



OPEN ALLIANCE FOR CLOUD ADOPTION

TOPIC: AI GUIDANCE

Contributors

- Andre Schwan — Gijima
- Harish Mekerira — Clean Harbors
- Igor Andrade Araujo — Insi
- Jarrod Mexted — Canadian Animal Protection Services
- Mark Williams — TachTech
- Matt Estes — The Walt Disney Company (retired)
- Roman Macak — T-Systems
- Ryan Skipp — T-Systems
- Shamir Charania — Clyfar Solutions
- Thomas Costello — Booz Allen Hamilton

CONTENTS

Contributors	1
LEGAL NOTICE.....	3
Executive Summary	4
Background	4
The “Return on X” (RoX) Framework	4
Introduction	6
Premises of this guidance	8
The four key macro-level areas of consideration	8
Strategy.....	9
Structure	10
Capability	12
Legal and Compliance	13
Implementation Approaches.....	15
Implementation Case 1: Commercial Structure (Marketing, Finance)	15
Implementation Case 2: Customer Service	16
Implementation Case 3: Product Development	17
Implementation Case 4: Supply Chain.....	18
Implementation Case 5: Engineering and Risk Management.....	19
Conclusion.....	20
References	21

LEGAL NOTICE

© 2026 Open Alliance for Cloud Adoption – A Linux Foundation Initiative, Inc. ALL RIGHTS RESERVED.

This “OACA AI Guidance – White Paper” is proprietary to the Open Alliance for Cloud Adoption – A Linux Foundation Initiative (the “Alliance”) and/or its successors and assigns.

This OACA document is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0). To view a copy of the license, visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>

If any derivatives of this document are published, the following statement must be identified: *“This document is based on the OACA AI Guidance – White Paper document created by the Open Alliance for Cloud Adoption – A Linux Foundation Initiative (OACA) but may contain changes to the original OACA document which have not been reviewed or approved by the OACA.”*

LEGAL DISCLAIMER:

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN “AS IS” BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, THE ALLIANCE (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NONINFRINGEMENT. THE INFORMATION CONTAINED IN THIS DOCUMENT IS FOR INFORMATIONAL PURPOSES ONLY AND THE ALLIANCE MAKES NO REPRESENTATIONS, WARRANTIES AND/OR COVENANTS AS TO THE RESULTS THAT MAY BE OBTAINED FROM THE USE OF, OR RELIANCE ON, ANY INFORMATION SET FORTH IN THIS DOCUMENT, OR AS TO THE ACCURACY OR RELIABILITY OF SUCH INFORMATION. EXCEPT AS OTHERWISE EXPRESSLY SET FORTH HEREIN, NOTHING CONTAINED IN THIS DOCUMENT SHALL BE DEEMED AS GRANTING YOU ANY KIND OF LICENSE IN THE DOCUMENT, OR ANY OF ITS CONTENTS, EITHER EXPRESSLY OR IMPLIEDLY, OR TO ANY INTELLECTUAL PROPERTY OWNED OR CONTROLLED BY THE ALLIANCE, INCLUDING, WITHOUT LIMITATION, ANY TRADEMARKS OF THE ALLIANCE.

TRADEMARKS: OPEN ALLIANCE FOR CLOUD ADOPTIONSM, OACASM, and the OPEN ALLIANCE FOR CLOUD ADOPTION logo[®] are trade names, trademarks, and/or service marks (collectively “Marks”) owned by Open Alliance for Cloud Adoption – A Linux Foundation Initiative, Inc. and all rights are reserved therein. Unauthorized use is strictly prohibited. This document does not grant any user of this document any rights to use any of the OACA’s Marks. All other service marks, trademarks, and trade names reference herein are those of their respective owners.

Executive Summary

Artificial Intelligence (AI) is having one of those moments where the marketing vocabulary ("transformation", "agents", "revolution") is regularly louder than the organizational reality. In that kind of noise, the most common mistake is not that leaders choose the wrong model, but that they choose a posture that treats AI as inevitable, universally beneficial, and exempt from the normal disciplines of strategy, governance, and capital allocation. This paper is written as a counterweight to that gravitational pull. Not anti-AI, not pro-AI, but pragmatic about how and where to integrate AI.

"Fear of missing out" is a terrible operating model. It pushes organizations into blanket adoption ("everyone gets a copilot!") before they know what is actually good, safe, governable, and worth funding. Here we reintroduce basic executive questions that, before we build or scale anything, attempt to answer:

1. What could we be doing with AI?
2. What can we be successful at doing with AI?
3. What organizational changes do we need to make to effectively use AI?

The split (1 and 2 vs. 3) matters because these groupings represent two fundamentally different types of organizational discovery. The first grouping is about exploring the options space without committing to it. The second grouping is about understanding the required organization changes to be successful in adopting AI and is the subject of this paper.

Effectively, we are building toward a practical outcome: a maturity model for AI use and adoption within organizations. This is rooted in (and builds upon) our existing work in creating the OACA Cloud Maturity Model (CMM).

Background

Some important background to help frame the rest of this work.

The “Return on X” (RoX) Framework

If we only measure AI with return on investment (ROI), we will predictably fund the kinds of AI work that are easiest to spreadsheet and hardest to live with. ROI is necessary but not sufficient to effectively evaluate AI use cases. Adoption tends to follow J-curves, and many effects of AI integration (cognitive load shifts, flow changes) cannot be quantitatively measured ahead of time.

To counter this, we introduce the RoX: Return on X as a deliberately expandable frame.

What kind of return do we actually care about, and what evidence would convince us it's real?

In this work we focus on three “Xs” (described below) because they’re pragmatic anchors often overlooked in AI hype cycles. You can layer in others over time (such as return on trust), but starting here helps avoid familiar failure modes of delivering solutions that appear efficient on paper while quietly eroding organizational capacity and long-term potential.

1) Return on Investment (ROI)

ROI is the financial lens: net benefits relative to cost which is typically expressed as a percentage. Leadership has to allocate scarce capital, but ROI varies over time and the biggest gains might not be realized by actions that generate “immediate profits”.

2) Return on Employee (ROE)

ROE asks: how does AI change the value employees create relative to the effort required to employ them?

AI is an amplifier of what is already there. In healthy systems, this shows up as real gains (in flow, happiness, etc.) but in brittle systems it can amplify friction, rework, and increase hidden labor (example: babysitting AI processes). ROE isn’t about squeezing people, it’s about understanding when AI reduces operational drag and cognitive tax, or when it quietly adds to it.

3) Return on Future (ROF)

ROF is the long-horizon lens: value created by investing in adaptability, not just immediate results. In AI, it’s what separates organizations that scale safely from those that keep starting over.

Signals of ROF include:

- Lower cost of future-adoption (for example, of the next technology wave)
- Ease of movement (for example, reduced vendor)
- Built-in governability

The point isn’t to get a perfect score. The point is to stop pretending all initiatives are equally feasible and to avoid investing where success depends on miracles. If you’ve done this part well, you end up with a small set of initiatives that are not merely exciting, but winnable.

Introduction

Most organizations don't have an "AI strategy" problem. They have a perception problem. The current discourse encourages leaders to treat AI like a universal solvent: pour it on every workflow to dissolve inefficiency and then declare victory because AI is used. That story is emotionally satisfying (and conveniently sells tools and consulting services) but it's also how you end up with a portfolio of initiatives that look impressive in demos and quietly underperform in reality because they aren't seeing AI as a tool to help achieve goals.

What could we be doing with AI?

This question exists to slow the reflex to jump to solutions. Not to block progress, but to make progress less naive. It shifts the conversation from what AI can do in theory to what's worth doing here: in this organization, this market, under these constraints, and with these people.

Here's the key move: instead of asking AI to generate solutions, you use AI to extract signals you'd otherwise miss. These signals can be about demand, constraints, and organizational capabilities.

For example, much of AI opportunities is just resolving unacknowledged friction. Feed AI the artifacts you already have (support tickets, pain points, lessons learned, process documents) and ask it to synthesize:

- Top job stories (what are people trying to do)
- Recurring friction and rework points
- 3-5 opportunity hypotheses framed as "If we improved X, Y would become easier because Z"
- What evidence would falsify each hypothesis.

Notice what this does: it forces specificity and it forces humility. If you can't name what would falsify an idea, you are not sensing, but rather campaigning!

Apply the same approach to explore constraints and capabilities signals. Once you are done, you are left with something lightweight but powerful. Not a roadmap, but a small, testable set of hypotheses grounded in real demand, organizational constraints, and organizational capabilities. A modest but powerful "field of possibilities".

What can we be successful at doing with AI?

The first question defines what's possible (and viable). It's where ideas meet the realities of operability, governability, and economics. You can prototype almost anything... you can't run almost anything.

For each hypothesis, explore what "must be true":

- Must be true operationally: We can monitor outputs, detect failures, roll back quickly, and handle incidents without improvisation
- Must be true organizationally: Someone is accountable, people know how to work with it, and it doesn't create hidden labor that cancels the benefit
- Must be true economically: The cost curve makes sense at real usage volumes, the risk is acceptable to the bottom line

Then pressure test them. Here is an example of how AI can be prompted to explore this question.

For each opportunity hypothesis below:

- 1) List the "must be true" prerequisites for us to deploy this safely and sustainably.
- 2) Identify likely failure modes (technical, workflow, human, reputational).
- 3) Specify what we would need to monitor to detect those failures early.
- 4) Propose a rollback / kill-switch strategy that can be executed quickly.

This surfaces a key insight: most ideas fail not because they're bad, but because capabilities to deliver and operate are missing. At this stage, the goal isn't a roadmap. It's a ranked set of feasible bets, grounded in what you can actually operate.

<i>What organizational changes do we need to make to effectively use AI?</i>

Well, that is the subject of the rest of this paper. Enjoy!


Premises of this guidance

At the organizational level, leadership defines the overall strategic objectives that guide how AI should be used across the business. These macro-level objectives create a common direction for all business units. Senior decision-makers set the top-level “prompt model” by defining the key priorities, questions, and outcomes AI should support. Each business unit then adds more specific prompts based on its area of responsibility, while teams define the data and metrics that show whether they are contributing to the broader objectives.

Benefits should be measured against these macro-level goals. AI can collect, combine, and analyze data from across the organization within this context, helping surface the most relevant insights. This allows people to spend less time gathering and analyzing data manually, and more time on **interpretation, judgment, problem-solving, and decision-making**.

The four key macro-level areas of consideration

There are four overarching elements which must be aligned in context of AI use, and measured as described above, in terms of benefits to People, Investment, and the organizations’ Future.



We analyze these further next:

Strategy

Challenge	Guidance
Organizations typically lack a strategic framework for adopting artificial intelligence that ensures: • alignment with organizational goals, • mitigates risks, and • maintains a structured yet agile and rapid implementation approach	The organization should adopt a clear AI strategy that supports business goals, manages risks, and enables continuous improvement. • Align AI initiatives with core business objectives. • Identify and manage implementation risks. • Use a structured framework that supports agile delivery. • Respond quickly to changing needs and challenges. • Provide regular training and stakeholder engagement. • Continuously evaluate and refine AI solutions.

Return on Investment

- Successful AI adoption can lead to increased **operational efficiency**, resulting in significant **cost savings** and enhanced **revenue generation** through improved decision-making and innovation. (Doing the job that it was meant to do, avoidance of features void of real value, matching to customer wants and needs).

Return on Employee

- Empowering employees with AI tools and training fosters a culture of continuous learning and collaboration, **enhancing productivity** and job **satisfaction** while **attracting top talent**.

Return on Future

- By strategically implementing AI, the organization positions itself as an industry leader, ready to **adapt to** future technological **advancements and market shifts**, ensuring **long-term sustainability and growth**.

Structure

Challenge	Guidance
Traditional organizational structures are misaligned in terms of: • Traditional hierarchies slow AI-driven decision-making and responsiveness • Centralized control limits experimentation and value realization • Lack of a defined AI operating model leads to fragmented initiatives and duplication • Separation between business, data, and technology functions creates misalignment • Accountability across the AI lifecycle is often unclear • AI is frequently delivered as isolated projects instead of reusable platforms • Limited governance and distributed capability hinder scalable AI adoption	AI is inherently cross-functional, and organizational structures must reflect this by: • Adopting a federated AI operating model balancing governance with decentralized execution • Establish a lean central AI function to define strategy, standards, and guardrails • Enable empowered cross-functional teams aligned to business domains • Create clear ownership and accountability for data, models, and business outcomes • Distribute AI capability across the organization to support real-time decision-making enabling teams to integrate business, data and technology capabilities. • Structure and roles must accommodate flexible real-time decision making and adjustment to product or service development requirements, as an enabler. • Promote reusable platforms, shared services, and embedded governance • Evolve leadership from control-focused management to coaching and enablement • Continuously refine structures to reduce silos and improve agility.

Return on Investment

A well-defined AI operating model provides:

- **Reduced duplication**, rework, and fragmented delivery
- Faster **time-to-value** through improved cross-functional execution
- **Lower delivery and scaling costs** through reuse of platforms and models
- Improved consistency, governance, and **execution predictability**
- Reduced operational and implementation **risk**
- and more **consistent outcomes** from AI investments.

Return on Employee

Clear structures, roles, and delegated decision-making authority:

- Reduced organizational **friction and decision bottlenecks**
- Empowered teams able to act on *real-time insights* confidently
- Increased **productivity, innovation**, and accountability
- Less time spent navigating **approvals and complexity**
- Improved engagement, trust, and **organizational agility**
- **Leadership focus** shifts toward enablement and support

Return on Future

A scalable and adaptable AI operating model:

- Enables easier **integration of future technologies** and capabilities
- **Simplifies scaling** of successful AI use cases across the organization
- Reduces future transformation **cost and complexity**
- **Improves agility** in responding to market and regulatory change
- Strengthens long-term competitiveness and **stakeholder confidence**

Capability

Challenge	Guidance
Members of the organization, along with contractors, partners, or consultants, require adequate knowledge of AI technologies to • assess suitable solutions for adoption, • implement effective change management strategies, and • evaluate the business value of AI initiatives?	The organization must • invest in assessment and training programs that enhance employees' knowledge of AI solutions. • Regular evaluation of the skills and knowledge of employees, partners, and consultants will ensure: ○ all stakeholders are equipped to support the organization effectively while also ○ assessing the impact and business value of AI initiatives to refine approaches aligned with organizational goals.

Return on Investment

Successful enhancement of employees' AI knowledge will lead to:

- improved **operational efficiency and innovation**,
- ultimately resulting in **cost savings and increased revenue** through more effective AI solutions.

Return on Employee

Empowering employees with AI training:

- fosters a culture of **continuous learning and collaboration**,
- boosting job **satisfaction and retention** while
- equipping them with valuable skills for **career advancement**.

Return on Future

By strategically investing in AI capabilities, the organization positions itself as

- a **market leader**,
- **ready to adapt** to future technological advancements and emerging trends,
- ensuring long-term **sustainability and competitiveness**.

Legal and Compliance

Challenge	Guidance
In the AI landscape, both technological capabilities and regulatory controls are evolving rapidly and not always in sync which leads to: • Organizations frequently lack the legal expertise, governance structures, and practical experience required to guide responsible AI adoption, use and development. • Regulatory fragmentation and differing AI risk classifications create complexity across jurisdictions and use cases. • Accountability for AI-driven decisions is often unclear, particularly where human and machine decision-making intersect. • Reliance on third-party AI models, platforms, and data introduces ecosystem, contractual, and dependency risks. • Measuring and demonstrating compliance remains difficult and can expose organizations to regulatory, legal, reputational, and operational risk. • Key risks include intellectual property disputes, licensing breaches, bias and discrimination, privacy and security failures, regulatory penalties, and costly rework.	Organizations should: • Establish a cross-functional AI Governance Committee comprising legal, risk, compliance, business, and technical leadership. • Identify and assess all AI use cases across the organization and apply a risk-based governance framework. • Define clear accountability, oversight, validation, and auditability across the full AI lifecycle. • Govern third-party AI dependencies contractually and operationally, including data usage, IP rights, liability, and service boundaries. • Adopt a “compliance by design” approach, embedding legal, ethical, and regulatory controls early in delivery lifecycles. • Use targeted external expertise where required to address evolving regulatory requirements and specialist risks. • Implement ongoing training, compliance monitoring, tooling, and independent audits to maintain visibility and accountability.

Return on Investment

A proactive, well-governed approach to legal and compliance enables:

- **Reduced** exposure to regulatory fines, litigation, reputational damage, and costly remediation or rework.
- Embedding of compliance early lowers implementation risk and **reduces long-term delivery costs**.
- Clear governance and streamlined legal review processes **accelerate time-to-market** and improve confidence in AI investments.
- Improved **predictability and sustainability of financial returns** from AI-driven initiatives.

Return on Employee

Clear governance, policies, and accountability structures support:

- Reduced ambiguity and **decision friction**.
- Employees to adopt and use AI with greater confidence, consistency, and ethical clarity.
- The shifting of **legal and compliance functions** from reactive gatekeeping to strategic enablement and advisory roles.
- Teams to spend **less time navigating** uncertainty and more time focused on innovation, decision-making, and **value creation**.
- Improved **capacity for scaling**

Return on Future

By embedding robust legal and compliance capabilities early, organizations create a foundation that:

- Establishes a **scalable and adaptable foundation** for operating within evolving regulatory and ethical environments.
- Reduces future **disruption, retrofit costs, and complexity** as AI regulations and technologies mature.
- **Strengthens** trust, brand **credibility**, and stakeholder confidence through responsible AI adoption.
- **Improves long-term agility, scalability, and competitiveness** across markets, partners, and jurisdictions.

Implementation Approaches

It helps to consider some practical implementations of the above:

Implementation Case 1: Commercial Structure (Marketing, Finance)

Challenge	Guidance
Traditional structures do not adjust dynamically to the new opportunities that AI present and try to handle it in the traditional manner – losing potential momentum.	Dynamic decision-making data and tools must be available to all teams easily and aligned to their business areas, to make and test decisions, and equally to reverse decisions if the sense-and-respond data provides negative feedback.

Return on Investment

Increased number of potential product trials with fast prioritization and reduced people time and meetings. Each area is evaluated an optimized based on their own appropriate metrics, rather than centrally standardized metrics, allowing for better analysis of the business value they produce.

Return on Employee

People are able to make fast decisions based on real-time sense-and-respond data, with suggested options, saving them having to go and collect and analyze data individually. People in the team leveraging appropriate AI modeling per function can leverage data more appropriately, and present it more effectively and quickly, focusing on problem items and recommended corrective actions to take based on specialist technology-based outputs per area. They can also become more deeply integrated into the bus areas, rather than on background data analysis, collection, and presentation.

Return on Future

More products and opportunities can be identified and prioritized early on, so as to shape and prepare the organization properly in advance. Recommendations and subtle trend identification (early on) drive course corrections before they become problems. This happens because even small deflections have data backing them, and finance teams don't have to wait for them to become problems, before they actually raise them for attention.

Implementation Case 2: Customer Service

Challenge	Guidance
Customer Service is an area that has received an outsized focus for AI investment, primarily driven by cost reduction potential. However, most of it is external facing and not internal to the organization.	The lens of Customer Service should be broadened for not just cost reduction and external facing but also incorporating advancements in multi-modal AI development to enhance overall customer satisfaction across both internal and external service interactions.

Return on Investment

Mature customer service functions leverage conversational AI, intelligent routing, and predictive analytics to significantly reduce cost-to-serve while increasing speed and personalization at scale. AI-driven automation handles routine inquiries, allowing human agents to focus on high-value interactions that drive retention and upsell opportunities.

Return on Employee

AI augments employees through real-time assistance, automated knowledge retrieval, sentiment analysis, and next-best-action recommendations, reducing cognitive load and burnout. Rather than replacing agents, mature AI-enabled environments elevate them into problem solvers and relationship managers supported by intelligent copilots. This increases confidence, skill development, engagement, and retention while redefining roles toward higher-value human work.

Return on Future

Customer service becomes a continuous learning system where AI transforms customer interactions into predictive insight for product, risk, and growth strategy. Generative AI and machine learning surface emerging trends, anticipate churn, and enable hyper-personalized experiences that evolve with customer behavior. This positions the organization to adapt faster than competitors, build long-term trust, and create durable strategic advantage in an AI-driven marketplace.

Implementation Case 3: Product Development

Challenge	Guidance
Current-to-target state of the business's development plan is missing.	With rapid advancements in tools (software and hardware capabilities) companies around the world are hard-pressed to make quick decisions to demonstrate AI readiness and value. However, this is creating undue stress leading to missed opportunities to look at the holistic business value of adoption AI for product development. Create an enterprise level prompt framework to guide and accelerate development in line with industry and compliance, etc.

Return on Investment

In AI-mature product organizations, machine learning accelerates market research, demand forecasting, and feature prioritization, reducing costly misaligned investments. Generative AI speeds up prototyping, design iteration, code generation, and testing cycles, shortening time-to-market while lowering development costs. As maturity increases, AI-driven portfolio optimization ensures capital is continuously allocated toward the highest-value, data-validated opportunities.

Return on Employee

AI augments product managers, designers, and engineers with intelligent copilots that support ideation, simulation, automated testing, and documentation, reducing repetitive tasks and cognitive overload. Teams can shift from manual analysis to strategic problem-solving and creative innovation. This elevates skill development, cross-functional collaboration, and engagement while attracting top talent eager to work in AI-empowered environments.

Return on Future

At advanced maturity levels, AI transforms product development into a continuous learning ecosystem where real-time user data feeds predictive design and adaptive product evolution. Digital twins, scenario modeling, and predictive analytics enable organizations to anticipate market shifts before competitors. This future-ready capability builds resilient product portfolios, accelerates innovation cycles, and positions the enterprise to lead in emerging markets rather than react to them.

Implementation Case 4: Supply Chain

Challenge	Guidance
There is misalignment to the rest of the industry, and if leading, you may diverge too far, but if following, you lose revenue potential	How to determine whether to prioritize on leading and setting the path, or risk false investments which increase opportunity cost, vs optimizing current costs

Return on Investment

In AI-mature supply chains, machine learning improves demand forecasting accuracy, inventory optimization, and dynamic pricing, significantly reducing working capital and stockouts. Intelligent automation streamlines procurement, logistics routing, and warehouse operations, lowering operational costs while improving service levels. As maturity advances, predictive risk modeling minimizes disruption impact, protecting margin and stabilizing cash flow.

Return on Employee

AI augments planners and operators with real-time decision intelligence, anomaly detection, and scenario simulation tools that reduce manual firefighting and spreadsheet dependency. Employees shift from reactive coordination to strategic orchestration, supported by predictive insights and automated workflows. This reduces stress, enhances analytical skill development, and increases engagement by enabling higher-value, decision-oriented roles.

Return on Future

At advanced maturity stages, AI enables autonomous and self-optimizing supply networks that continuously learn from global data signals, market shifts, and geopolitical risks. Predictive analytics allow enterprises to simulate future disruptions and proactively redesign sourcing and distribution strategies. This builds a resilient, adaptive supply ecosystem capable of sustaining competitive advantage in an increasingly volatile and AI-driven global economy.

Implementation Case 5: Engineering and Risk Management

Challenge	Guidance
AI tools add complexity to stakeholder accountability, as owners must take responsibility for systems whose behavior may not be fully predictable or explainable. This challenges traditional GRC models focused on periodic assessments; instead, continuous monitoring is required to manage model drift, evolving use cases, and increasingly dynamic risk profiles.	AI Risk Management – Action Plan • Categorize: Tier all AI use cases by risk and assign a business owner accountable for outcomes. • Define: Establish and publish AI policies upfront; schedule regular policy reviews. • Monitor: Implement continuous monitoring for model drift, bias, and control effectiveness. • Validate: Track vendor updates and confirm systems remain aligned with approved implementations. • Control: Review user behavior regularly to ensure adherence to approved use cases. • Assess: Map integrations and quantify failure impact; align dependencies with risk appetite. • Automate: Use AI to create feedback loops for risk signals, regulatory monitoring, and policy updates.

Return on Investment

- **Lower costs:** Detection of drift early, enforcement of controls, and risk-tiered guardrails reduce incidents and regulatory exposure.
- **Faster, safer adoption:** Standardized approvals and decision paths limit escalations.
- **Proactive risk control:** Continuously monitored models prevent high-impact failures and enable rapid containment.
- **Reduced audit effort:** Measurable controls and regular reviews shorten audits and minimize remediation.
- **Better decisions & spend:** Leverages performance insights optimize AI investments and eliminate underperforming use cases.

Return on Employee

- **Clear accountability:** Leaders confidently approve AI use cases with well-defined ownership of outcomes and minimal escalation.
- **Confident adoption:** Employees use AI consistently within understood boundaries, reducing hesitation and misuse risk.
- **Ongoing visibility:** Leadership remains informed through continuous risk transparency, avoiding blind spots from one-time approvals.
- **Reduced burden:** System owners operate with clear metrics and thresholds, lowering ambiguity and stress around accountability.
- **Mature risk culture:** Employees recognize that risk is actively monitored and managed, reinforcing responsible AI usage.
- **Safer operating environment:** Proactive detection improves safety across operations, including risks, fraud, and insider threats.

Return on Future

- **Regulatory readiness:** Organization consistently operates ahead of evolving AI regulations, with compliance handled incrementally rather than through disruptive changes.
- **Trusted governance posture:** Demonstrated due diligence builds sustained trust with regulators, customers, and stakeholders.
- **Predictable risk containment:** Issues are identified early and contained within defined boundaries, reducing escalation and systemic impact.
- **Non-reactive operations:** The organization responds to risk with structured, controlled processes rather than ad hoc reactions.
- **Adaptive compliance capability:** Continuous visibility and telemetry enable rapid alignment to new global regulatory requirements.
- **Enhanced resilience to threats:** Strong defense posture against increasingly sophisticated AI-driven external risks.

Conclusion

By grounding AI adoption in the question, *“What kind of return do we actually care about, and what evidence would convince us it’s real?”*, organizations can move beyond hype-driven decisions toward disciplined, outcome-oriented execution.

This guidance reinforces that successful AI adoption is not defined by how much is deployed, but by how well it aligns with three critical dimensions of return:

- **Return on Investment (ROI):** ensuring that AI initiatives deliver measurable financial value, with a clear understanding of cost, scalability, and sustainable impact.
- **Return on Employee (ROE):** ensuring that AI enhances—not erodes—human capability by reducing friction, improving flow, and enabling employees to operate at a higher level of effectiveness and engagement.

- **Return on Future (ROF):** ensuring that today’s investments build adaptability, governability, and strategic flexibility, lowering the cost and risk of future transformation.

When these three lenses are applied together, organizations are better equipped to answer three foundational questions with clarity and discipline:

1. What could we be doing with AI?
2. What can we be successful at doing with AI?
3. What organizational changes are required to make that success real and sustainable?

This integrated approach prevents a familiar failure mode: delivering AI solutions that appear efficient on paper yet quietly degrade organizational capacity, increase hidden complexity, or constrain future options.

Instead, by balancing financial outcomes, human impact, and long-term adaptability, organizations can focus their efforts on a smaller set of initiatives that are not only innovative, but truly winnable, scalable, and enduring.

We hope that this lens has provided value and thought material, and would love to hear your feedback via our Linux Foundation Initiative page:

<https://oaca-project.org/contact/> or look for more material at <https://www.linuxfoundation.org/Projects/oaca>

References

Current OACA and other material addressing this topic:

ID	Name	Location
1	Karl Popper: Theory of Falsification	https://www.simplypsychology.org/karl-popper.html
2	Gartner: Outcomes that matter: ROE, ROI, ROF	https://www.gartner.com/en/documents/6330147
3	OACA: AI Responsibly – Position Paper	https://oaca-project.org/assets/whitepaper/AI-Responsibly-Position-Paper-2023.pdf