AI Inventories: Practical Challenges for Organizational Risk Management

Co-Authored by Kent Sokoloff (Chevron) and Hadassah Drukarch, Sez Harmon, and Patrick McAndrew (Responsible AI Institute)

AI

the human energy company®

www.responsible.ai

1. Introduction

For organizations committed to harnessing the potential of AI or considering its integration into their operations, maintaining a comprehensive inventory of AI use cases is essential. Such an inventory enables effective risk management and helps ensure financial oversight, as AI development and deployment can be resource intensive. In simple terms, you cannot control what you do not know exists within your organization.

This guide aims to provide business leaders and practitioners with practical insights into the following key areas:

1. **Understanding AI Use Case Inventories:** This guide will clarify what an AI use case inventory is and explain why creating one is more complex than a traditional IT application inventory.

2. **Challenges in Managing AI Inventories:** As AI technology and adoption rapidly accelerate, maintaining an up-to-date inventory has become increasingly difficult for medium and large organizations. This guide explores how these changes present unique barriers at a time when tracking AI use is more critical than ever.

3. **Designing a Practical AI Inventory:** This guide outlines what an effective AI inventory could look like and discusses how it plays a vital role in mitigating risks and managing financial impacts.

4. **The Role of Regulations and Standards:** Regulations and standards can provide valuable frameworks for making AI inventory management more achievable. This guide analyzes how these can support organizations in overcoming inventory management challenges.

By exploring these four areas, this guide equips industry professionals with the knowledge needed to build and maintain successful AI inventories. Opportunities for regulatory bodies to improve AI inventory management processes across the U.S. are also highlighted.

2. AI Use Cases & Inventories

2.1 Defining an AI Use Case

Maintaining an IT application inventory is a well-established practice in medium to large organizations. Typically, companies track vendor-provided software and internally developed applications and their outputs through internal review processes.[1] While there may be variations in the level of detail included, most IT inventories consistently capture core elements such as software code, programs, documentation, hardware, network resources, and user materials.[2] However, even with IT inventories, challenges can arise especially when dealing with edge cases where new applications are built on identified applications in the inventory.

For instance, in a typical IT setting, a commonly used tool like Excel can serve as a base for specialized applications, such as a custom-built spreadsheet designed to manage a particular financial issue. This scenario raises a critical question: should the inventory track Excel itself as the application, or should it also track the specialized spreadsheet as a unique entry? While large organizations may include both in their inventories, medium-sized organizations often struggle with defining the boundaries. This same ambiguity carries over when developing AI use case inventories but becomes even more pronounced due to the unique characteristics of AI systems.

AI applications differ from conventional IT systems in several ways. One of the key distinctions is the broad adaptability of AI algorithms and large language models (LLMs), which are often not limited to specific domains. For example, an AI model initially designed to analyze geospatial data for oil drilling could also be applied to assess patent opportunities. This versatility complicates the inventory process, as the value and risk of an AI model cannot be easily tracked at the algorithm level alone. Tracking the algorithm itself does not provide sufficient insight into how it is applied, what risks it posesses, or the economic impact it might have.

To manage AI effectively, it is incumbent on organizations to move beyond a simple cataloging of algorithms and instead focus on documenting AI use cases. By capturing the full context of each AI deployment, companies can better understand the specific risks, challenges, and opportunities presented by each use. An AI use case, in this context, is defined not just by the algorithm or model being used, but by how the AI is applied to address a specific business problem or opportunity.

For the purposes of this guide, **AI use case** is defined by three critical factors:

1. **Purpose of Use:** This refers to the challenge or opportunity the AI system is designed to address.[5] Understanding the problem the AI is solving within the organization helps to clarify its role, potential value, and associated risks.

2. **Model Type:** The type of AI model deployed is a key component of the use case. Whether it is a machine learning algorithm, natural language processing model, or another type of AI system, knowing the model type provides a baseline for assessing its behavior and performance.

3. **Scope of Data:** The datasets used to train and validate the AI model must be tracked. AI systems rely on vast amounts of data to function, and the source, quality, and sensitivity of that data can have significant implications for compliance, security, and ethical considerations.

By capturing key elements such as purpose of use, model type, and data scope, organizations can build an AI use case inventory that provides a more detailed and actionable understanding of how AI is integrated into their operations. However, identifying specific AI use cases within an organization can be particularly challenging. Despite this difficulty, it is a crucial step in elevating the maturity of your AI governance program. A well-structured AI use case inventory serves as the foundation for managing operational and financial risks, offering leadership greater visibility into AI's evolving role in the company. It also can help the organization comply with emerging regulations requiring transparency and accountability in the use of AI systems.

This shift from merely tracking algorithms to building use case-based inventories represents a significant evolution in managing organizational risk. As AI adoption accelerates, the complexity of identifying and documenting AI use cases may grow, but it remains essential to strengthening AI governance and ensuring that the benefits of AI are realized without introducing undue risks.[6] Creating a comprehensive and robust AI use case inventory is not just about risk management; it's a critical step in aligning AI systems with business objectives and regulatory expectations.

1. "IT Inventories Definition."
2. "IT Inventories Definition."
3. von Hollen, "The Importance of Accurate IT Inventory."
4. Rana, "Generic LLMs vs. Domain-Specific LLMs: What's the Difference?"
5. "AI Guide for Government: A Living and Evolving Guide to the Application of Artificial Intelligence for the U.S. Federal Government."

2.2. The Challenge of Comprehensive AI Inventories

The rapidly evolving AI landscape has made it increasingly challenging to maintain a comprehensive inventory of AI use cases within an organization. **The process of identifying and cataloging AI use cases is not as straightforward as it may seem, and the complexity grows as organizations scale their AI operations across different departments and business units.** As AI technologies advance, the ability to track and manage every instance of AI deployment becomes more complex, and in some cases, impractical.

The introduction of OpenAI's ChatGPT-3 in 2022 marked a significant turning point in the AI technology space.[7] Its release, followed by further advancements like ChatGPT-4o, Microsoft Copilot, Llama 3.1and Gemini, dramatically expanded AI capabilities, placing powerful AI tools in the hands of non-technical users. These tools driven by LLMs enable anyone not just data scientists to perform complex tasks such as querying datasets, generating reports, creating images, and even making predictions with minimal effort.

While this democratization of AI brings exciting opportunities for businesses, it also introduces new challenges in identifying, controlling, and inventorying AI use cases. Without proper oversight, anyone in an organization could potentially use these AI tools to query sensitive datasets and generate outputs that bypass traditional data governance controls. For example, a marketing team might use a tool to draft customer insights, while an HR department could rely on it to plan team-building events. However, the ease of access to such tools also opens the door to more risky uses, such as unintentionally exposing restricted data through AI prompts or sharing sensitive outputs without proper security measures.

This risk is particularly pronounced when it comes to data-related issues. For instance:

- Restricted data could be used to train an AI model and inadvertently become part of public outputs, leading to potential data breaches.
- Restricted data could be entered into AI prompts, which might then be used to train models, again raising the risk of unauthorized data exposure.
- Even if restricted data does not make it to the public domain, it could still be shared within an organization without proper access controls, leading to sensitive information reaching individuals who should not have access to it.

The improved access to these powerful AI models creates a high-risk environment, making it difficult to track every AI use case across the organization. While it may be acceptable for an HR team to use an AI chatbot to brainstorm event ideas, using the same tool to process HIPAA-protected health data could represent a serious breach of privacy. The pervasive availability and affordability of these tools complicate the process of identifying, tracking, and cataloging every AI use case, significantly increasing the potential risks to the organization.

6. Evans, "What Is an AI Inventory, and Why Do You Need One?"
7. Metz, "OpenAI Unveils New ChatGPT That Can Reason Through Math and Science."
8. "Introducing GPT-4o and More Tools to ChatGPT Free Users."
9. "Empower Your Organization with Copilot."
10. "Meet Llama 3.1."
11. "Gemini: Supercharge Your Creativity and Productivity."

For more, refer to this case study in flexibility and risk: Empower Your Organization with Copilot

The introduction of generative AI tools presents a fundamental challenge to the very concept of comprehensive AI inventories. As these tools become more widely adopted across organizations, they blur the lines between formal AI deployments and ad-hoc, user-driven AI interactions. This shift makes it increasingly difficult to maintain a fully accurate inventory of AI use cases.

This challenge is exacerbated by the fact that generative AI tools are constantly evolving, with new features and capabilities being added regularly. As these tools continue to improve, they are likely to become even more embedded in everyday business operations, further complicating the task of tracking their usage. In the absence of strict policies and robust tracking mechanisms, organizations risk falling behind in their ability to manage AI risks effectively.

2.3. Who is Selling You Applications with AI-Integrations?

One of the most significant challenges organizations face when building an AI inventory is tracking AI integrations introduced by third-party providers. Vendors may embed AI functionality into their products or services without the recipient organization's full awareness, making it difficult to track these systems and their associated risks.[12] This challenge is not limited to AI-specific service providers but extends to any third-party vendor that leverages AI in their offerings. The problem can be illustrated through several distinct scenarios that expose the complexities of managing AI use cases introduced through third parties.

Opaque AI Operations from Subcontracted Services

In many cases, an AI vendor may offer a specialized AI analysis service but subcontract portions of that service to a third party, making the system's operations opaque even to the primary vendor. For example, imagine an AI vendor providing fraud detection services to a financial institution. Unknown to the institution, the vendor relies on a third-party company to handle parts of the AI analysis process, using models or data sources that are not fully transparent. The recipient organization remains unaware of how these subcontracted AI systems operate, what data they use, or how they manage potential risks.

In this scenario, the organization's AI inventory would be incomplete because it would fail to capture the full scope of the third-party AI integrations. This lack of transparency creates blind spots, making it difficult to manage operational risks or ensure compliance with data governance policies.

Non-AI Vendors Leveraging AI in Products

Another common challenge occurs when vendors outside of the AI services industry integrate AI into their products without explicit disclosure. For instance, an educational program developer might leverage a publicly available AI image-generation system to create visual content for their organization's e-learning platform. In this case, the educational product itself is not marketed as AI-driven, but the AI system used to generate content plays a critical role. If the organization is not aware of the AI component, it may fail to include this AI use case in its inventory.

Without tracking these hidden AI integrations, organizations may overlook the potential risks associated with AI-generated content, such as copyright violations, inappropriate use of sensitive data, or even ethical concerns over the authenticity and quality of AI-produced materials.

Subcontracted AI Use with Data Leakage Risks

In yet another example, a consultant hired to develop an educational program may subcontract a company that uses AI to generate voice recordings. If this subcontracted company uses an open-source AI platform to create or modify voice recordings, there are potential risks that may go unnoticed. What happens if the voice of an executive or employee used in the program ends up in the public domain due to inadequate control over the AI's data-handling processes?

Similarly, it can be nearly impossible for an organization to know whether third-party software developers use tools on highly proprietary projects. These tools, which assist in code writing through AI-generated suggestions, could expose confidential project data without the company's knowledge. It might also risk the intellectual property position of your company if the code is largely the product of AI-generated activities.

These scenarios illustrate a fundamental question for organizations: should all these AI integrations be treated as use cases that need to be inventoried? And more importantly, can the organization realistically track and document all such instances? The sheer complexity of monitoring every third-party AI integration, especially when vendors do not always disclose the use of AI, can make it seem like an overwhelming task.

Beyond these vendor-specific complications, the growing ubiquity of AI in everyday products and services adds to the challenge. AI is now embedded in a wide range of technologies, from the soda machine that uses AI to predict restocking needs, to your coffee-ordering app that recommends your favorite drinks. As AI capabilities become more integrated into even mundane technologies, organizations are faced with a daunting task: deciding how to capture and manage the vast array of AI functionalities scattered across their operations. Building an all-encompassing AI inventory could quickly become unmanageable, requiring significant resources and constant updates to keep pace with evolving AI integrations.

While the task of building a complete AI inventory may seem overwhelming, organizations do not need to give up or wring their hands in frustration. Instead, they can adopt a practical approach that focuses on building a robust AI inventory that delivers real value and addresses the key issues of risk management and fiscal responsibility.

Rather than trying to inventory every single AI integration organizations can prioritize AI use cases based on risk, impact, and visibility. This approach involves focusing on areas where AI is used in ways that directly affect critical business operations, data security, compliance, or customer-facing activities. By concentrating on high-impact AI use cases, organizations can ensure they manage the most significant risks while maintaining a reasonable scope for their inventory efforts.

As a best practice to achieve this, organizations should:

- **Establish clear criteria for what constitutes an AI use case that needs to be inventoried.** Not every AI functionality requires the same level of scrutiny. AI used in critical business functions, such as decision-making, critical public infrastructure for energy companies, customer interactions, or data analysis, should take priority in the inventory.

12. Brown, "Third-Party AI Tools Pose Increasing Risks for Organizations."

- **Require transparency and disclosure from third-party vendors.** Organizations should work with vendors to ensure that any AI integration, even those that are embedded or subcontracted, is clearly disclosed. A best practice would be to have vendors contractually obligated to provide information on how their AI models operate, what data they use, and how risks are managed.

- **Monitor high-risk AI use cases**. For vendors that subcontract AI services or use public AI models, as a best practice, organizations should monitor these integrations. This could include regular audits of the vendor's AI operations or establishing protocols for tracking the flow of sensitive data.

- **Focus on dynamic and high-value AI systems.** Instead of trying to track AI in every soda machine or app, organizations can focus on AI systems that evolve, learn, and directly influence decision-making. Prioritizing high-value and high-risk AI use cases ensures that inventory efforts are manageable and impactful.

By taking this strategic approach, organizations can build an AI inventory that not only keeps them compliant with regulatory requirements but also ensures that they manage the most critical AI risks in their operations. While it may not be realistic to inventory every AI system embedded in third-party products, organizations can still create a robust framework for managing AI use cases that matter most to their risk and fiscal management strategies.

3. How Does Risk Play into This?

3.1 Identifying Risk

If one of the primary goals of maintaining an AI use case inventory is managing risk, the next logical question is: **What types of risks need to be prioritized?** The answer depends on the sector in which your organization operates, as AI risks can vary significantly across industries. In the energy sector, for instance, priority might be given to risks associated with the physical safety of employees and the community, while in the financial services sector, the focus could be on ensuring AI systems do not lead to biased decisions, such as unfairly denying loan applications to protected classes.

AI risks are diverse and cut across sectors, meaning there is no magical "aha" moment when identifying potential dangers. Instead, as a best practice, organizations should consider adopting a comprehensive approach to risk identification, integrating sector-specific concerns into their overall risk management framework. Several open-source frameworks offer valuable insights into AI risks and can be used to guide this process. Some key resources include:

- NIST AI Risk Management Framework
- EU AI Act
- ISO 42001
- Microsoft Responsible AI Principles and Approach
- MIT AI Risk Repository

Each of these frameworks highlights various high-risk consequences of AI applications, helping organizations assess which AI use cases require closer scrutiny. For example, in the energy sector, risks might include threats to human safety, environmental hazards, and challenges in validating the accuracy or safety of AI-generated results by subject matter experts (SMEs). Additionally, cascading failures from poor critical infrastructure security, misuse of sensitive customer data, or disruptions caused by inaccurate AI recommendations could all have far-reaching consequences.

To illustrate, here are some potential **high-risk areas in industrial sectors** that constitute best practice in AI inventories:

- **Risks to human safety and the environment**, from ineffective or incorrect AI systems that assist in the management of critical safety protocols.
- **Difficult-to-validate results,** where SMEs may struggle to assess the accuracy or safety of AI-driven decisions in complex energy systems.
- **Dependencies and cascading failures** from vulnerabilities in critical infrastructure security.
- **Mismanagement of assets**, stemming from insufficiently tested AI models used in system development.
- **Misuse of sensitive customer data,** either internally or externally, leading to unethical or illegal decision-making.
- **Market manipulation**, including threats like data poisoning or model manipulation that could distort energy pricing or supply.
- **Biased recommendations,** resulting from models that are not sufficiently tailored to specific industrial sector use cases.
- **Misidentification of human errors,** particularly in complex systems like smart grids, which rely on AI for monitoring and management.
- **Disruptions to company operations**, due to inaccurate or untrustworthy AI executions or prompt injections that impact essential functions.

In sectors where identifying every single AI use case may not be feasible, consider a best practice of prioritizing an inventory of use cases that fall into these high-risk categories. By focusing on AI systems that have the most potential to cause harm—whether through safety risks, ethical breaches, or operational disruptions—organizations can ensure their AI inventories provide meaningful insight and support effective risk management.

Ultimately, understanding the specific risks associated with AI use in your industry allows for a targeted approach to inventory management, ensuring that the most critical risks are identified, assessed and mitigated.

3.2. Risk Mitigation

While this guide does not delve deeply into specific methods of risk mitigation, it is crucial to recognize that effective organizations must implement a comprehensive risk assessment and mitigation process for AI systems. Without such protocols, organizations leave themselves vulnerable to various risks, including ethical violations, data security breaches, and operational failures. A well-structured risk mitigation approach addresses these concerns through several key strategies that ensure AI systems are deployed responsibly and safely.

One of the foundational aspects of mitigating AI risks involves **training the users of AI systems.** Employees and stakeholders need to be educated on AI's capabilities, limitations, and potential biases. This ensures users understand the broader implications of AI-driven decisions and are equipped to recognize when intervention may be necessary. Regular training programs also help foster a culture of responsible AI use, where individuals are aware of both the opportunities and risks that AI introduces.

Transparency is another critical component of effective risk mitigation. Organizations need to ensure that the use of AI is clearly communicated, both internally and externally, so that all stakeholders understand when and how AI is influencing decision-making processes. This transparency builds trust and allows for better scrutiny of AI outputs. It is particularly important when AI is used in sensitive areas like hiring, financial services, or healthcare, where the consequences of AI errors or biases can be severe.

Integrating a **"human-in-the-loop"** mechanism is a widely recommended strategy for AI risk mitigation. In this approach, AI-driven recommendations are not automatically accepted but are subject to human review and validation before any final decision is made. This layer of human oversight helps to catch potential errors, biases, or misinterpretations produced by AI models, ensuring that AI is used as a tool to assist decision-making rather than fully automating critical processes. Human-in-the-loop systems are especially important in sectors like finance, healthcare, and law, where AI outputs can have profound impacts on individuals and society.

When using **third-party AI systems**, risk mitigation becomes even more complex, and organizations must go beyond internal protocols. To manage these risks effectively, appropriate **contract language** is essential in partnerships with AI vendors. A best practice would be that contracts require third-party providers to disclose how their algorithms are trained, how they use the organization's data, and whether they subcontract any part of their AI services. This level of transparency helps organizations ensure that third-party AI systems align with their governance standards and risk management frameworks. Additionally, organizations need to understand whether external subcontractors are involved, as these secondary relationships can introduce further risks, including data security vulnerabilities or the use of unapproved algorithms.

In summary, risk mitigation for AI systems is a multi-faceted process that involves training, transparency, human oversight, and strong vendor agreements. By embedding these strategies into their AI governance frameworks, organizations can better manage the risks associated with AI use, especially when dealing with third-party providers. A proactive approach to risk assessment and mitigation ensures that AI systems are used responsibly and effectively, safeguarding both the organization and its stakeholders from unintended consequences.

4. Where to Draw the limit?

4.1 Decision-Making in AI Investment and Inventories

One of the most critical factors when deciding whether to greenlight an AI initiative is recognizing that, much like managing an AI inventory, these decisions are **use case-based** rather than application-based. This distinction is crucial because the value and risks associated with AI vary significantly depending on how the AI system will be used, rather than the specific technology itself. Whether evaluating a new internal AI project or acquiring a license for a third-party AI solution, the decision process should be grounded in a thorough understanding of how the AI will address a specific business need, solve a problem, or create value.

This process mirrors the approach organizations use to assess traditional IT projects, where the focus is on the specific goals and outcomes rather than the technology alone. When reviewing an AI initiative, business leaders should consider whether the proposed use case aligns with the organization's strategic objectives, whether it addresses a significant business challenge, and whether the benefits outweigh the costs and potential risks. This use case-centered approach ensures that AI investments are purposeful and drive measurable value.

To help organizations navigate the complexities of AI decision-making, the World Economic Forum recently published a report titled Unlocking Value from Generative AI, which provides valuable insights into how companies can effectively leverage AI for success. By adopting a use case-based approach to AI investment decisions, organizations can ensure that they are not just following trends but making deliberate, strategic choices about where and how to invest in AI. This method allows companies to prioritize projects that directly contribute to their success, while also maintaining control over the scope of their AI inventories and investments.

4.2 Special Considerations

When deciding whether to proceed with an AI project, there are several **granular considerations** that organizations should take into account beyond the high-level use case approach. These factors, although often implied in broader discussions, play a critical role in ensuring that AI investments are both valuable and sustainable. By addressing these considerations, organizations can improve the effectiveness of their decision-making process and mitigate potential risks.

Risk Assessment and Mitigation

While it may not always be explicitly stated, the World Economic Forum (WEF) strongly implies that any AI investment decision should include a thorough risk assessment. Risk is an inherent factor in AI projects, whether related to data privacy, security, operational failure, or ethical concerns like bias. It is a best practice to have a formal **risk assessment process** in place to evaluate the potential risks of an AI project before giving it the green light.

This risk assessment should consider:

- Data security and privacy risks: How sensitive is the data being used, and what are the risks if this data is exposed or misused?
- Operational risks: What is the potential impact of AI system failure or inaccuracies, particularly in critical functions?
- Ethical and legal risks: Does the AI project introduce risks of bias, unfair treatment, or violations of existing or emerging regulations?

Beyond identifying these risks, it also follows that organizations have a clear understanding of **risk mitigation strategies.** This may include deploying human-in-the-loop systems, establishing clear accountability measures, and ensuring that the AI model is transparent and explainable. Organizations that prioritize risk assessments and implement comprehensive mitigation approaches are better positioned to make informed AI investment decisions.

Avoiding Duplicative AI Projects

When the AI inventory is use case-based, there is a risk of green-lighting projects that, while addressing different business areas, are fundamentally duplicative in their core technology. For example, different departments may propose AI solutions for optimizing various processes— such as logistics, customer service, or supply chain management—that use similar underlying models, leading to unnecessary duplication of efforts.

To avoid this, a best practice is to involve skilled data scientists in the evaluation process. A data scientist can assess whether proposed AI projects are truly differentiated in terms of their business value or whether they are duplicative in nature. This helps ensure better prioritization of AI initiatives and avoids wasted resources. A data scientist's expertise can also help align AI projects with overall business goals, ensuring the technology is applied efficiently across different departments. By focusing on business value and avoiding duplication, organizations can ensure that they are leveraging AI technology to its fullest potential without overinvesting in redundant solutions.

Anticipating Regulatory Changes

The rapid pace of AI advancement is matched by the quickly evolving regulatory landscape. Internationally, governments and regulatory bodies are actively developing frameworks and guidelines to govern the ethical use of AI, with a growing focus on transparency, accountability, and fairness. Before embarking on an AI project, it is prudent for organizations to **anticipate where regulation may be heading** in their sector or region.

Asking hypothetical questions about the future of AI regulation helps organizations prepare for potential compliance challenges down the line. Some of the questions to consider are:

- What are the potential requirements for transparency in AI decision-making?
- How might data privacy laws impact the use of AI models that rely on personal or sensitive data?
- Could emerging laws on AI bias or fairness affect the deployment of certain AI systems?

Engaging in these forward-looking discussions helps ensure that AI investments are resilient to regulatory changes, preventing costly redesigns or compliance failures in the future. Furthermore, involving legal and compliance teams in the initial stages of AI project planning can help organizations navigate these regulatory challenges effectively.

5. Conclusion and Key Takeaways

Maintaining a comprehensive inventory of AI use cases is essential for ensuring the successful adoption and management of AI within an organization. As AI technologies become more versatile, it is increasingly important to focus on the use case rather than the underlying application, model, or algorithm. This use case-driven approach helps organizations identify where AI can offer the most value and, more importantly, where risks and data sensitivities are highest.

However, as AI becomes more ubiquitous, it may soon be unrealistic to maintain a complete inventory of every AI use case across an organization. Therefore, **prioritizing inventories in areas of high risk, high data sensitivity, and where economic returns are greatest.** This guide also outlines how organizations can **build a robust risk assessment and management program** and how **strategic decision-making can guide AI investments to maximize value and minimize risk.**

Launching or advancing an AI initiative within any organization requires a thoughtful and strategic approach to tracking how AI is used. This enables risk assessment, risk mitigation, and the ability to make informed financial decisions regarding AI investments. Given the complexity and rapid evolution of AI, there is a strong need for clear, industry-wide government regulations and standards to guide organizations on what to prioritize in their AI inventories and investments.

Internationally recognized standards from bodies like ISO and NIST can play a pivotal role in speeding up AI adoption by providing a common framework for addressing AI risks. Just as the SOC 2 compliance certification provided clarity for Software as a Service (SaaS) offerings, emerging standardized guidelines for AI can empower organizations to invest confidently in AI technologies while managing third-party vendor risks effectively.

Key Takeaways:

1. Use Case-Driven Inventories: Focus on cataloging AI use cases rather than applications, algorithms, or models. This approach helps ensure that the organization targets areas of high impact, risk, and opportunity.

2. Prioritize High-Risk Areas: Given the increasing ubiquity of AI, it is unrealistic to track every AI use case. Instead, concentrate on areas with high risk, high data sensitivity, and the potential for significant economic returns.

3 Risk Management and Investment: Develop a structured risk assessment and management program that informs decisions about AI investments, ensuring that they are aligned with business goals and mitigate potential pitfalls.

4. Regulatory Frameworks: Embrace internationally recognized standards and regulations to guide AI adoption. Standards from organizations like ISO and NIST can provide the clarity

By integrating these strategies, organizations can harness the full potential of AI while safeguarding against the risks associated with its deployment.

Bibliography

"AI Guide for Government:A Living and Evolving Guide to the Application of Artificial Intelligence for the U.S. Federal Government," Identifying AI Use Cases in Your Organization. IT Modernization: Centers of Excellence, 2024. https://coe.gsa.gov/coe/ai-guide-for-government/identifying-ai-use-cases-in-your organization/index.html#:~:text=A%20use%20case%20is%20a,opportunity%20that%20AI%20may%20solve.

"SOC 2® - SOC for Service Organizations:Trust Services Criteria," AICPA & CIMA, Association of International Certified Professional Accountants, 2024. https://www.aicpa-cima.com/topic/audit-assurance/audit-and-assurance-greater-than-soc-2.

Brown, Sara, "Third-Party AI Tools Pose Increasing Risks for Organizations," MIT Management Sloan School, Ideas Made to Matter: Artificial Intelligence (blog), 21 September 2023. https://mitsloan.mit.edu/ideas-made-to-matter/third-party-ai-tools-pose-increasing-risks-organizations.

Evans, Bex, "What Is an AI Inventory, and Why Do You Need One?" Onetrust (blog), 15 November 2023. https://www.onetrust.com/blog/what-is-an-ai-inventory-and-why-do-you-need-one/.

"Gemini: Supercharge Your Creativity and Productivity," Google, 2024, https://gemini.google.com/.

Hollen, Erik von, "The Importance of Accurate IT Inventory," UCSLogistics. UCSL Services: Articles (blog), 11 March 2024. https://www.ucslogistics.com/post/the-importance-of-accurate-it-inventory.

"IT Inventories Definition", Law Insider, 2024. https://www.lawinsider.com/dictionary/it-inventories.

"Meet Llama 3.1", Meta, 2024. https://www.llama.com/.

Metz, Cade. "OpenAI Unveils New ChatGPT That Can Reason Through Math and Science. " The New York Times, 12 September 2024. https://www.nytimes.com/2024/09/12/technology/openai-chatgpt-math.html.

"Empower Your Organization with Copilot." Microsoft, 2024. https://www.microsoft.com/en/microsoft-copilot.

"Introducing GPT-4o and More Tools to ChatGPT Free Users." Microsoft, 13 May 2024. https://openai.com/index/gpt-4o-and-more-tools-to-chatgpt-free/.

Rana, Hiral. "Generic LLMs vs. Domain-Specific LLMs: What's the Difference?" Dataversity: Smart Data News, Articles, & Education (blog), 10 May 2024. https://www.dataversity.net/generic-llms-vs-domain-specific-llms-whats-the-difference/.

About the Authors

Kent Sokoloff, Ph.D., is a Senior Data Architect at Chevron where he helped launch the Responsible AI (RAI) program and is a member of the RAI Operations Team. Kent has extensive experience in data and information management for machine learning and artificial intelligence algorithms.

Chevron is a leading international energy company and a member of the Responsible AI Institute.

About Responsible AI Institute

Since 2016, Responsible AI Institute (RAI Institute) has been at the forefront of advancing responsible AI adoption across industries. As a non-profi t organization, RAI Institute partners with policymakers, industry leaders, and technology providers to develop responsible AI benchmarks, governance frameworks, and best practices. RAI Institute equips organizations with expert-led training, real-time assessments, and implementation toolkits to strengthen AI governance, enhance transparency, and drive innovation at scale.

Members include leading companies such as Boston Consulting Group, Genpact, KPMG, Kennedys, Ally, ATB Financial and many others dedicated to bringing responsible AI to all industry sectors.

Ready to get started? Become a Responsible AI Institute member.

Media Contact

Nicole McCaffrey
Head of Strategy & Marketing, Responsible AI Institute
nicole@responsible.ai
+1 (440) 785-3588

Connect with RAI Institute

RAI Hub
LinkedIn
Slack
YouTube