against the programme, ready for a security audit or an export-licence review.

When The Requester Is An Agent

An automated agent gets its own named identity and scope and passes the same six steps. It cannot inherit more access than the person or service it acts for.

Why It Reads The Same Everywhere

Only the identity source and the rule set change by sector. The control points — access, policy, processing, evidence — are identical, which is what makes the pattern portable.

How The Two Platforms Fit Together

REFERENCE TOPOLOGY

Every request, file, and prompt passes through one governed control plane before it reaches AI. The intelligence runs inside the customer's own perimeter; the evidence comes out the other side.



01 People & Systems

Employees, customers, partners, suppliers, and third-party systems sit outside the ring.

02 Channels

Secure email, managed file transfer, file sharing, web forms, APIs, and agents.

03 Control Plane · Kiteworks

Policy, identity, encryption, inspection, and audit apply before anything reaches AI.

04 Sovereign Core · Eggsplain

Orchestration, agents, AI apps, and private models, on-premises or hybrid.

05 Evidence Out

GDPR and NIS2 reporting, AI Act evidence, board, and auditor oversight.

Ready to build the business case for trusted, governed AI?

www.kiteworks.com

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.