CYERA

GUIDE

Laying a Foundation for Secure AI

How Cyera Supports Compliance with the NIST AI Risk Management Framework

Table of Contents

Introduction

About the NIST AI Risk Management Framework

What data do I have?

The NIST AI RMF is a voluntary framework designed by experts from industry, government, academia, and civil society. It helps organizations identify and mitigate risks associated with AI system development and deployment, and consists of four core functions - divided into various categories and subcategories - that aim to ensure the trustworthiness of AI systems.

The four functions are Govern, Map, Measure, and Manage.

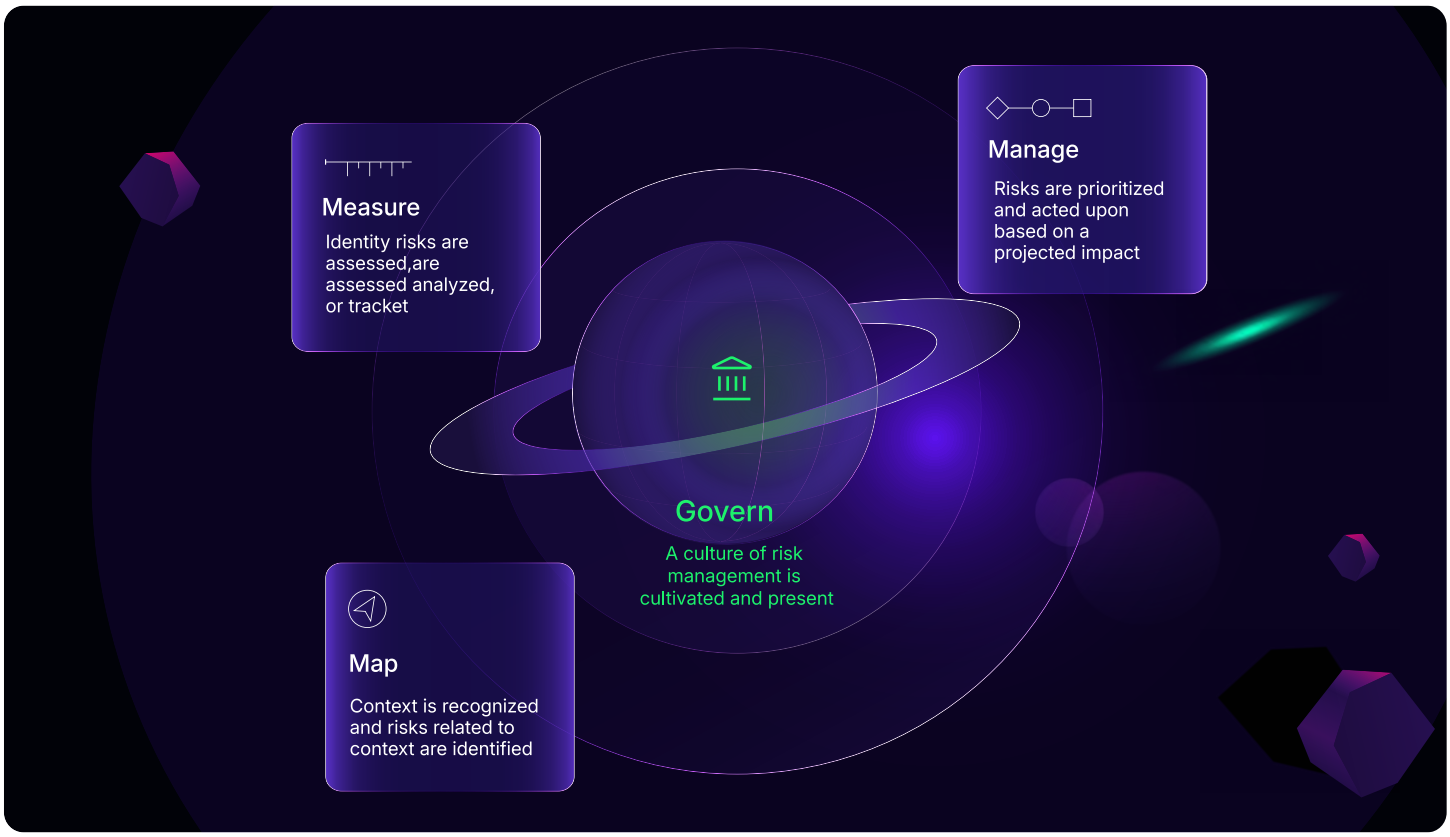

Govern

Function involves understanding applicable compliance obligations, defining roles and responsibilities, assessing risks, and defining organizational risk tolerance.

Map

Applies general risk management principles to the specific context of AI model design and development.

Measure

function focuses on developing metrics and methodologies to track the trustworthiness, efficacy, and potential risks of AI models over time.

Manage

function prioritizes and responds to the risks and benefits of AI systems throughout their lifecycle.

Introduction

About Cyera

What data do I have?

Cyera is a unified, AI-native data security platform that empowers organizations to discover, classify and protect data. It enables security leaders to manage sensitive data across highly permissive and widely distributed environments with high precision and efficiency.

The platform's agentless, fully automated data discovery provides a comprehensive inventory of sensitive data across structured and unstructured sources, including IaaS, SaaS, DBaaS and on-premises environments. This capability addresses critical data challenges like data proliferation and drift. Powered by AI-native classification, Cyera goes beyond traditional methods by also understanding context, intent, and nuance - decoding data down to the DNA level. This deep insight uncovers ghost data, reveals sensitive data risks, reduces false positives, and mitigates threats like data breaches and ransomware — areas where conventional DLP and data governance tools fall short.

Cyera goes beyond traditional methods by also understanding context, intent, and nuance - decoding data down to the DNA level

Cyera's AI Guardian solution supports all four NIST AI RMF functions through a tightly integrated suite:

- **AI Security Posture Management (AISPM):** Inventories AI tools and usage patterns, mapping organizational exposure to AI misuse and Shadow AI.
- **Runtime Protection:** Detects and enforces data policies in real-time during AI tool interactions, blocking unsafe prompts or outputs.
- **Omni DLP**: Classifies structured and unstructured data at creation and use, with over 95% precision, and ensures PII, PHI, and IP are safeguarded from LLM ingestion or AI leaks.
- **AI Risk Assessment & Breach Readiness:** Prepares stakeholders for AI misuse or exfiltration events through risk modeling, simulation, and response automation.

These capabilities directly enable organizations to operationalize the NIST AI RMF and demonstrate measurable, risk-aware AI governance.

Govern

Governance Control	How Cyera Supports it
GV-1.1 **Legal and regulatory requirements involving AI are understood, managed, and documented.**	• Cyera's AI-native data classification engine comes pre-trained with classifiers aligned to major regulatory and industry frameworks (e.g., GDPR, HIPAA, and PCI DSS). Out-of-the box policies, also aligned with these frameworks, trigger alerts, notify data owners with remediation instructions, and integrate with workflow tools for automated remediation. • Cyera ensures AI models do not ingest sensitive personal information and that any applicable data sovereignty requirements are respected.
GV-1.6 **Mechanisms are in place to inventory AI systems and are resourced according to organizational risk priorities.**	• Cyera's DSPM provides unparalleled visibility into all your data, including its location, who's using it, and what they're doing with it. It continuously monitors your entire data estate, detecting changes to data and new data creation. • Cyera's AI Security Posture Management (AISPM) provides automated discovery and inventory of all AI systems, including Shadow AI apps, plugins, and LLM-based tools, supporting continuous system inventory and AI lifecycle tracking.
GV-1.7 **Processes and procedures are in place for decommissioning and phasing out AI systems safely and in a manner that does not increase risks or decrease the organization's trustworthiness.**	• Cyera can help you validate the decommissioning and phasing out of AI systems by confirming whether organizational data still resides within them, or if they retain access to organizational data stores.
GV-2.1 **Roles and responsibilities and lines of communication related to mapping, measuring, and managing AI risks are documented and are clear to individuals and teams throughout the organization.**	• Cyera's Data Risk Assessment, AI Risk Assessment, and Breach Readiness services can help you better understand the gaps in your data security posture, including developing strategies for risk mitigation and improving your incident response plan. • Through asset discovery and classification, Cyera can also help your organization identify data owners and their responsibilities.

Governance Control	How Cyera Supports it
GV-2.2 **The organization's personnel and partners receive AI risk management training to enable them to perform their duties and responsibilities consistent with related policies, procedures, and agreements.**	• Cyera's AI Guardian raises awareness among personnel regarding responsible AI usage by notifying them of policy violations, blocking sensitive prompts, and redacting regulated data from Copilot interactions.
GV-2.3 **Executive leadership of the organization takes responsibility for decisions about risks associated with AI system development and deployment.**	• Cyera's Data Risk Assessment and AI Risk Assessment services help organizational leadership better understand the risks that AI poses to their data estate, enabling them to formulate a risk management strategy for secure AI adoption.
GV-3.2 **Policies and procedures are in place to define and differentiate roles and responsibilities for human-AI configurations and oversight of AI systems.**	• Cyera supports this control by cataloging all human and non-human identities with access to your data estate, determining their access privileges. This information enables enforcement of role-based access controls for human and AI users, and provides visibility into which users have access to both sensitive data and AI applications.

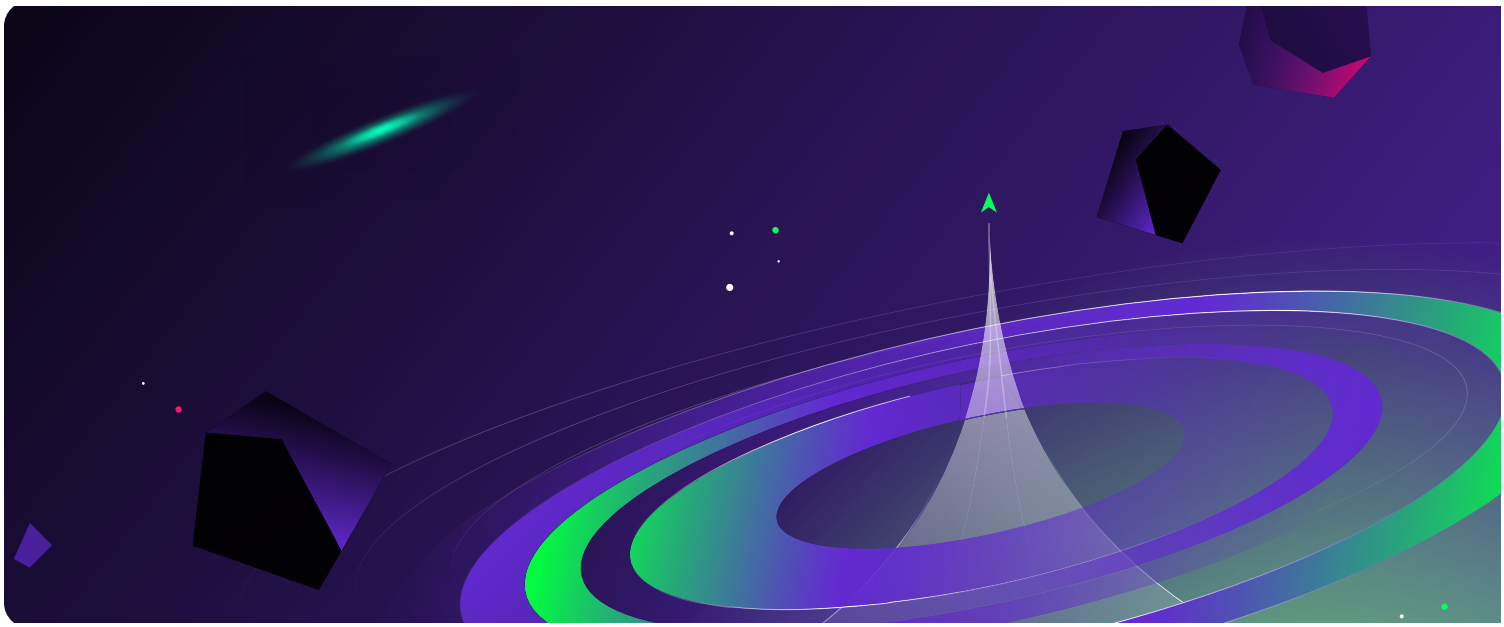

Governance Control	How Cyera Supports it
GV-4.1 **Organizational policies and practices are in place to foster a critical thinking and safety-first mindset in the design, development, deployment, and uses of AI systems to minimize potential negative impacts.**	• Cyera continuously monitors your data estate, enforcing responsible AI use by blocking sensitive prompts and redacting regulated data from Copilot interactions. Cyera monitors AI outputs for policy violations and classifies AI-generated data, making it easier to audit them. • Cyera helps raise awareness of data security issues among AI actors. When a policy violation occurs, Cyera notifies the data owner via email, Slack, or other channels, providing instructions for remediation.
GV-4.2 **Organizational teams document the risks and potential impacts of the AI technology they design, develop, deploy, evaluate, and use, and they communicate about the impacts more broadly.**	• Cyera can inform impact assessments related to AI system development and deployment. Cyera's AI Risk Assessment and Breach Readiness services now simulate AI misuse scenarios (e.g., prompt injection, inferencing violations) and support tabletop exercises, enabling rigorous pre-deployment and post-deployment risk planning. • Additionally, Cyera integrates with workflow tools for automated incident response.
GV-4.3 **Organizational practices are in place to enable AI testing, identification of incidents, and information sharing.**	
GV-6.1 **Policies and procedures are in place that address AI risks associated with third-party entities, including risks of infringement of a third-party's intellectual property or other rights.**	• Cyera's data discovery capabilities identify third-party AI applications within your IT ecosystem. In addition to out-of-the-box classifiers, Cyera's AI-native design learns classifications unique to your environment, using deep context to identify sensitive categories of data - including intellectual property - with 95 percent precision. • By leveraging Cyera's AISPM and AI Shield, you gain the visibility and control needed to prevent intellectual property or other sensitive data from entering AI training datasets or being ingested by AI systems.
GV-6.2 **Contingency processes are in place to handle failures or incidents in third-party data or AI systems deemed to be high-risk.**	• Finally, Cyera's advanced services - including Data Risk Assessment, AI Risk Assessment, and Breach Readiness services - help you gauge data risks posed by third party AI products and plan your incident response strategy.

Map

Map Control	How Cyera Supports it
MP-1.1 **Intended purposes, potentially beneficial uses, context specific laws, norms and expectations, and prospective settings in which the AI system will be deployed are understood and documented. Considerations include: the specific set or types of users along with their expectations; potential positive and negative impacts of system uses to individuals, communities, organizations, society, and the planet; assumptions and related limitations about AI system purposes, uses, and risks across the development or product AI lifecycle; and related TEVV and system metrics.**	• This category requires organizations to consider unintended consequences of deploying AI systems, as even accurate systems can be misused, or used in unforeseen contexts. Therefore, designing AI systems with built-in guardrails is essential. • Cyera's AI Guardian supports responsible AI deployment by identifying AI systems and users, integrating with DLP tools to monitor session usage, and linking alerts to AI assets. • Cyera's Runtime Protection enables prompt-level monitoring of how AI is used and what data it accesses. It blocks or redacts inappropriate uses, and supports human oversight and documentation for impact reviews.
MP-1.3 **The organization's mission and relevant goals for AI technology are understood and documented.** **MP-1.4** **The business value or context of business use has been clearly defined or – in the case of assessing existing AI systems – re-evaluated.**	• These categories emphasize transparency and aligning AI development and use with organizational and social values. Cyera supports these goals by providing AI developers and deployers visibility and control over the data used by AI systems, mapping the flow of data from human and AI users to AI tools. This enables auditing and explanation of AI outputs, preventing them from exceeding organizational guardrails. • Cyera's Runtime Protection enables prompt-level monitoring of how AI is used and what data it accesses. It blocks or redacts inappropriate uses, and supports human oversight and documentation for impact assessments.

Map Control	How Cyera Supports it
MP-2.2 **Information about the AI system's knowledge limits and how system output may be utilized and overseen by humans is documented. Documentation provides sufficient information to assist relevant AI Actors when making decisions and taking subsequent actions.**	• Cyera supports this category by tracing the lineage of data upstream from your deployed AI systems, and mapping the flow of AI-generated data to external entities.
MP-2.3 **Scientific integrity and TEVV considerations are identified and documented, including those related to experimental design, data collection and selection (e.g., availability, representativeness, suitability), system trustworthiness, and construct validation. especially when the information comes from multiple (or unknown) sources.**	• By identifying and tracing the lineage of AI training datasets, Cyera provides your organization with the necessary visibility into the data collection and curation process to evaluate the scientific integrity of AI models.
MP-3.2 **Potential costs, including non-monetary costs, which result from expected or realized AI errors or system functionality and trustworthiness – as connected to organizational risk tolerance – are examined and documented.**	• Cyera's Data Risk Assessment, AI Risk Assessment, and Breach Readiness services help your organization better understand the risks to your data estate. This includes assessing the qualitative and quantitative costs of an AI system failure in your operating environment and evaluating the adequacy of your response plans. • Cyera tracks data lineage from ingestion through output to assess potential misuse or harm from AI outputs. This data supports quantitative impact models required by NIST for residual risk evaluation.

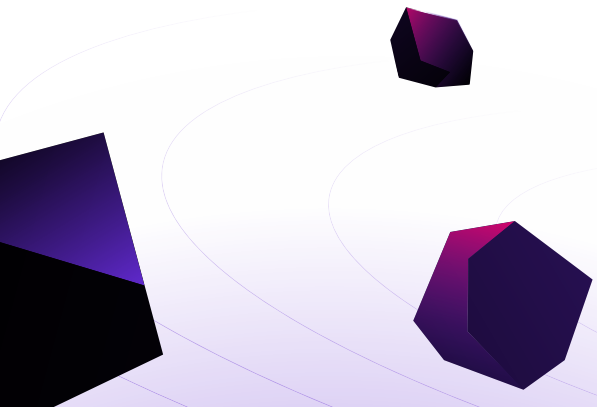

Map Control	How Cyera Supports it
MP-3.3 **Targeted application scope is specified and documented based on the system's capability, established context, and AI system categorization.**	• This category aims to narrow the deployment context of AI to better manage associated risks. Cyera assists by identifying entities with access to AI tools, enabling organizations to restrict access to specific users, business silos, or geographic regions.
MP-3.4 **Processes for operator and practitioner proficiency with AI system performance and trustworthiness – and relevant technical standards and certifications – are defined, assessed, and documented.**	• This category addresses the need to vet AI system users, ensuring they possess the requisite skills and knowledge to appropriately and responsibly interface with AI and handle AI outputs. • Cyera assists by cataloging identities with access to AI systems and by monitoring for risky usage, such as prompts that contain sensitive data or outputs that violate organizational acceptable use policies.
MP-3.5 **Processes for human oversight are defined, assessed, and documented in accordance with organizational policies from GOVERN function.**	• This category concerns AI system oversight, including testing AI systems in secure environments before deployment and training users on best practices. • Cyera supports this requirement by identifying data migration from test/ development to production environments. Upon policy violations, Cyera alerts data owners via email, Slack, or other channels with instructions for remediation, and integrates with workflow tools for automated remediation.

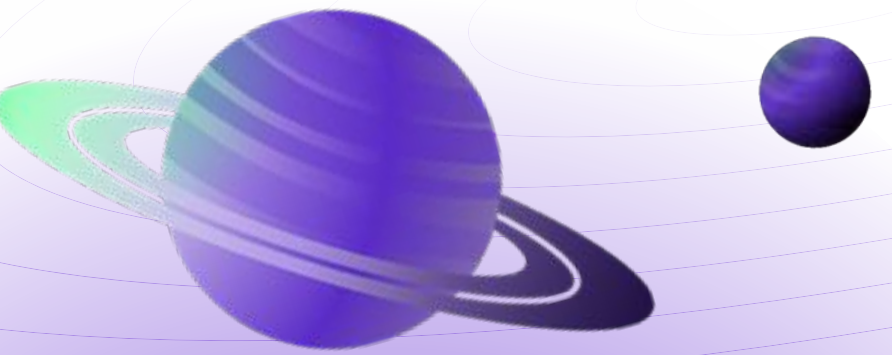

Map Control	How Cyera Supports it
MP-4.1 **Approaches for mapping AI technology and legal risks of its components – including the use of third-party data or software – are in place, followed, and documented, as are risks of infringement of a third-party's intellectual property or other rights.**	• Cyera understands your organization's data from its context as well as its content, enabling it to identify sensitive data categories like trade secrets or intellectual property. • Cyera's data discovery and classification capabilities also provide visibility into your data's location, who's using it, and for what. This information enables you to build an inventory of all third-party AI technologies in use in your IT ecosystem. • Cyera's AISPM detects and monitors third-party integrations (e.g., Copilot, browser AI tools, AI chat interfaces), ensuring that AI components used in your stack are evaluated and governed per AI RMF guidelines.
MP-4.2 **Internal risk controls for components of the AI system, including third-party AI technologies, are identified and documented.**	
MP-5.1 **Likelihood and magnitude of each identified impact (both potentially beneficial and harmful) based on expected use, past uses of AI systems in similar contexts, public incident reports, feedback from those external to the team that developed or deployed the AI system, or other data are identified and documented.**	• Insights gleaned from data discovery, classification, and minimization efforts with Cyera's DSPM and AISPM, along with its Data Risk Assessment service, help your organization better understand the likelihood and magnitude of beneficial or harmful impacts from deploying AI systems. • Cyera clarifies your data security posture, the extent of your attack surface, and the greatest risks to data privacy and data security, giving you the confidence to safely and securely leverage your most valuable data to drive business growth,

Measurment

Measurment Control	How Cyera Supports it
MS-1.1 **Approaches and metrics for measurement of AI risks enumerated during the MAP function are selected for implementation starting with the most significant AI risks. The risks or trustworthiness characteristics that will not – or cannot – be measured are properly documented.**	• Cyera can help identify and mitigate data drift, which can erode the appropriateness and effectiveness of AI metrics and controls. • Cyera's Data Risk Assessment and AI Risk Assessment services help you identify the biggest risks to your data estate, enabling prioritized risk mitigation and minimization of your attack surface before deploying AI systems that could exacerbate underlying data risks.
MS-1.2 **Appropriateness of AI metrics and effectiveness of existing controls are regularly assessed and updated, including reports of errors and potential impacts on affected communities.**	• Cyera's AISPM identifies Shadow AI applications, and its AI Runtime Protection tracks AI inputs and outputs to uncover sensitive data exposure and prevent misuse. Reports generated by Cyera help assess the effectiveness of controls over time.
MS-1.3 **Internal experts who did not serve as front-line developers for the system and/or independent assessors are involved in regular assessments and updates. Domain experts, users, AI Actors external to the team that developed or deployed the AI system, and affected communities are consulted in support of assessments as necessary per organizational risk tolerance.**	• Cyera MCP dashboards now offer AI-specific views showing policy violations, prompt behavior anomalies, and AI-inferred data risks, providing structured metrics for AI trustworthiness.
MS-2.1 **Test sets, metrics, and details about the tools used during TEVV are documented.**	• Cyera helps your organization identify validation and test data sets for TEVV documentation

Measurment Control	How Cyera Supports it
MS-2.2 **Evaluations involving human subjects meet applicable requirements (including human subject protection) and are representative of the relevant population.**	• Cyera automatically maps AI-generated data to various compliance frameworks such as GDPR, enabling your organization to verify compliance with human data subjects' privacy rights.
MS-2.3 **AI system performance or assurance criteria are measured qualitatively or quantitatively and demonstrated for conditions similar to deployment setting(s). Measures are documented.**	• Cyera supports these categories by identifying AI training datasets, enabling your organization to more accurately measure the representativeness of training data relative to the operational environment over time.
MS-2.4 **The functionality and behavior of the AI system and its components – as identified in the MAP function – are monitored when in production.**	• Cyera's Runtime Protection logs unsafe prompts, detects AI system drift (e.g., when AI outputs deviate from original purpose), and auto-remediates content leaks—ensuring AI systems remain trustworthy and within safety bounds.
MS-2.5 The AI system to be deployed is demonstrated to be valid and reliable. Limitations of the generalizability beyond the conditions under which the technology was developed are documented.	• Cyera assists with these categories by monitoring input data and alerting on and logging policy violations.
MS-2.6 The AI system is evaluated regularly for safety risks – as identified in the MAP function. The AI system to be deployed is demonstrated to be safe, its residual negative risk does not exceed the risk tolerance, and it can fail safely, particularly if made to operate beyond its knowledge limits. Safety metrics reflect system reliability and robustness, real-time monitoring, and response times for AI system failures.	• Cyera's AI Runtime Protection can block sensitive or dangerous prompts and mask sensitive outputs. It also integrates with workflow tools to facilitate automated incident response.
MS-2.7 AI system security and resilience – as identified in the MAP function – are evaluated and documented.	

Measurment Control	How Cyera Supports it
MS-2.9 **The AI model is explained, validated, and documented, and AI system output is interpreted within its context – as identified in the MAP function – to inform responsible use and governance.**	• Cyera assists by classifying and tracing the lineage of data ingested by AI models.
MS-2.10 **Privacy risk of the AI system – as identified in the MAP function – is examined and documented.**	• Cyera discovers and classifies PII across your data estate, including within AI training datasets. Cyera's Omni DLP now applies real-time privacy controls on AI inputs and outputs, including redaction, prompt blocking, and enforcement of GDPR and HIPAA rules for PII/PHI handling in AI environments. • Audit logs help track privacy risks from specific systems and users, while Cyera's Data Risk Assessment and AI Risk Assessment services helps you better understand the nature and scope of risks that AI poses to PII in your data estate.
MS-2.11 **Fairness and bias – as identified in the MAP function – are evaluated and results are documented.**	• Cyera's data discovery and classification capability help you better understand training and TEVV datasets. With Cyera's insights, you can develop a clearer picture of the completeness, representativeness, and balance of data sources.
MS-2.12 **Environmental impact and sustainability of AI model training and management activities – as identified in the MAP function – are assessed and documented.**	• Leveraging Cyera's DSPM for data minimization efforts can significantly reduce your organization's energy consumption and carbon footprint. Many Cyera customers save tens of thousands of dollars per month on data storage costs, cutting hundreds of thousands of kilowatt hours per year, and hundreds of metric tons of carbon dioxide.

Measurment Control	How Cyera Supports it
MS-2.13 **Effectiveness of the employed TEVV metrics and processes in the MEASURE function are evaluated and documented.**	• This category focuses on continuous improvement, including corrective actions to enhance the quality, accuracy, reliability, and representativeness of the data fed into AI models. Cyera assists by identifying and classifying all data fed into and generated by AI systems.
MS-3.1 **Approaches, personnel, and documentation are in place to regularly identify and track existing, unanticipated, and emergent AI risks based on factors such as intended and actual performance in deployed contexts**	• Cyera's AI Risk Assessment service helps identify potential data security and compliance risks related to the use of AI. Reports generated by Cyera help track AI risks by highlighting the most frequent policy violations and riskiest users.

Manage

Manage Control	How Cyera Supports it
MG-1.2 **Treatment of documented AI risks is prioritized based on impact, likelihood, and available resources or methods.**	• Cyera's AI Risk Assessment service leverages its DSPM in conjunction with virtual CISO-led exercises to evaluate your organization's data security posture relative to 30 different controls from frameworks such as ISO 42001 and the NIST AI Risk Management Framework. This service can help you identify and prioritize the most serious risks relating to the use of AI systems.
MG-1.3 **Responses to the AI risks deemed high priority, as identified by the MAP function, are developed, planned, and documented. Risk response options can include mitigating, transferring, avoiding, or accepting.**	
MG-1.4 **Negative residual risks (defined as the sum of all unmitigated risks) to both downstream acquirers of AI systems and end users are documented.**	
MG-2.2 **Mechanisms are in place and applied to sustain the value of deployed AI systems.**	• Cyera continuously monitors your data estate, identifying, classifying, and applying DLP policies to AI-generated data. It also integrates with workflow tools to support automated remediation, sending alerts to data owners via email, Slack, or other channels with instructions for remediation.
MG-2.3 **Procedures are followed to respond to and recover from a previously unknown risk when it is identified.**	
MG-2.4 **Mechanisms are in place and applied, and responsibilities are assigned and understood, to supersede, disengage, or deactivate AI systems that demonstrate performance or outcomes inconsistent with intended use.**	• Cyera catalogs all identities with access to your data estate, including AI entities. It also identifies which identities pose the greatest risk to your data, helping to determine whether an AI system is beginning to behave in a riskier fashion that may require human intervention.

Manage Control	How Cyera Supports it
MG-3.1 **AI risks and benefits from third-party resources are regularly monitored, and risk controls are applied and documented.**	• Cyera identifies and classifies your data wherever it resides, allowing you to build an inventory of AI applications in your IT ecosystem, as well as an inventory of human and non-human users with access to your data.
MG-3.2 **Pre-trained models which are used for development are monitored as part of AI system regular monitoring and maintenance.**	• This visibility and control allows you to continuously monitor third-party AI tools such as Microsoft Copilot or ChatGPT, and take actions such as blocking sensitive inputs or masking data in AI outputs that violate privacy regulations or organizational policies.
MG-4.1 **Post-deployment AI system monitoring plans are implemented, including mechanisms for capturing and evaluating input from users and other relevant AI Actors, appeal and override, decommissioning, incident response, recovery, and change management.**	• Cyera continuously monitors your data estate, providing insights that can be leveraged to assess AI's trustworthiness over time.
MG-4.2 **Measurable activities for continual improvements are integrated into AI system updates and include regular engagement with interested parties, including relevant AI Actors.**	• For example, Cyera detects dataset modifications and monitors human interactions with AI, blocking sensitive prompts and classifying and protecting regulated data in AI outputs. Cyera also identifies AI entities with overbroad permissions. Audit reports and event logs can inform impact analyses, post-deployment validation, or decommissioning.
MG-4.3 **Incidents and errors are communicated to relevant AI Actors, including affected communities. Processes for tracking, responding to, and recovering from incidents and errors are followed and documented.**	

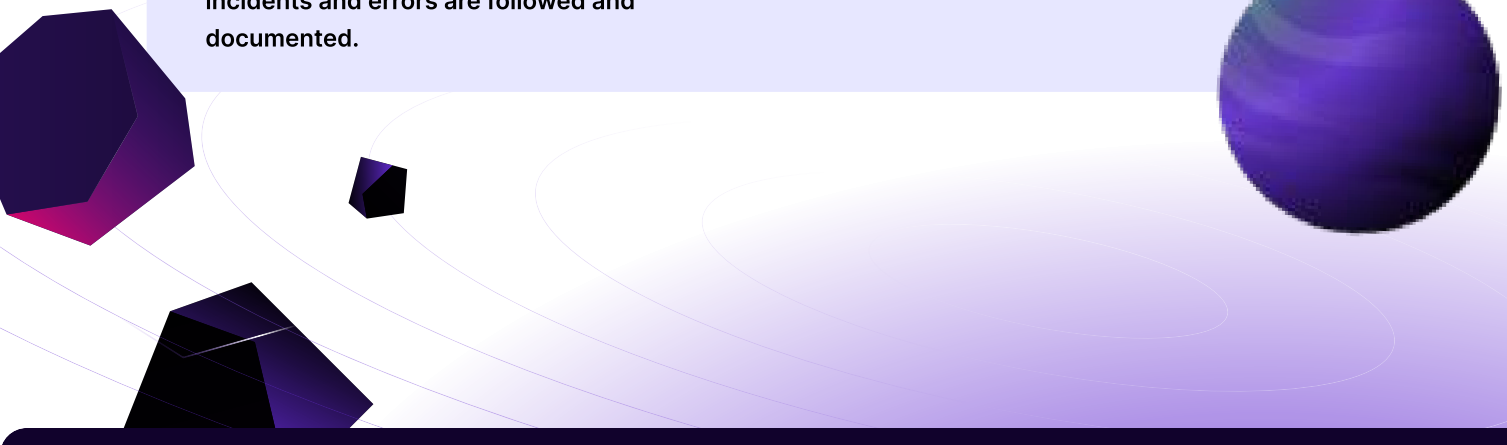


CYERA