
ArxNimbus
ThrivacaAIQ™
PRODUCT SUMMARY

WHAT IS AIQ™?

ArxNimbus ThrivacaAIQ (AI Exposure Quantification) is the only platform that translates AI exposure into defensible financial impact—connecting AI models, business processes, and revenue to quantified EBITDA risk.

Powered by the patented Thrivaca™ actuarial engine, AIQ uniquely combines:

- actuarial-grade financial exposure modeling
- AI Bill of Materials (AIBOM) intelligence

to quantify how AI risk propagates from technical systems to business impact to financial loss.

AIQ translates complex AI vulnerabilities, dependencies, and failure modes into defensible, NIST AI RMF and ISO 42001-aligned financial exposure metrics that executives, boards, regulators, and security leaders can understand and act on.

- Not a governance checklist. Not a scoring model.
- AIQ delivers a single, actuarial source of truth for AI exposure—in dollars, not just scores.

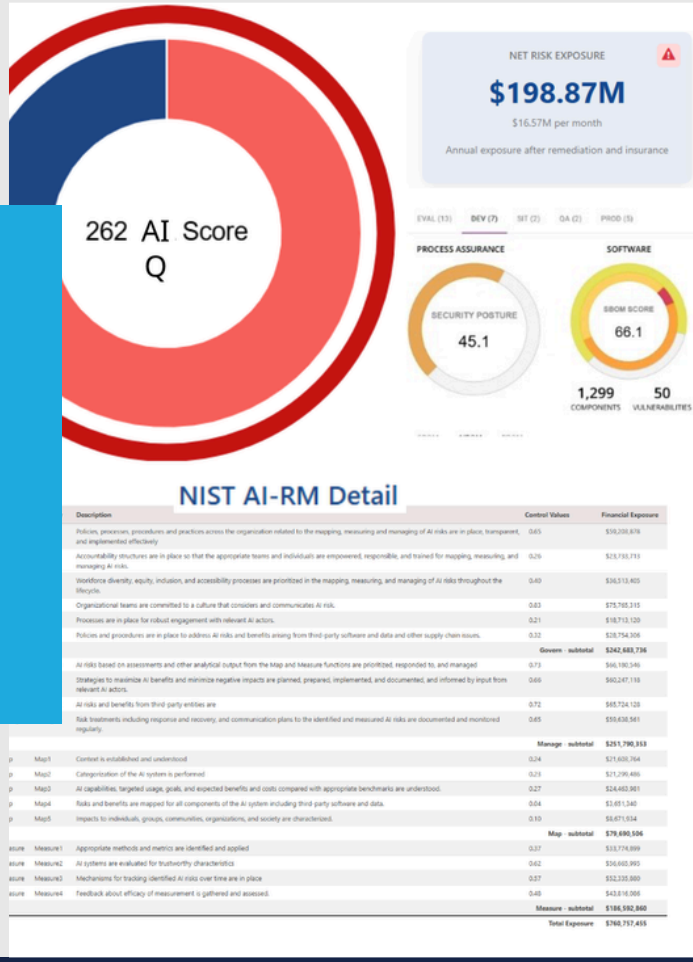

1. Insurance-Grade Actuarial
2. AI Exposure Quantification

Trusted by Insurers Underwriting 35% of the S&P 500

Only Solution Quantifying
AI Exposure in Financial Terms

ArxNimbus

ThrivacaAIQ™

Achieve 1-3% EBITDA improvement through quantified AI exposure reduction

Powered by the Thrivaca Actuarial Engine

Quantifies AI Exposure in CFO Terms:
Net Risk, ROI, EBITDA Impact

Aligned to NIST AI RMF & ISO 42001

AI IS SCALING FASTER THAN IT CAN BE MEASURED. AIQ CHANGES THAT.

AI is being deployed faster than it can be measured in financial terms. ***Boards Ask:*** ***"How much EBITDA is at risk from AI—and where?"*** Most organizations cannot answer with financial precision.

PROBLEM:	AIQ SOLVE:
AI Shadow IT: AI initiatives launched without security oversight or accountability ("keep the suits out").	Knowledge graph inventory of AI Models, with hashed date stamped auditable record.
Regulatory Tsunami: 426 AI regulations already adopted globally, 400+ more under discussion	Compliance framework aligning to Federal Reserve SR 11-7, NAIC ORSA and ASOP 56. Cross walk standard automation between existing and emerging NIST, EU and other sector compliance frameworks.
Subjective Scoring: Existing governance still relies on static, opinion-based frameworks (FAIR/Monte Carlo)	NIST-aligned actuarial-based scoring of existing systems, models, and compliance.

THE 84%-11% GOVERNANCE GAP

- 84% of Fortune 500 CIOs/CEOs say AI is critical to the future economic success *(KPMG 2025)*
- Only 11% have adequate AI governance in place
- **This is the largest unquantified financial risk in the enterprise today**

Without actuarial quantification, AI exposure remains unmeasured, unmanaged, and financially indefensible.

TURNING AI EXPOSURE INTO FINANCIAL CERTAINTY.

ThrivacaAIQ™ applies the Thrivaca actuarial engine across two dimensions of AI exposure:

Component 1: Financial Exposure Quantification

- NIST AI RMF and ISO 42001-aligned financial exposure assessment
- Actuarial-based quantification (not FAIR/Monte Carlo)
- Net Risk Exposure, Accessible ROI, EBITDA Assurance
- 11-question AI initiative framework

Component 2: Technical Exposure Intelligence

- AIBOM extraction from Hugging Face and model registries
- Business Entity-Process mapping (models → processes → P&L)
- Knowledge Graph visualization
- Trustworthy AI scoring (accuracy, bias, model validation)

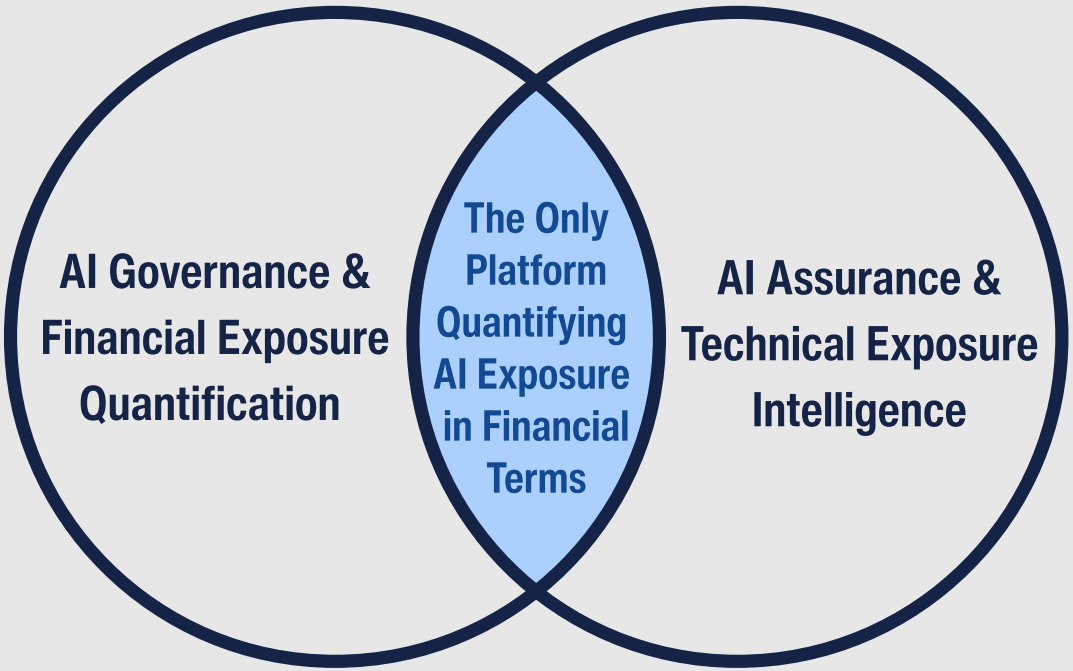

One Actuarial Engine. Insurance-Grade Rigor.

ThrivacaAIQ™ applies the Thrivaca actuarial engine—already validated across insurance and enterprise markets—to AI exposure.

AIQ CORE CAPABILITIES.

- ACTUARIAL AI EXPOSURE QUANTIFICATION
- AI BILL OF MATERIALS (AIBOM) INTELLIGENCE
- BUSINESS ENTITY PROCESS MAPPING
- KNOWLEDGE GRAPH VISUALIZATION
- MULTI-FRAMEWORK SUPPORT
- ENTERPRISE AI PORTFOLIO VIEW

ACTUARIAL AI EXPOSURE QUANTIFICATION

FUNCTION	WHAT IT DOES	WHY IT MATTERS
Framework Alignment	Maps AI risk quantification to both leading AI Risk Management frameworks NIST AI RMF and ISO 42001	Aligning AI risk exposure in financial terms to NIST AI RMF and ISO 42001 helps tie AI controls directly to economic impact, improving governance, prioritization, board reporting, and regulatory/legal defensibility.
Probability-Based Modeling	Applies actuarial loss distributions validated by insurance industry data to quantify AI risk—eliminating subjective scoring and opinion-based assessments	Produces audit-defensible, repeatable AI risk estimates that regulators and boards trust—reducing bias, improving capital allocation, and providing transparency that subjective scoring and opaque simulations cannot deliver.
Financial Impact Translation	Translates AI threats into dollar-denominated exposure ranges, expressing risk in the same financial language used for all other business decisions	Translate AI risks into quantified dollar exposure so they can be compared directly to revenue, EBITDA, margins, and capital allocation—enabling data-driven prioritization, budget optimization, improved risk visibility, and clear executive/board accountability.
Defensible Loss Estimates	Produces AI risk estimates using auditable actuarial methodology validated and accepted by insurance carriers and regulatory bodies	Defensible actuarial loss estimates aligned to leading frameworks give boards and regulators auditable, trusted numbers, improving capital planning, insurance purchasing, compliance evidence, and reducing disputes over AI-related incident costs and disclosures.

AI BILL OF MATERIALS (AIBOM) INTELLIGENCE

FUNCTION	WHAT IT DOES	WHY IT MATTERS
Business Process Identity	Links each AI model to its specific business context—which entity owns it, what process it supports, and what data it accesses	Enables Responsible AI governance and automated compliance reporting by establishing
Automated ModelIngestion	Automatically catalogs AI model components and assigns static vulnerability scores, identifying risks in open-source libraries, dependencies, and model architecture	Detects vulnerable open-source components that could compromise model outputs or leak sensitive data—before they're exploited in production environments
Automated Model Dynamic Testing	Runs user-selected “Red Teaming” automatically to verify model bias and pedigree	Verifiable results will document the model resistance to bias or altercation drift
Automated AIBOM Generation	Automatically generates comprehensive AI Bills of Materials documenting complete technical lineage—every component, version, dependency, and known vulnerability	Automates internal governance and external regulatory reporting requirements—providing audit-ready documentation of AI composition and supply chain integrity
AIBOM hashed and date stamp audit file	Creates cryptographically hashed, timestamped audit records for each AI model release in standardized JSON format (CycloneDX and SPDX compliant)	Delivers tamper-resistant evidence of model composition at every release—supporting continuous compliance monitoring, incident response and investigation, and regulator-ready audit trails. CycloneDX/SPDX alignment ensures AIBOMs integrate cleanly with enterprise security tooling, GRC platforms, and regulatory reporting workflows

AI BILL OF MATERIALS (AIBOM) INTELLIGENCE

FUNCTION	WHAT IT DOES	WHY IT MATTERS
Continuous Monitoring	Continuously monitors deployed AI models, automatically rescoring assurance levels and alerting stakeholders when changes breach user-defined thresholds	Catches model degradation, drift, or new vulnerabilities in production—triggering mitigation before compromised AI outputs impact business operations or compliance posture
Integration with Existing Systems	Exports AIBOMs in standard formats to integrate with existing GRC platforms, security tools, and compliance management systems	Embeds AI assurance into existing governance workflows rather than creating separate processes—AI risk management becomes part of enterprise GRC, not a standalone silo

BUSINESS ENTITY PROCESS MAPPING

FUNCTION	WHAT IT DOES	WHY IT MATTERS
Model-to-Process Linkage	Maps each AI model to the business processes it supports, creating visibility into which models affect critical operations (e.g., GPT-4 used in Loan Underwriting, Customer Service, and Fraud Detection)	Enables impact analysis when AI models fail or change—immediately identify which business processes and revenue streams are affected, ensuring continuity planning and compliance
Process-Level Risk Quantification	Quantifies AI risk by business process, showing exposure concentration (e.g., Personnel Security: 58% of risk = $7.7M total, with $4M concentrated in Mortgage Processing)	Identifies highest-impact AI risk concentrations so teams can prioritize remediation efforts. AIBOMs provide technical detail for targeted mitigation and measured improvement"
Revenue Stream Attribution	Maps AI risk exposure to specific revenue streams, showing financial impact by product line or business unit (e.g., $12M exposure in Consumer Lending vs. $3M in Wealth Management)	Connects risk mitigation directly to EBITDA improvement, showing CFOs how reducing AI risk in specific business units protects and enhances bottom-line performance.
Remediation ROI	Calculates cost versus risk reduction for each mitigation action, showing financial return per remediation investment (e.g., $50K control spend eliminates $2M exposure = 40:1 ROI	Justifies AI assurance budgets by demonstrating measurable financial returns, enabling CFOs to treat AI risk mitigation as an investment rather than a cost center

KNOWLEDGE GRAPH VISUALIZATION

FUNCTION	WHAT IT DOES	WHY IT MATTERS
Interactive Relationship Mapping	Models → Data Sources → Business Processes → Financial Impact	The cumulative process assurance history over a selected time period supported with a projected EBITDA save can be visualized.
Dependency Chain Visibility	Visualizes how AI models connect across business processes using knowledge graphs, revealing upstream and downstream ripple effects when any model changes or fails	Automatically alerts when AI model changes affect business processes. Traces dependency chains to show cascading impacts—one compromised model flags all connected operations and updates aggregate risk scores
Version Tracking	Tracks AI model versions over time, detecting when models drift from baseline performance or when unannounced changes alter behavior (e.g., GPT-4 updated from v1.0 to v1.2 without notification)	Prevents silent failures by detecting model drift and undocumented changes before they impact business operations, maintaining assurance scores based on tested versions
Stakeholder View	Presents the same AI risk data through role-specific lenses—executives see business impact and EBITDA exposure, while technical teams access model configurations and remediation details	Enables cross-functional alignment by giving each stakeholder the view they need—CFOs see financial risk, CISOs see security controls, operations see process impact—all from one source of truth

MULTI-FRAMEWORK SUPPORT

FUNCTION	WHAT IT DOES	WHY IT MATTERS
NIST AI RMF	Structures AI risk assessment around NIST AI RMF's Govern, Map, Measure, Manage framework with automated control family evaluation and NIST 800-53 alignment	AI model components are continuously assessed through vulnerability analysis, implementation verification, and static/dynamic testing, then automatically mapped to NIST AI RMF requirements via knowledge-graph intelligence—aggregating evidence and generating compliance reports without manual documentation.
ISO 42001	Aligns to ISO 42001 international AI management system standards, supporting global compliance requirements	Enables automated ISO 42001 reporting by mapping model components, NIST 800-53 controls, and CWE vulnerability data to international standard requirements—critical for multinational operations
Automated Crosswalk	Automatically maps AI assessments across evolving frameworks using third-party crosswalk automation—eliminating manual framework translation as requirements change	With 428 global AI frameworks today and 400+ emerging in 2026, automated crosswalking is essential to maintain compliance without exponentially increasing assessment workload
Regulatory Adaptability	Maintains compliance mappings for EU AI Act, U.S. state regulations, and sector-specific requirements (Banking SR 11-7, Insurance ORSA/NAIC AI Bulletin/ASOP 56)	Adapts and supports sector-specific assessments to EU AI Act, emerging state laws, and evolving sector regulations (financial services SR 11-7, insurance ORSA Model #505, NAIC AI guidance, ASOP 56)

ENTERPRISE AI PORTFOLIO VIEW

FUNCTION	WHAT IT DOES	WHY IT MATTERS
Multi-Initiative Analysis	Analyzes AI risk at any scale—from individual AI projects to enterprise-wide portfolios across multiple business units with unified visibility	Enables both tactical project-level decisions and strategic portfolio governance—CISOs see enterprise exposure while business units manage their own AI initiatives independently.
Materiality-Based Prioritization	Ranks AI initiatives by quantified financial exposure rather than subjective risk scores—prioritizing based on actual dollar impact to revenue and operations	Enables data-driven resource allocation by measuring, not approximating, business unit risk—invest in mitigating the $10M exposure before the $500K exposure
Portfolio Optimization	Calculates exposure-to-value ratios for each AI initiative—identifying which projects carry disproportionate risk relative to their business contribution	Reveals misaligned risk-reward profiles across the portfolio—enables strategic decisions to remediate high-risk/low-value initiatives or sunset AI projects that don't justify their exposure
Board-Ready Dashboard	Delivers executive-level AI risk dashboard with key performance indicators, trend analysis, and industry benchmark comparisons—designed for board and C-suite consumption	The knowledge graph foundation provides board-appropriate high-level insights with ability to drill into supporting detail—executives get strategic visibility without technical noise, but can investigate when needed

KEY OUTCOMES WITH ThrivacaAIQ

- **Unified Financial Truth for AI Exposure**
  Defensible, dollarized exposure metrics trusted by insurers and enterprise leaders.
- **Board-Ready Clarity**
  AI exposure communicated in financial terms that executives can govern and measure.
- **Regulatory & Insurance Confidence**
  Auditable, NIST-aligned AI exposure assurance supporting compliance and underwriting.
- **Scalable AI Assurance**
  Extends seamlessly from individual AI initiatives to enterprise-wide AI environments.

ThrivacaAIQ IDEAL FOR...

- Enterprises scaling AI across business functions
- CISOs and risk leaders quantifying AI exposure in financial terms
- CFOs accountable for AI-related financial exposure and insurance impact
- AI initiative leaders responsible for governance and deployment confidence
- Boards fulfilling fiduciary responsibility in the age of AI

ThrivacaAIQ DEPLOYMENT SNAPSHOT

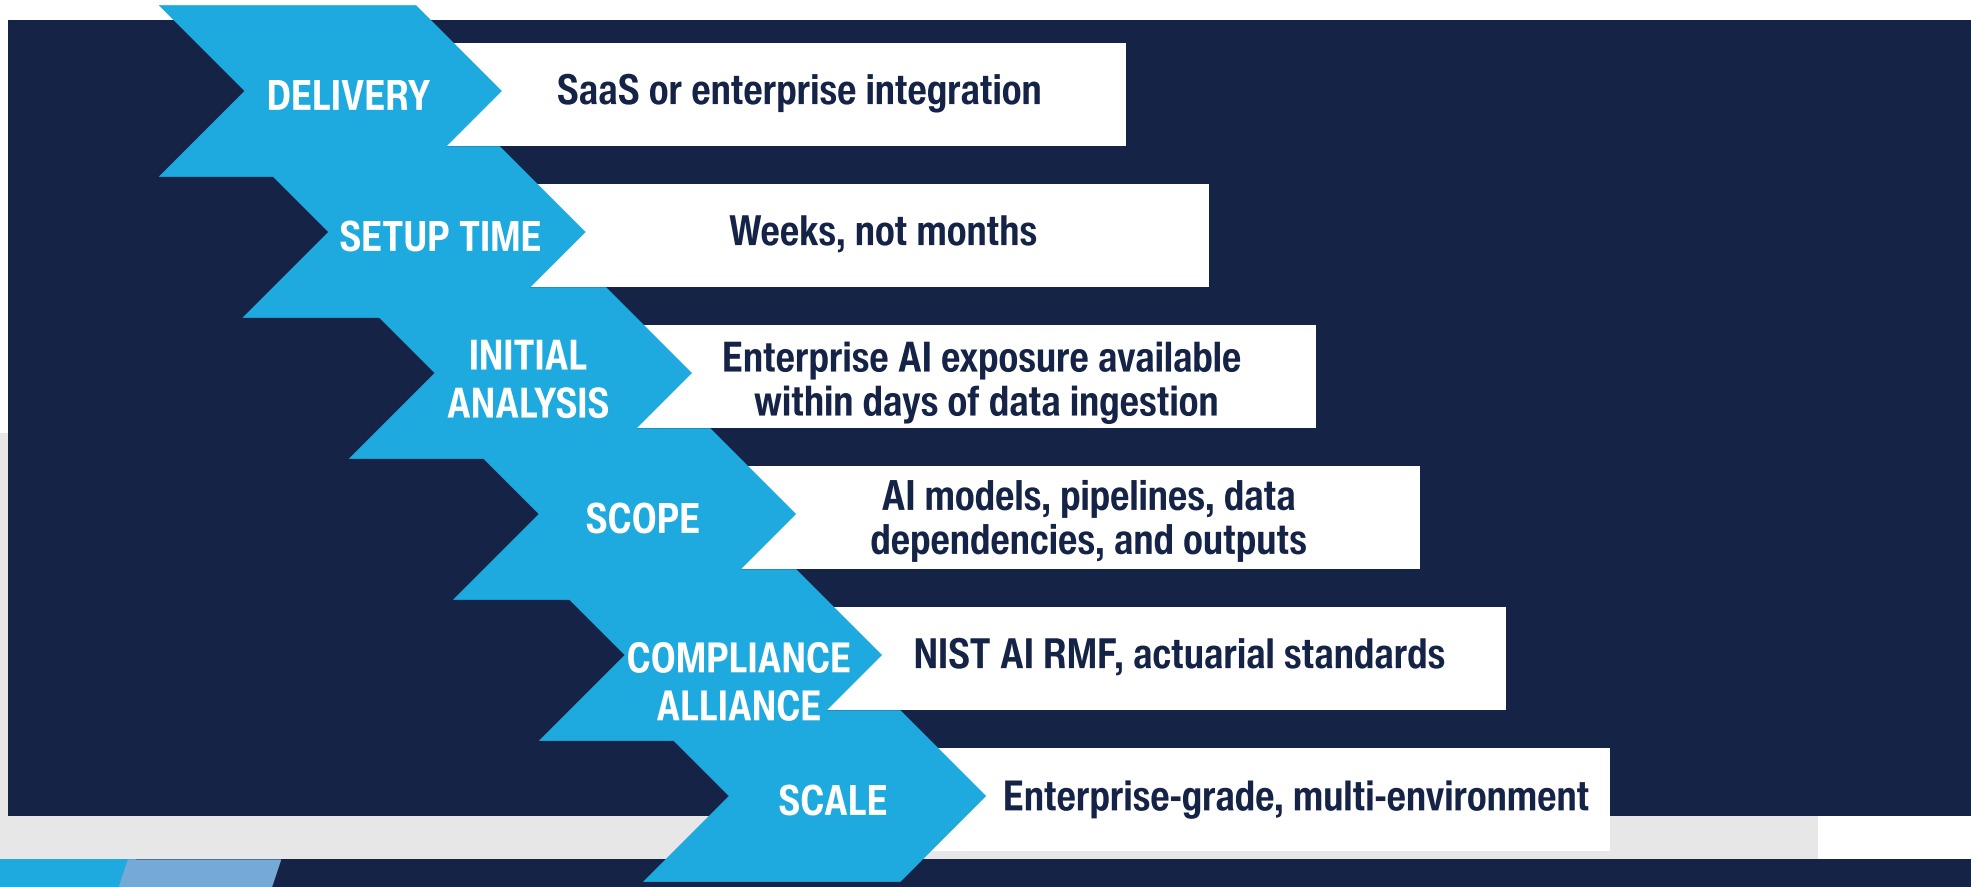

© 2026 ArxNimbus ThrivacaAIQ | Product Summary | p 16 of 17

KEY DIFFERENTIATORS

Knowledge Graph Architecture:
Scalable Neo4J-powered platform organizing relationships between vulnerabilities, assets, and business processes

Proprietary RBOM:
Unique Release Bill of Materials tracking—no capability other platform offers this operational assurance

Industry Evolving AIBOM:
Most comprehensive AI Bill of Materials with direct CISA partnership defining emerging standards. Knowledge graph extensibility of additional AIBOM elements as technology matures.

Actuarial-Grade Financial Translation:
Using 300+ years of insurance industry actuarial science for unprecedented accuracy

MITRE | ATT&CK®

Recognized by Gartner, Momentum Cyber, and leading academic institutions for innovation in cyber and AI risk.

ArxNimbus

ArxNimbus was built on a simple idea: AI + actuarial modeling is the most reliable way to measure exposure in financial terms.
Since 2016, we've used AI to normalize fragmented security data—and actuarial science to quantify its financial impact.
The result: a single, defensible financial system of record for cyber and AI exposure.

Built in collaboration with leaders across national security, academia, and industry—including MITRE, the University of Chicago, and the U.S. Department of Defense—the Thrivaca™ engine reflects how risk is actually modeled and validated in the real world.

ThrivacaAIQ™ applies that engine to AI exposure—bringing insurer-grade financial rigor to AI-driven decision-making.
Nine years of AI. Not nine months.

888-422-6584
info@arxnimbus.com
arxnimbus.com/AIQ