CAVALRYAI
INNOVATION MEETS SECURITY

CAVALRYAI INSIGHTS

AI Changes the Compliance Equation:

What CIOs and CISOs Need to Know About GDPR, HIPAA and SOC 2

As AI moves deeper into enterprise workflows, protecting data is no longer enough. Organizations need to govern who can use it, which models can process it, where it can travel - and prove those controls are working.

By David Whitford, CTO, CavalryAI

The Compliance Boundary Has Moved

For years, enterprise security leaders built controls around a relatively understandable world. They knew where sensitive information lived, which applications processed it, who could access it, and how it was retained, monitored and audited.

Then AI changed the equation. An employee can now take information from a controlled system, place it into an AI interface and - in seconds - introduce another model, another processor, another data flow and potentially another jurisdiction into what had previously been a well-understood security architecture.

The underlying compliance principles have not disappeared. But the number of places where organizations must enforce them has multiplied.

If your organization cannot control which data reaches which AI model, can it truly say it controls the data?

GDPR: AI Doesn't Make Personal Data Less Personal

The European Data Protection Board has clarified that existing GDPR principles continue to apply to the development and deployment of AI models. Its Opinion 28/2024 addresses when an AI model may be considered anonymous, whether legitimate interest can serve as a legal basis in certain circumstances, and the consequences of developing a model with unlawfully processed personal data.

The practical implication is that transforming information into an AI workflow does not make the privacy question disappear. CIOs and CISOs still need to understand why personal data is processed, which models can access it, whether information can be extracted or inferred, where processing occurs and how those decisions can be demonstrated.

The CIO takeaway

Data residency is only one component of AI sovereignty. True control requires visibility into people, data, models, infrastructure, external services and governance.

HIPAA: AI Expands the Surface Area Around Protected Health Information

The HIPAA Security Rule already requires administrative, physical and technical safeguards for electronic protected health information, including controls around access, authentication, integrity and transmission. AI does not remove those obligations; it creates new places where they must be considered.

HHS has also proposed updates to strengthen the HIPAA Security Rule, including more explicit requirements around asset inventories, network mapping, risk analysis, encryption, multifactor authentication, vulnerability management, network segmentation, incident response and compliance auditing. Those changes remain proposed, but they point to a more rigorous cybersecurity posture around ePHI.

HHS guidance now explicitly includes a third-party AI chatbot on a provider patient portal as an example of a business associate when it provides services involving PHI. For technology leaders, the message is clear: AI vendors and AI data flows have to be evaluated as part of the same compliance ecosystem as other systems that create, receive, maintain or transmit PHI.

The useful question is not simply, "Is this AI application HIPAA compliant?" It is, "What happens to PHI throughout the entire AI interaction?"

SOC 2: AI Introduces New Systems That Need Evidence

SOC 2 is not a regulation. It is an assurance framework used to evaluate controls at service organizations against the AICPA Trust Services Criteria. As organizations increasingly rely on third-party services, customers and business partners want evidence that relevant controls are designed and operating effectively.

AI complicates that assurance conversation. An organization may have strong access controls around its systems, but an auditor or customer will increasingly want to know whether those controls extend into AI workflows: who can retrieve which data, which models can process it, how model providers are governed, what gets retained and whether interactions are logged and reconstructable.

Having an AI policy is not the same as having enforceable AI controls - and having controls is different from having evidence that those controls operate effectively.

CAVALRYAI
INNOVATION MEETS SECURITY

The Common Thread: Control the Data Path

GDPR, HIPAA and SOC 2 come from different legal and assurance contexts, but for enterprise AI they point toward a common architectural principle: know where sensitive information goes and control what can happen to it.

That means visibility cannot stop at the database or application. Organizations need to understand the entire AI data path - from user identity and authorization, through policy and model routing, to private or external processing, and finally to logging, monitoring and governance.

Different frameworks. One architectural requirement: control.

CAVALRYAI
INNOVATION MEETS SECURITY

AI Changes the Compliance Equation: What CIOs and CISOs Need to Know About GDPR, HIPAA and SOC 2

As AI moves deeper into enterprise workflows, protecting data is no longer enough. Organizations need to govern who can use it, which models can process it, where it can travel—and prove those controls are working.

By David Whitford, CTO, CavalryAI

GDPR

AI doesn't make personal data less personal.

Organizations must understand why personal data is processed, which models can access it, and whether it can be extracted or inferred. Sovereignty and lawful processing remain essential.

HIPAA

AI expands the surface area around PHI.

Safeguards must follow the data through every AI interaction—including third-party models and services. HHS is proposing expanded security requirements that raise the bar even higher.

SOC 2

AI introduces new systems that need evidence.

Auditors will look for proof that AI access is controlled, monitored, and governed. Policies aren't enough—organizations need evidence that controls operate effectively.

THE AI GOVERNANCE CONTROL PLANE

Security follows the request all the way to the data.

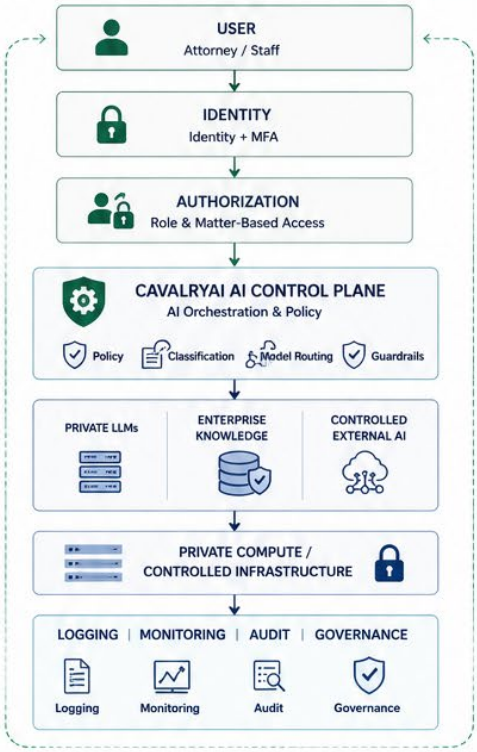

GDPR

- Purpose limitation
- Lawful processing
- Data minimization
- Data subject rights
- Data sovereignty

HIPAA

- Access controls
- Minimum necessary
- Authentication
- Integrity & security
- Auditability

SOC 2

- Security controls
- Availability
- Confidentiality
- Change management
- Evidence & reporting

Different frameworks. One architectural requirement: control.

COMPLIANCE IS BECOMING DYNAMIC

Traditional compliance was system-centric. AI requires the security decision to follow the request—based on the user, data, purpose, model and context.

PRIVATE AI IS NOT AUTOMATICALLY COMPLIANT AI

A private model doesn't automatically provide authorization, classification, retention, auditability, or governance. **Private AI is an architecture. Not a location.**

THE RISE OF PERMISSION-AWARE AI

AI must answer a new question first:

"Should I answer this question for this user using this data and this model?"

Authorization, ethical walls and data boundaries must be enforced in every interaction.

FIVE QUESTIONS EVERY CIO & CISO SHOULD BE ASKING

- What data can our AI actually reach?
- Does authorization follow the user into AI?
- Which models are allowed to process which information?
- Can we reconstruct what happened?
- Can we change the model without losing governance?

COMPLIANCE AS A COMPETITIVE ADVANTAGE

The greatest value comes from the information organizations are most reluctant to expose. Strong governance enables trust, accelerates innovation and turns security into an advantage—not a brake.

THE CAVALRYAI PERSPECTIVE

We design Private AI environments around your data, access requirements, model strategy and governance needs—so you can innovate with confidence and control.

IS YOUR AI ARCHITECTURE READY FOR THE NEXT COMPLIANCE CONVERSATION?

CavalryAI helps organizations assess AI workflows, data, access, models and governance—so you can put your collective intelligence to work. Securely.

SCHEDULE A PRIVATE AI ASSESSMENT

312-729-5400
sales@thecavalry.ai
thecavalry.ai

CAVALRYAI
INNOVATION MEETS SECURITY

CAVALRYAI
INNOVATION MEETS SECURITY

Compliance Is Becoming Dynamic

Traditional compliance has often been system-centric: identify the application containing sensitive data, protect it, control access, monitor it and audit it.

AI changes that model because the same AI interface can interact with radically different classes of information from one request to the next. A public-information query, a confidential contract, PHI and privileged legal work product may all flow through the same user experience.

The security decision therefore needs to follow the request. AI governance increasingly has to understand the user, the data, the purpose, the model and the context of the interaction - not merely the application being used.

Five Questions Every CIO and CISO Should Be Asking

1. What data can our AI actually reach? Not what management believes it can reach. What can it technically retrieve?

2. Does authorization follow the user into AI? If someone cannot access a matter, patient record, financial record or restricted document normally, AI should not become an alternative route to it.

3. Which models are allowed to process which information? Public information may have one policy, sensitive proprietary information another, and highly regulated data still stricter boundaries.

4. Can we reconstruct what happened? Who asked? What information was retrieved? Which model processed it? What policy applied? What response was returned?

5. Can we change the model without losing governance? Models will change. The organization should not need to rebuild its security and compliance architecture every time they do.

Models are replaceable. Governance cannot be.

Private AI Is Not Automatically Compliant AI

Running an LLM privately can be valuable, but a private model does not automatically provide appropriate authorization, data classification, purpose limitation, retention controls, auditability, model governance, ethical walls, threat monitoring or evidence for assurance.

Private AI is an architecture, not a location. The goal is not simply to place a model behind a firewall. It is to create an environment in which identity, data, models, compute and policy operate inside a coherent governance structure.

The Rise of Permission-Aware AI

Traditional AI asks, "Can I answer this question?" Enterprise AI needs to ask another question first:

Should I answer this question for this user, using this data and this model?

The words in the prompt may be identical for a physician, administrator, billing specialist or outside consultant. The authorized answer may not be. The same principle applies in law, finance, government and other regulated industries.

That is why role-based access control, matter and document permissions, data classification and policy-based routing are foundational to enterprise AI rather than administrative afterthoughts.

Compliance Can Become a Competitive Advantage

The objective is not to prevent organizations from using AI. It is to make AI trusted enough to use where it matters most.

The greatest enterprise value often lives in the information organizations are most reluctant to expose: client knowledge, patient information, financial intelligence, trade secrets, intellectual property and years of institutional expertise. Organizations that establish trusted AI boundaries can begin putting more of that information to work.

Security does not have to be the brake on AI. Done correctly, it can become the accelerator.

What the Architecture Should Look Like

A mature Private AI environment should allow organizations to establish different processing rules for different classes of information:

Information Class	Illustrative AI Policy
Public information	Approved external AI may be appropriate.
Internal business information	Approved models based on organizational policy.
Sensitive proprietary information	Private models and controlled infrastructure where required.
Highly regulated information	Additional authorization, processing, logging and governance controls.

Security should be sophisticated. The user experience should be simple.

The CavalryAI Perspective

At CavalryAI, we believe organizations should not have to choose between AI innovation and control of their information.

Our approach begins with the environment: what data exists, where it lives, who should have access, which workflows create value, which information must remain private, when external AI may be appropriate and what must be logged and governed.

Then the AI architecture is built around those answers - not the other way around.

AI Should Fit Your Security Architecture. Not the Other Way Around.

The Board-Level Question

Boards are increasingly likely to ask, "Are we using AI?" That is no longer the hard question.

The harder questions are: What is AI allowed to know? Who determines that? Where does our information go? Which third parties process it? Can we prove our controls worked? What happens when something goes wrong?

Organizations that can answer those questions confidently will be positioned to move faster as AI becomes embedded across the enterprise. Those that cannot may discover that their biggest AI problem is not model performance. It is trust.

The Principle That Does Not Change

The compliance landscape will evolve. Models will change. Regulations will change. Threats and workflows will change.

But one principle is unlikely to change: organizations need control over their data and the intelligence created from it.

The future of enterprise AI is therefore not simply about choosing the most powerful model. It is about creating an architecture where privacy, security, governance and innovation can coexist.

Your Data. Your Domain. Your AI.

IS YOUR AI ARCHITECTURE READY FOR THE NEXT COMPLIANCE CONVERSATION?

CavalryAI works with organizations to assess AI workflows, sensitive data, access requirements, model strategy, infrastructure and governance - and design a Private AI environment around the organization's security requirements.

Schedule a Private AI Assessment

312-729-5400 | sales@thecavalry.ai | thecavalry.ai

Sources & Further Reading

- European Data Protection Board, Opinion 28/2024 on certain data protection aspects related to AI models
- U.S. HHS, HIPAA Security Rule Notice of Proposed Rulemaking Fact Sheet
- U.S. HHS, Business Associates guidance
- AICPA & CIMA, SOC for Service Organizations Engagements - Overview