snare
A PROPHECY SOLUTION

EU AI ACT LOGGING & AUDITABILITY FRAMEWORK

Powered by Snare & AskSnare

From AI Governance to Investigation Readiness

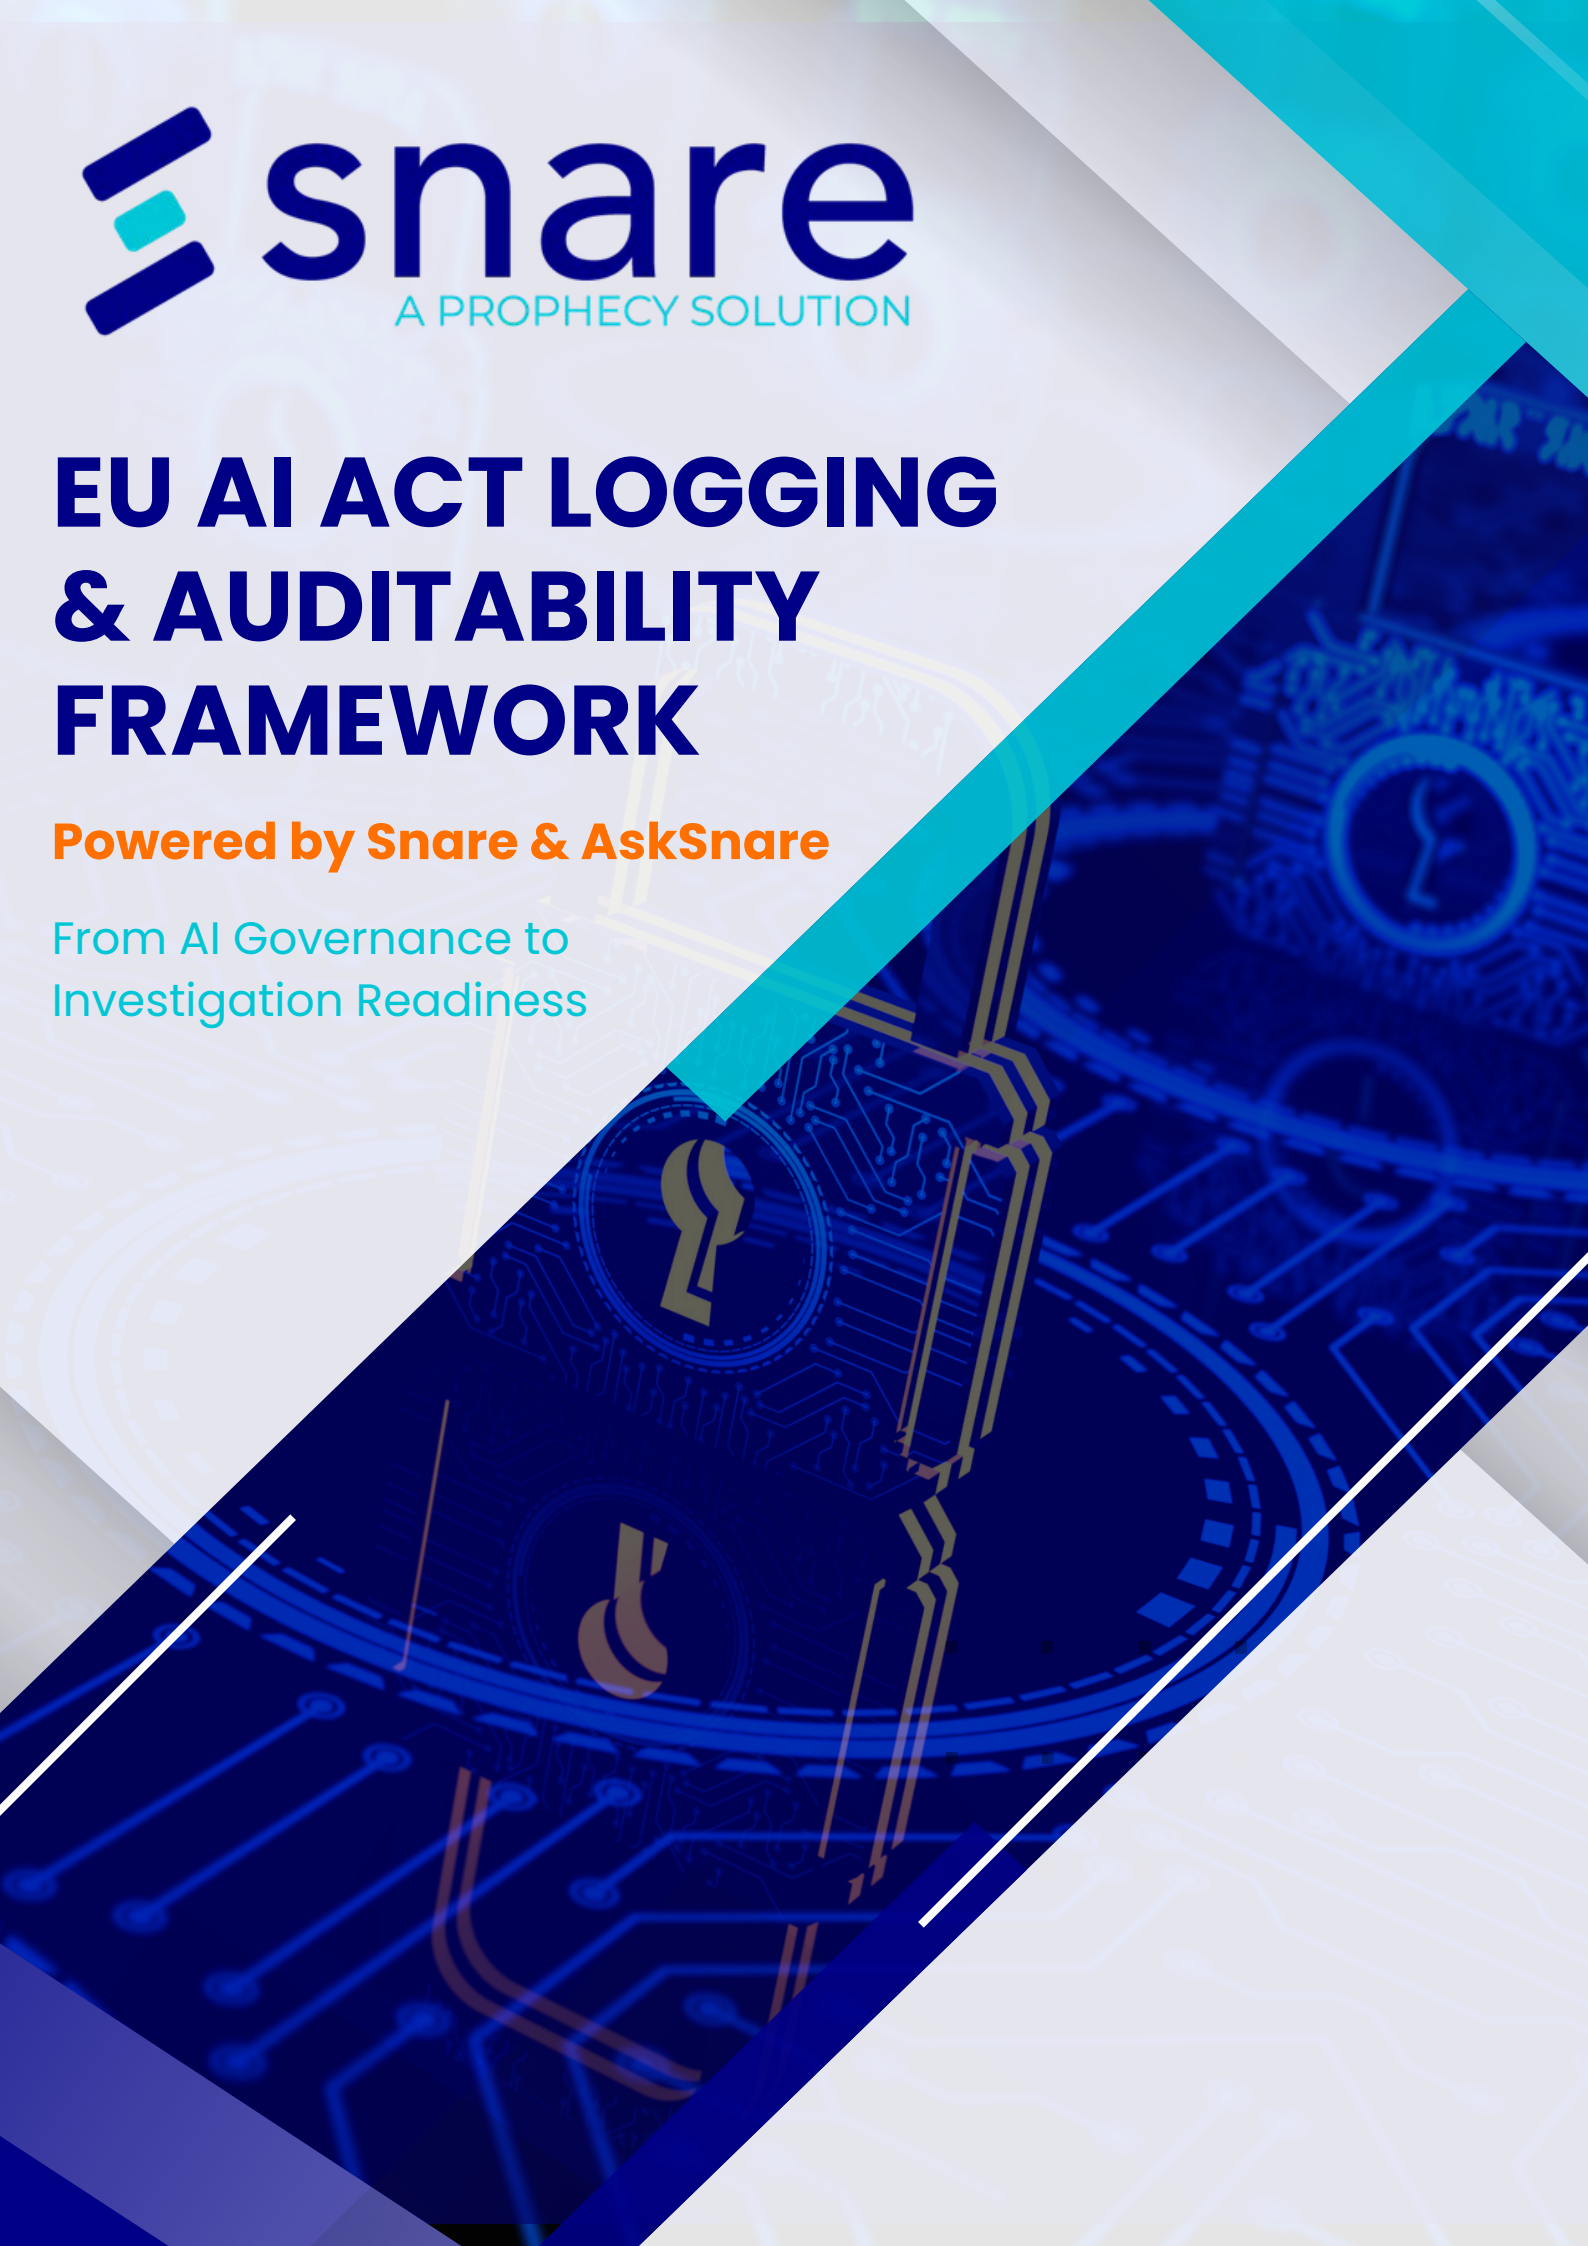

Introduction

Artificial Intelligence is rapidly becoming embedded within business operations, customer interactions, security platforms, analytics systems, and decision-making processes.

As AI adoption grows, regulators are shifting their focus from what AI can do to how organisations govern, monitor, and explain AI-driven outcomes. The European Union AI Act represents the world's first comprehensive AI regulation and is widely expected to influence AI governance practices globally, much like GDPR transformed privacy and data protection standards.

While much of the conversation focuses on risk classifications and compliance obligations, there is another requirement that security and technology leaders cannot ignore:

Auditability.

Organisations must be able to demonstrate transparency, accountability, traceability, and human oversight for AI-assisted systems and decisions. This requires more than policies and procedures.

It requires evidence.

And evidence starts with logging.

Why Logging Matters in the Age of AI

AI systems create new operational risks:

- Automated decision making
- Data access across multiple systems
- Prompt and query activity
- Model interactions
- Administrative changes
- Security incidents
- Unintended outputs
- Regulatory scrutiny
-

When questions arise, organisations must be able to answer:

- What happened?
- When did it happen?
- Who initiated the activity?
- What data was involved?
- What systems interacted?
- What changed?
- Can the event be reconstructed?

Without comprehensive logging and audit trails, these questions become difficult—if not impossible—to answer.

The EU AI Act reinforces the importance of maintaining sufficient records to support transparency, accountability, and oversight throughout the AI lifecycle.

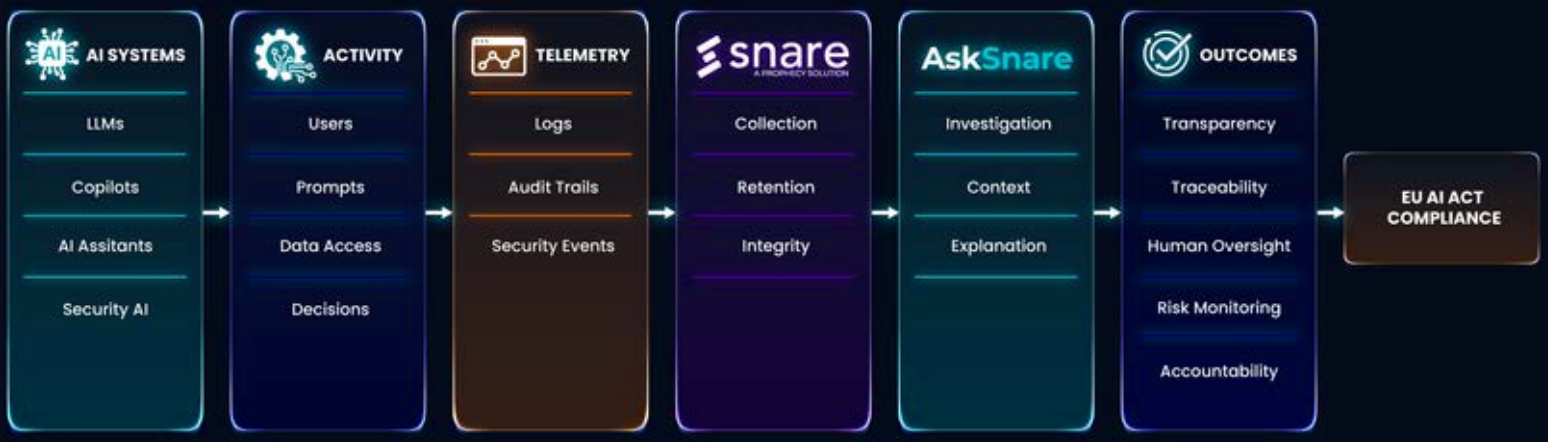

EU AI Act Requirements and Logging Implications

Transparency

Requirement

Organisations must understand how AI systems are used and be able to explain their purpose and outputs.

Logging Implications

Capture:

- User interactions
- AI system access
- Administrative activity
- Data source interactions
- Model execution events

Mini Compliance Checklist

□ Maintain an inventory of AI systems and applications
□ Log all user access to AI platforms
□ Record administrative and configuration changes
□ Monitor interactions between AI systems and enterprise data sources
□ Retain records showing when AI systems were used and by whom
□ Establish ownership and accountability for AI services

Snare Alignment

Snare's forensic-grade logging creates detailed event records that support investigations and compliance reviews long after an event has occurred.

snare
A PROPHECY SOLUTION

Traceability

Requirement

AI-related actions and outcomes must be reviewable and reconstructable.

Logging Implications

Capture:

- User authentication events
- System activity
- Configuration changes
- Data access events
- Administrative actions

Mini Compliance Checklist

□ Enable audit logging across all AI-related systems
□ Capture authentication and privilege changes
□ Record all configuration modifications
□ Maintain logs of data access and movement
□ Synchronise timestamps across all systems
□ Ensure retained logs can reconstruct key decisions and events

Snare Alignment
Snare's forensic-grade logging creates detailed event records that support investigations and compliance reviews long after an event has occurred.

snare
A PROPHECY SOLUTION

Human Oversight

Requirement

Humans must be able to review, challenge, and validate AI-assisted outcomes.

Logging Implications

Organisations need:

- Investigation-ready records
- Searchable audit trails
- Historical event visibility
- Correlated activity data

Mini Compliance Checklist

□ Define escalation procedures for AI-related incidents
□ Ensure authorised personnel can access audit records
□ Maintain searchable investigation records
□ Record actions taken following AI-generated recommendations
□ Validate AI outputs against human review processes
□ Document decision overrides and exceptions

AskSnare Alignment
AskSnare enables analysts and governance teams to investigate activity through natural language queries, accelerating access to the information needed to validate AI-driven decisions.

Risk Monitoring

Requirement

AI systems must be monitored for abnormal behaviour and operational risks.

Logging Implications

Monitor:

- Failed access attempts
- Privilege escalation
- Unusual behaviour patterns
- Service interruptions
- Missing telemetry
- Administrative changes

Mini Compliance Checklist

□ Establish monitoring for AI infrastructure and supporting systems
□ Alert on failed access attempts and suspicious activity
□ Detect unexpected configuration changes
□ Monitor logging health and missing telemetry
□ Review privileged account activity regularly
□ Conduct periodic risk assessments using collected telemetry

Snare Alignment
Snare provides continuous visibility into security events, system activity, and logging health across the environment.

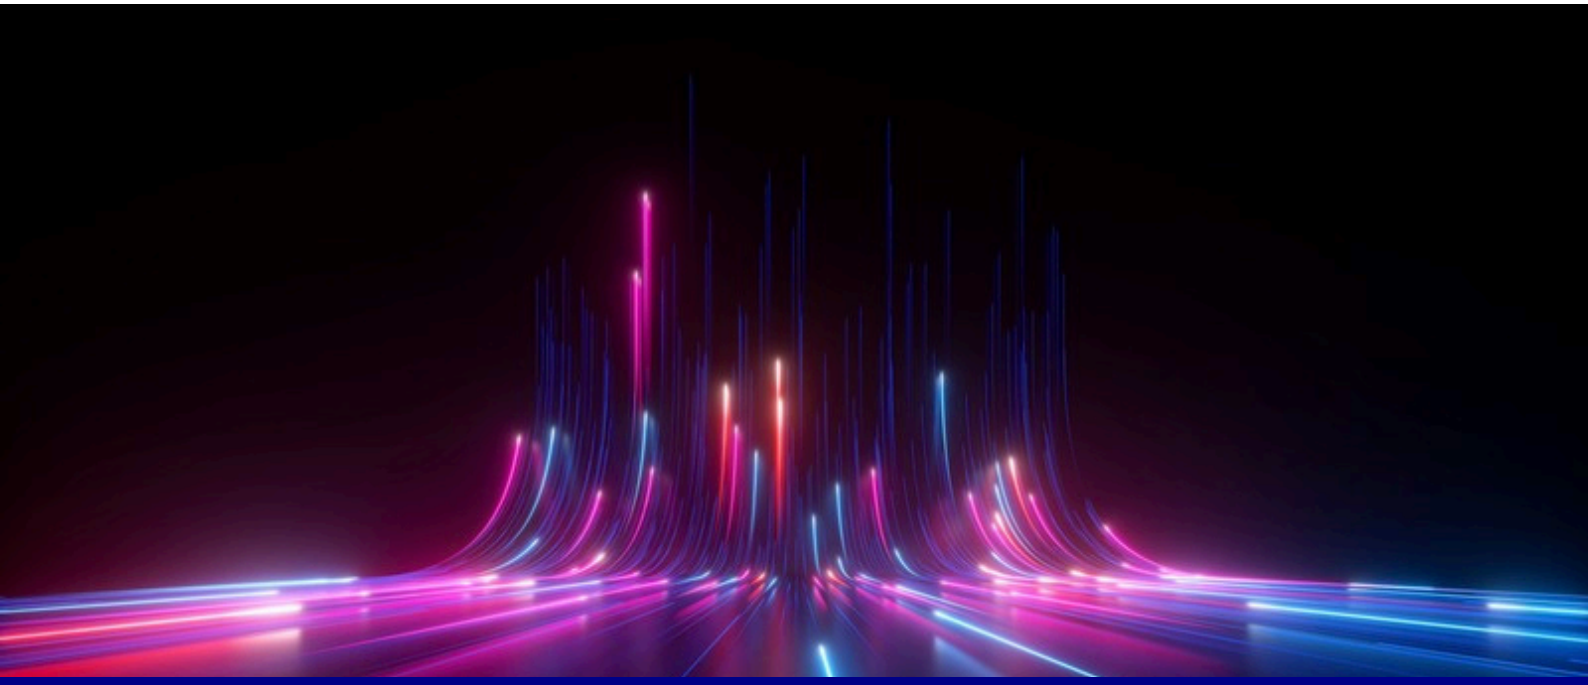

snare
A PROPHECY SOLUTION

Record Keeping

Requirement

Organisations must maintain records sufficient to demonstrate compliance and support audits.

Logging Implications

Requirements include:

- Long-term retention
- Secure storage
- Searchability
- Data integrity
- Evidence preservation

Mini Compliance Checklist

□ Define retention periods for AI-related audit records
□ Implement secure storage and access controls
□ Protect logs against unauthorised modification
□ Ensure records are searchable and retrievable
□ Test evidence retrieval processes regularly
□ Maintain documented retention and deletion policies

Snare Alignment
Snare's long-term retention, compression, replay, and archival capabilities help organisations maintain investigation-ready records while reducing storage costs.

Accountability

Requirement

Organisations must be able to demonstrate who was responsible for AI system operation, oversight, and decision-making.

Logging Implications

Capture:

- User identity
- Administrative activity
- Approval workflows
- Privileged actions
- Policy changes

Mini Compliance Checklist

□ Define ownership for each AI system
□ Record all privileged and administrative actions
□ Maintain approval and change records
□ Monitor access to AI governance controls
□ Preserve evidence of human review activities
□ Produce audit reports on demand

Snare & AskSnare Alignment

Snare captures the evidence required to establish accountability, while AskSnare enables teams to quickly identify who performed specific actions, when they occurred, and what systems were affected.

EU AI Act Logging Maturity Model

Use this maturity model to assess your organisation's readiness to support AI governance, auditability, and compliance requirements under the EU AI Act.

Maturity Heat Map

How Mature Is Your AI Governance Logging Strategy?

Capability	L1	L2	L3	L4	L5
Centralised Logging	●	●	●	●	●
Audit Trails		●	●	●	●
Long-term Retention			●	●	●
Investigation Readiness			●	●	●
AI Activity Monitoring				●	●
Governance Reporting				●	●
Natural Language Investigation					●
AI-Assisted Insights					●

Level 1 – Minimal Visibility

Characteristics

- Basic operating system logs only
- Limited retention periods
- No centralised log management
- AI systems operating without governance visibility
- Investigations are largely manual and reactive

Compliance Checklist

□ Define ownership for each AI system
□ Record all privileged and administrative actions
□ Maintain approval and change records
□ Monitor access to AI governance controls
□ Preserve evidence of human review activities
□ Produce audit reports on demand

Key Risks

- Basic operating system logs only
- Limited retention periods
- No centralised log management
- AI systems operating without governance visibility
- Investigations are largely manual and reactive

Next Step
Establish centralised log collection and begin identifying AI-related systems and activities that require monitoring.

Level 2 – Compliance Awareness

Characteristics

- Centralised logging implemented
- Extended retention periods
- Basic compliance reporting
- Audit trails available but fragmented
- Limited AI-specific monitoring

Compliance Checklist

□ Centralised log collection deployed
□ Retention policies formally documented
□ User access events captured
□ Configuration changes monitored
□ Critical systems included in logging strategy
□ Compliance reporting available

Key Risks

- Limited visibility into AI activity
- Manual audit preparation
- Inconsistent governance controls
- Difficult cross-system investigations

Next Step
Expand visibility beyond infrastructure logs to include AI platforms, applications, data sources, and user interactions.

Level 3 – Investigation Ready

Characteristics

- Forensic-grade logging established
- Consistent retention across environments
- Searchable audit trails
- Historical investigations supported
- AI-related activity visible

Compliance Checklist

□ AI system access logging enabled
□ Data access and movement tracked
□ Audit trails searchable
□ Long-term retention implemented
□ Administrative actions monitored
□ Investigation procedures documented

Key Risks

- Basic operating system logs only
- Limited retention periods
- No centralised log management
- AI systems operating without governance visibility
- Investigations are largely manual and reactive

Next Step
Introduce governance monitoring and continuous oversight of AI systems and supporting infrastructure.

Level 4 – Governance Aligned

Characteristics

- AI governance processes established
- Continuous monitoring of AI environments
- Audit evidence readily available
- Risk monitoring integrated into operations
- Human oversight supported through telemetry

Compliance Checklist

□ AI governance framework documented
□ AI-related activity monitored continuously
□ Privileged actions reviewed regularly
□ Audit evidence available on demand
□ Risk monitoring and alerting configured
□ Human review processes documented

Key Risks

- Large data volumes may slow investigations
- Contextual understanding may require specialist resources
- Governance activities remain labour intensive

Next Step
Leverage AI-assisted investigation and contextual analysis to improve operational efficiency and accelerate governance workflows.

Level 5 – AI Intelligence Driven

Characteristics

- Natural language investigation capabilities
- Context-rich audit and compliance reporting
- Automated risk identification
- Governance visibility across AI environments
- Continuous operational intelligence

Compliance Checklist

□ AI-assisted investigation workflows implemented
□ Natural language access to telemetry available
□ Behavioural anomalies identified automatically
□ Governance reporting automated
□ Investigation readiness regularly tested
□ Continuous improvement program established

Key Outcomes

✔️ Demonstrable AI governance maturity
✔️ Faster investigations and audits
✔️ Reduced compliance burden
✔️ Improved transparency and accountability
✔️ Strong evidence foundation for regulatory reviews

Snare & AskSnare Alignment

At this level, organisations move beyond simply collecting logs. They can rapidly interrogate telemetry, investigate AI-related activity, and demonstrate accountability through trusted evidence and contextual intelligence.

Snare provides the forensic-grade telemetry foundation, while AskSnare enables security, governance, and compliance teams to interact with that evidence using natural language, dramatically reducing the time required to answer critical questions.

Quick Self-Assessment

Maturity Level	Score
Level 1 – Minimal Visibility	1 Point
Level 2 – Compliance Awareness	2 Points
Level 3 – Investigation Ready	3 Points
Level 4 – Governance Aligned	4 Points
Level 5 – AI Intelligence Driven	5 Points

How to Score

- Mostly Level 1 responses = High Risk
- Mostly Level 2 responses = Compliance Developing
- Mostly Level 3 responses = Investigation Ready
- Mostly Level 4 responses = Governance Mature
- Mostly Level 5 responses = AI Governance Leader

Recommended Goal

For organisations implementing AI systems today, the recommended target is Level 4 or above, ensuring sufficient transparency, traceability, oversight, and auditability to support emerging AI governance requirements.

Where AskSnare Changes the Game

Most organisations can collect logs.
Far fewer can understand them quickly.
AskSnare transforms forensic telemetry into actionable intelligence by allowing teams to investigate using natural language.

Examples include:

- What activity occurred before this AI recommendation was generated?
- Which users interacted with this system during the incident window?
- What configuration changes occurred in the last seven days?
- Which critical systems stopped logging?
- Show unusual authentication behaviour relating to AI workloads.

By reducing investigation complexity and improving access to telemetry, AskSnare helps organisations move from simply collecting evidence to actively understanding it.

snare
A PROPHECY SOLUTION

Building Your AI Governance Logging Strategy

Organisations beginning their AI governance journey should focus on five priorities:

1. Identify AI Systems

Document AI platforms, models, and decision-support tools
currently in use.

2. Define Audit Requirements

Determine what evidence must be retained to demonstrate accountability.

3. Establish Logging Standards

Ensure AI-related activity is consistently captured across environments.

4. Centralise Telemetry

Avoid fragmented audit trails across disconnected systems.

5. Improve Investigation Readiness

Ensure teams can quickly access and analyse historical activity
when required.

The Snare Perspective

The EU AI Act reinforces a principle cybersecurity professionals have understood for years: **You cannot investigate what you did not record.**

As organisations adopt AI, trusted telemetry becomes a strategic asset, not just for security operations, but for governance, compliance, and organisational accountability.

Snare provides the forensic-grade logging foundation required to capture and preserve that evidence.

AskSnare helps transform that evidence into accessible, contextual intelligence.

Together, they help organisations prepare for a future where AI governance depends on transparency, accountability, and investigation readiness.

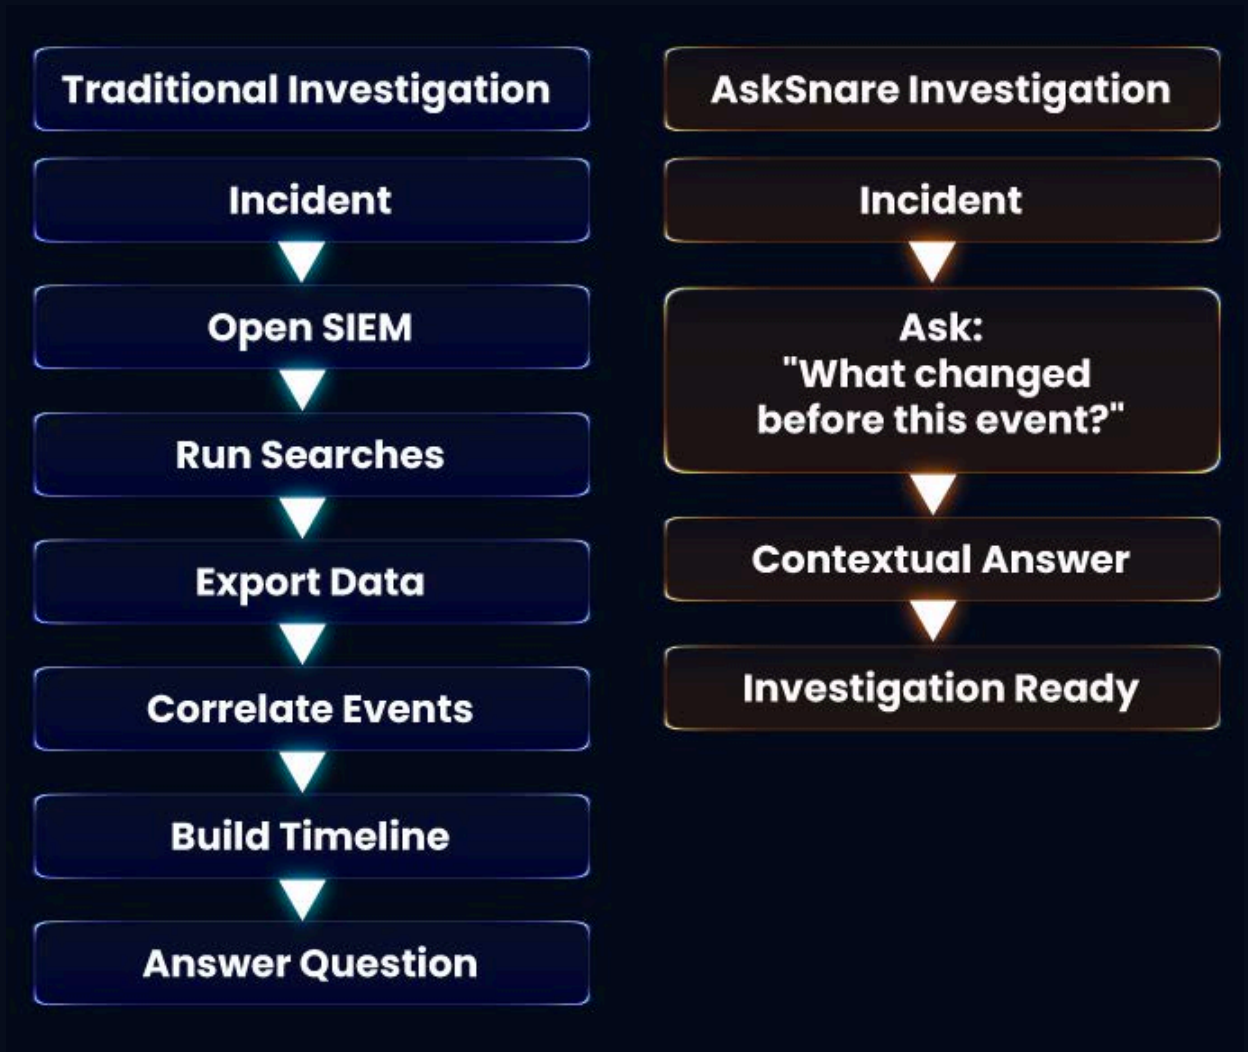

Key Takeaways & Recommended Next Steps

The EU AI Act represents a significant shift in how organisations think about AI.

For many years, logging has primarily been viewed as a cybersecurity, compliance, or operational requirement. The emergence of AI governance changes that perspective.

As organisations increasingly rely on AI to support decisions, automate processes, analyse data, and augment security operations, the ability to explain, audit, and investigate those systems becomes critical.

The organisations best positioned for the AI era will not necessarily be those with the most advanced AI models.

They will be the organisations that can confidently answer three fundamental questions:

What happened? - Can you reconstruct the events, interactions, and activities that occurred across your AI ecosystem?

Why did it happen?- Can you understand the context behind decisions, recommendations, alerts,
and actions?

Can you prove it?- Can you provide evidence to auditors, regulators, customers, and internal stakeholders when accountability is required?

Without trusted telemetry, those questions become difficult to answer.
Without long-term retention, critical evidence may no longer exist.
Without investigation readiness, compliance quickly becomes an operational burden.

The EU AI Act reinforces a principle security teams have understood for years:
You cannot govern what you cannot see, and you cannot investigate what you did not log.

Recommended Next Steps

Whether your organisation is just beginning its AI journey or already deploying AI across the enterprise, there are several practical steps you can take today.

1. Inventory Your AI Environment

Identify:

- AI platforms
- AI-enabled applications
- Large Language Models (LLMs)
- Security AI tools
- Customer-facing AI systems
- Internal AI assistants

Document where AI is being used, what data it accesses, and who is responsible for oversight.

---

2. Assess Your Current Logging Coverage

Review whether you currently capture:

- User access activity
- Administrative actions
- Configuration changes
- Data access events
- Authentication activity
- Security events
- AI-related system interactions

Look for visibility gaps that could impact investigations or audits.

3. Establish AI Governance Logging Standards

Develop consistent policies covering:

- Log collection
- Retention periods
- Audit requirements
- Access controls
- Evidence preservation
- Investigation procedures

Ensure AI-related activity is included within existing security and compliance frameworks.

---

4. Centralise and Retain Critical Telemetry

Fragmented audit trails create governance blind spots.

Establish a trusted telemetry foundation that enables:

- Centralised visibility
- Long-term retention
- Investigation readiness
- Regulatory reporting
- Evidence preservation

---

5. Improve Investigation Readiness

Being compliant is one thing.

Being able to rapidly answer difficult questions during an audit, incident, or investigation is another.

Regularly test your ability to answer:

- Who performed the action?
- What changed?
- When did it occur?
- What systems were involved?
- What evidence supports the conclusion?

6. Prepare for AI-Driven Governance

As AI adoption accelerates, organisations will need tools that help them understand and investigate increasingly complex environments.

Solutions such as Snare and AskSnare provide a foundation for:

- Trusted telemetry
- Auditability
- Traceability
- Investigation readiness
- AI governance visibility
- Faster access to evidence

Final Thought

AI is changing how organisations operate.

The EU AI Act is changing how organisations must govern.

Together, they create a new requirement for visibility, accountability, and evidence.

The organisations that succeed will not be those that simply deploy AI the fastest.

They will be the organisations that build trust in AI through transparency, traceability, and governance.

That trust begins with trusted telemetry.

And trusted telemetry begins with logging.

You cannot govern what you cannot see, and you cannot investigate what you did not log.

Prepare Your Organisation for the AI Governance Era

Discover how Snare and AskSnare help organisations build the transparency, traceability, and investigation readiness required for emerging AI regulations.

Book a Demo

Speak to a Logging Specialist

AMERsales@prophecyinternational.com
APACsales@prophecyinternational.com
EMEAsales@prophecyinternational.com

snare
A PROPHECY SOLUTION