VIEW POINT

AI BILL OF MATERIALS (AIBOM): SECURING THE GENERATIVE AI SUPPLY CHAIN

Abstract

Generative AI ecosystems now depend on interconnected models, datasets, APIs, orchestration layers, and autonomous agents. As enterprises scale AI adoption, governance maturity often lags behind expanding operational complexity. Traditional software governance frameworks cannot provide sufficient visibility into evolving AI dependencies, creating new enterprise risks around trust, compliance, and accountability. An AI Bill of Materials (AIBOM) offers a structured mechanism for improving traceability, runtime visibility, and lifecycle governance across modern AI systems. As organisations increasingly prioritise AI model provenance tracking, AIBOMs are emerging as a foundational capability for operationalising responsible AI, strengthening governance, and managing evolving AI supply-chain risk.

Infosys®
Navigate your next

Introduction: The visibility crisis in enterprise AI ecosystems

Enterprise AI adoption has moved rapidly from experimentation to enterprise-wide deployment. Organisations now embed generative AI into finance, procurement, compliance, HR, customer operations, and knowledge workflows. These systems increasingly rely on open-source models, external APIs, vector databases, and agentic orchestration frameworks, expanding LLM supply chain vulnerabilities across modern enterprises.

McKinsey's 2025 State of AI survey found that 88% of organisations already use AI in at least one business function, while 79% use generative AI capabilities. However, despite rapid adoption, many organisations still struggle to trace model origins, training datasets, runtime dependencies, and third-party AI integrations. This governance gap creates operational blind spots that increase regulatory exposure, reputational risk, and decision-making uncertainty.

As enterprises scale AI adoption, governance priorities are shifting from experimentation toward traceability, operational trust, and lifecycle accountability. This transition is accelerating interest in structured governance mechanisms such as the AI Bill of Materials (AIBOM).

The expanding attack surface of generative AI

Modern AI ecosystems operate as interconnected networks rather than isolated applications. Foundation models, retrieval systems, APIs, and autonomous agents now operate as connected decision ecosystems. This complexity creates governance blind spots that traditional software-centric security frameworks cannot fully address.

Enterprise leaders now face a growing governance challenge: AI dependencies increasingly influence business outcomes in ways organisations cannot fully trace.

From software dependencies to AI dependencies

Traditional software environments typically rely on deterministic dependencies and predictable release cycles. Generative AI ecosystems introduce more dynamic and opaque dependency structures. Modern enterprise AI stacks often include:

- Open-source foundation models
- Fine-tuned LLMs
- Retrieval-augmented generation pipelines
- Third-party inference APIs
- Vector databases
- Agentic orchestration frameworks
- Prompt engineering layers
- Tool integrations and plugins

Many of these dependencies remain externally governed or continuously updated, making operational traceability significantly harder. Unlike conventional applications, AI systems continuously evolve through prompt refinement, runtime orchestration, and adaptive workflows. This makes runtime dependency governance and operational traceability significantly more complex.

Why LLM supply chain vulnerabilities are increasing

The rise of open-source AI development has accelerated innovation, but it has also introduced new operational risks. Organisations increasingly face:

- Poisoned or manipulated training datasets
- Compromised model checkpoints
- Prompt injection attacks
- Insecure third-party plugins
- Hidden downstream dependencies
- Unverified fine-tuning practices
- Misconfigured retrieval systems
- Shadow AI deployments across departments

Many of these risks emerge across external ecosystems that enterprises neither fully govern nor continuously monitor. This creates a governance challenge that extends beyond traditional cybersecurity.

Verizon's 2026 Data Breach Investigations Report found that 31% of breaches now begin with vulnerability exploitation, while shadow AI has become a leading cause of accidental enterprise data exposure. These LLM supply chain vulnerabilities can also affect regulatory compliance, financial reporting accuracy, customer trust, intellectual property protection, procurement integrity, and enterprise decision-making reliability. For enterprise leaders, these risks increasingly represent enterprise governance concerns rather than isolated technology issues.

Why traditional governance frameworks fall short

While Software Bills of Materials (SBOMs) improve visibility across software supply chains, AI ecosystems introduce fundamentally different governance challenges. Unlike traditional applications, AI systems can:

- Adapt behaviour dynamically
- Generate non-deterministic outputs
- Change operational performance without code modifications
- Introduce hidden inference dependencies
- Evolve through continuous fine-tuning
- Rely on runtime orchestration decisions

As a result, governance teams may struggle to identify which dependency, dataset, or orchestration layer influenced a specific AI outcome.

Traditional governance models also struggle because AI systems increasingly operate through interconnected workflows. This shift requires governance models capable of tracking both static and runtime AI dependencies. As a result, enterprises increasingly view the AI Bill of Materials (AIBOM) as the next evolution of operational AI governance.

AI Bill of Materials (AIBOM): Building traceability into AI operations

As AI environments become increasingly distributed, enterprise leaders require clearer visibility into how models, orchestration layers, APIs, and external dependencies interact across workflows.

An AI Bill of Materials (AIBOM) establishes traceability across the full AI lifecycle. Rather than functioning as a static inventory, it enables enterprises to establish traceability, governance transparency, and operational accountability across interconnected AI ecosystems.

Defining the components of an AIBOM

An effective AIBOM should capture the critical elements influencing AI behaviour and governance exposure. The components that can help improve transparency and trust across AI ecosystems include:

- Model origin and ownership
- Training datasets and lineage
- Fine-tuning history
- Prompt orchestration logic
- Third-party APIs and plugins
- Runtime dependencies
- Governance guardrails
- Deployment history
- Model version histories

This level of traceability becomes increasingly important as organisations operationalise interconnected AI systems across enterprise workflows. Without clear dependency visibility, enterprises may struggle to answer critical governance questions such as:

- Which datasets influenced a model outcome?
- Which external APIs affected runtime decisions?
- Which model version generated a customer response?
- Which third-party dependency introduced operational risk?

Shifting from static inventories to continuous assurance

AI governance cannot rely on static documentation alone. Enterprise AI environments continuously evolve through real-time model updates, dynamic orchestration, adaptive prompts, autonomous workflows, third-party integrations, and continuous fine-tuning.

As organisations operationalise AI at scale, governance models must support continuous operational assurance. This is driving demand for dynamic AIBOMs, runtime telemetry, continuous monitoring, real-time governance visibility, lifecycle observability, and automated compliance validation.

Why provenance is becoming central to AI governance

As enterprise AI ecosystems become more interconnected, provenance visibility is emerging as a strategic governance priority. Strong AI model provenance tracking enables organisations to:

- Improve explainability
- Support reproducibility
- Strengthen audit readiness
- Investigate AI incidents faster
- Validate intellectual property ownership
- Increase governance transparency
- Reduce operational uncertainty
- Improve regulatory defensibility

In enterprise AI ecosystems, provenance is rapidly becoming the foundation of governance trust and accountability. Enterprise stakeholders increasingly require evidence-based governance rather than policy-based assurances. Decision-makers increasingly require verifiable visibility into how AI systems evolve, behave, and influence operational outcomes.

Why AI model provenance tracking is becoming mission-critical

Enterprise leaders increasingly recognise that governance policies alone cannot guarantee trustworthy AI outcomes. As generative AI becomes embedded into enterprise workflows, provenance visibility is becoming a core governance requirement rather than a technical enhancement. Organisations now need operational traceability across the full AI lifecycle, from training and deployment to runtime monitoring and updates.

Improving lineage visibility across AI ecosystems

Effective AI model provenance tracking provides visibility into the full AI lifecycle. This includes everything from dataset and model lineage, fine-tuning records, and prompt history to third-party dependency changes. This visibility helps organisations establish evidence-ready AI systems for governance, risk, and compliance.

For enterprise leaders, provenance visibility can improve:

- Enterprise accountability
- Vendor oversight
- Incident investigations
- Governance transparency
- Decision traceability
- Internal audit readiness

It also enables stronger coordination between technology, legal, procurement, risk, and compliance teams.

Supporting auditability and regulatory readiness

Protecto AI reported in 2025 that 77% of organisations are actively building AI governance programmes, yet only 36% have adopted formal governance frameworks such as the NIST AI RMF. This governance maturity gap highlights why enterprises increasingly require embedded operational governance rather than isolated compliance policies.

Regulatory expectations around AI governance continue evolving globally. Frameworks such as the EU AI Act, ISO/IEC 42001, and NIST AI RMF are increasing focus on accountability, explainability, and operational transparency. These frameworks increasingly depend on strong AI model provenance tracking to support AI audits, bias investigations, incident response, regulatory reporting, intellectual property validation, reproducibility requirements, and vendor accountability.

Managing governance complexity in agentic AI systems

McKinsey's 2025 AI research found that 62% of organisations are already experimenting with AI agents. Agentic AI introduces additional governance complexity because autonomous systems increasingly make runtime decisions independently.

Modern agentic ecosystems often involve:

- Tool calling
- Multi-model orchestration
- Adaptive workflows
- Autonomous reasoning loops
- Dynamic retrieval pipelines
- Cross-platform integrations

As these environments evolve, static governance models struggle to support adaptive AI decision environments. Governance visibility must evolve at the same pace as AI autonomy. This is accelerating enterprise demand for continuous governance controls, runtime telemetry, provenance-aware orchestration, structured AIBOM compliance frameworks, and lifecycle assurance systems.

From compliance policies to operational AI governance

AI governance has evolved beyond cybersecurity or compliance alone. Organisations now require operational models capable of embedding governance directly into enterprise AI lifecycles.

Breaking down governance silos

AI governance responsibilities increasingly span multiple business functions, including security, compliance, procurement, operations, and risk management. Siloed governance structures often create fragmented oversight and inconsistent accountability.

As enterprises scale AI adoption, governance models must become more integrated, operational, and continuous.

Enabling governance through AIBOM compliance frameworks

Modern AIBOM compliance frameworks support more than regulatory reporting. They also strengthen operational resilience by improving auditability, third-party risk visibility, governance automation, lifecycle monitoring, operational accountability, and AI transparency. This operational approach helps organisations move from reactive governance toward continuous AI assurance.

Operationalising responsible AI at enterprise scale

Responsible AI requires more than policy statements. It requires governance controls capable of monitoring how AI systems behave across evolving enterprise environments.

Effective governance programmes increasingly depend on:

- Continuous monitoring
- AI telemetry
- Lifecycle governance
- Provenance visibility
- Dependency traceability

As AI ecosystems continue to evolve, enterprises will require governance frameworks capable of embedding trust directly into operational workflows. Approaches focused on enterprise responsible AI governance can help organisations strengthen visibility, resilience, and accountability across the AI lifecycle.

Conclusion: The future of AI governance will depend on verifiable trust

Enterprise AI ecosystems will continue becoming more autonomous and interconnected. As organisations expand generative AI adoption, governance maturity will increasingly depend on provenance visibility, dependency traceability, runtime assurance, and continuous oversight.

The AI Bill of Materials (AIBOM) is emerging as more than a governance artefact. It is becoming foundational governance infrastructure for operationalising trust across modern AI environments. Strong AI model provenance tracking and mature AIBOM compliance frameworks can help organisations strengthen resilience, improve audit readiness, and reduce evolving LLM supply chain vulnerabilities.

As enterprises operationalise AI at scale, governance models that embed continuous traceability, runtime accountability, and provenance visibility will increasingly define responsible AI maturity.

For more information, contact infosysbpm@infosys.com

Infosys®
Navigate your next