



Learning Outcomes

AI & Technology Literacy

Core & Specialized Credentials

International Certification Alliance

Version 1.0 — August 2026



LICENSING INFORMATION

The work in this document was facilitated by the International Certification Alliance (ICA) and done through the contributions of various experts and practitioners. These learning objectives are intended to help the growing Tomorrow-Ready™ community worldwide.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/> or send a letter to:

Creative Commons,
P.O. Box 1866,
Mountain View, CA 94042, USA

YOU ARE FREE TO:

Share — copy and redistribute the material in any medium or format

UNDER THE FOLLOWING TERMS:

Attribution: You must give appropriate credit to The International Certification Alliance (ICA), provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

NonCommercial: You may not use the material for commercial purposes.

NoDerivatives: If you remix, transform, or build upon the material, you may not distribute the modified material.

NOTICES:

You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.

No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.



TABLE OF CONTENTS

LICENSING INFORMATION	2
TABLE OF CONTENTS	3
HOW TO READ THIS DOCUMENT	4
CORE CREDENTIAL LEARNING OUTCOMES	5
AI & Technology Literacy Core	5
SPECIALIZED CREDENTIAL LEARNING OUTCOMES	8
AI and Human Judgment	8
Leading in an AI-Augmented Organization	11
AI-Augmented Cybersecurity Awareness	13
AI for Customer Insights	15
AI for Product Discovery	17
AI for Product Strategy	19
AI for Product Metrics	21
AI for Product Planning	23
AI for Stakeholder Management	25
AI-Enabled Change	27
Data, AI, and Decision Accountability	29
SPECIAL THANKS	30

HOW TO READ THIS DOCUMENT

This document outlines the learning outcomes that accredited member organizations must address to offer ICA's AI & Technology Literacy certification. To achieve the certification, a learner must complete the core credential plus at least two specialized credentials.

Each learning outcome (LO) follows a particular pattern, described below.

O.O Learning Outcome Name

Additional Context, describing why this Learning Outcome is important or what it is intended to impart.

The learning outcome's purpose further describes what is expected to be imparted to the learner (e.g., a key point, framework, model, approach, technique, or skill).

CORE CREDENTIAL LEARNING OUTCOMES

AI & Technology Literacy Core

AI & Technology Literacy is an 8-hour Core Credential for professionals who want to develop a clear-eyed, confident understanding of how artificial intelligence is changing the nature of work and what that means for the people doing it. AI systems are reshaping how decisions get made, how information gets produced, and how organizations create and deliver value. Professionals who engage with AI only as users of tools – without understanding what AI can and cannot do, why human judgment and accountability remain essential, and what responsibilities come with deploying AI in organizational contexts – are poorly equipped for the work ahead.

1. What AI Is and What It Actually Does

1.1 How Generative AI Produces Outputs

Generative AI produces text by predicting likely continuations from patterns in the material it was trained on. Fluency is therefore not evidence of accuracy, and practitioners can mistake a confident response for a reliable one.

Establish an accurate, nontechnical understanding of generative AI so learners can use its outputs with appropriate judgment. Cover the difference between a deterministic system which follows rules they were given (traditional automation) and generative systems which can produce variable results each time it runs.

1.2 AI Capabilities and Reliability

AI capabilities vary widely and move quickly. What is dependable in one use may be unreliable in another, and unreliable today may be dependable within months. Capability fit and reliability requires ongoing judgement.

Develop learners' ability to judge where an AI capability adds value and where it creates risk for a particular task, using a vocabulary that remains useful as capabilities change. Cover how underlying data quality affects fit as much as underlying AI capability.

1.3 AI's Role in Work Activities

AI affects activities within a role differently. Some work can be automated, some is strengthened through augmentation, and some still depends on human presence, judgment, or accountability.

Help learners develop a differentiated view of how AI affects different work activities rather than treating jobs as simply replaceable or protected. Cover the risks on both sides: such as automation removing the judgement that would have caught errors, and augmentation degrading into rubber-stamping. Before either, cover whether the work activity needs to happen at all and AI is just making unnecessary work faster.

2. Human Capabilities in an AI-Augmented World

2.1 Human Capabilities in AI-Augmented Work

As AI takes on more routine cognitive work, greater value shifts toward capabilities such as judgment under ambiguity, ethical reasoning, relational intelligence, contextual interpretation, and creative synthesis. Which of these matters most depends on the role, the task, or the system being used.

Develop learners' ability to identify which human capabilities become most valuable in their own work with AI, and what is lost when those capabilities are weak.

2.2 Directing AI Work

As more work is composed through AI, the quality of a result depends on how well the work was set out before it began. And simply stating the task is not enough. Product literacy, the ability to explain what a piece of work is for, why it matters, how people will use it, and what a good result looks like, was once mainly a product role's concern but is now needed by everyone working with AI.

Develop learners' product literacy so they can set out a piece of work well enough for an AI system to produce something usable, and recognize when a weak result came from how the work was set out rather than from AI itself.

3. Working Critically with AI Outputs

3.1 Reviewing AI-Generated Work

When work is produced by AI, judgement is a separate and deliberate review step. Judgement which is further complicated with the fluent and convincing nature of AI output even when the work is inaccurate, biased, or poorly suited to the task.

Build learners' ability to review AI-generated work as an explicit step requiring verification and judgement against quality, accuracy, bias, and practical usefulness, rather than accepting output that reads well.

3.2 Human Judgment in AI-Assisted Work

Human accountability cannot be delegated to AI. Producing real AI-assisted work helps practitioners see where verification and judgment remain essential in the workflow.

Give learners direct experience using AI while maintaining clear human responsibility for the quality and consequences of the work.

4. Responsible AI in Organizational Contexts

4.1 Ethical and Data Considerations

AI use can create real consequences involving privacy, confidentiality, intellectual property, bias, consent, transparency, and harm to affected people or groups.

Ground AI ethics in practical organizational decisions by helping learners recognize the considerations most relevant to a specific use case.

4.2 AI Application Assessment

Human, ethical, and organizational consequences are easier to address before an AI application is deployed than after harm occurs. Pre-adoption evaluation supports more responsible innovation.

Build learners' ability to assess proposed AI applications before adoption using practical human, ethical, and organizational considerations.

SPECIALIZED CREDENTIAL LEARNING OUTCOMES

AI and Human Judgment

AI and Human Judgment is a 2–4 hour Specialized Credential for professionals who want to develop their ability to make clear, confident decisions about when to rely on AI outputs, when to override them, and when not to use AI at all. As AI systems become faster and more capable, the pressure to defer to their outputs increases – even in situations where human judgment is more reliable, more accountable, or more appropriate. This course builds the practical decision-making discipline needed to work with AI as a capable assistant rather than an unquestioned authority, including the ability to recognize the specific evidence that should prompt an override and to calibrate that judgment over time using real outcomes.

1.1 Conditions that Drive AI Over-Reliance

Over-reliance on AI is often driven by time pressure, confident AI outputs, social expectations, and the convenience of accepting a ready-made answer rather than by lack of knowledge alone. Automation is where the error cost is highest since nothing is left to catch it.

Help learners understand AI over-reliance as a behavioral and organizational challenge that can be countered through deliberate practices. Cognitive conditions such as automation bias or cognitive ease. Behavior conditions such as skill atrophy or diffusion of responsibility.

2.1 Human-AI Decision Boundaries

AI adds different value across different decisions. Sound judgment requires clear boundaries for when AI can lead, when it needs verification, when human judgment should take precedence, and when the decision should not involve AI at all.

Build learners' ability to make explicit, defensible decisions about the appropriate role of AI and the evidence that should change that role. For example, an AI's role in a given decision could be to lead, to verify, or to defer to human judgment, depending on the decision's stakes and the evidence available.

3.1 Maintaining and Calibrating Independent Judgment

Independent judgment is difficult to sustain when overriding AI feels slower, riskier, or socially discouraged. Judgment also improves when past decisions are compared with actual outcomes.

Develop a repeatable practice that helps learners maintain independent judgment and calibrate it over time based on actual outcomes.

4.1 Organizational Conditions for Sound Human-AI Judgment

Individual judgment is difficult to sustain when organizational norms reward speed over accuracy or treat AI outputs as authoritative by default.

Help learners see human-AI judgment as an organizational design issue as well as an individual capability.

Leading in an AI-Augmented Organization

Leading in an AI-Augmented Organization is a 2–4 hour Specialized Credential for professionals in leadership, management, and change roles who want to develop their ability to build the organizational conditions that enable teams to work with AI responsibly, effectively, and sustainably. AI adoption in organizations is often driven by tools and technology decisions made faster than the human and organizational capabilities needed to use them well. Leaders who do not address the cultural, structural, and accountability conditions surrounding AI, and who do not follow through with the workflow changes and training investment that responsible adoption requires, leave their teams to navigate complexity alone.

1.1 Leadership Practices in an AI-Augmented Organization

AI does not change the purpose of leadership, but it increases the importance of modeling critical AI use, creating safety for concerns, and making accountability explicit.

Establish what effective leadership looks like when AI becomes embedded in everyday work.

2.1 Psychological Safety Around AI

Teams make poorer decisions when people feel pressure to accept AI outputs, hide uncertainty, or avoid raising concerns. Psychological safety supports more honest human–AI collaboration.

Help leaders recognize the psychological safety conditions that shape AI use and the leadership response needed to strengthen them. For example, the impact of how AI can be framed as either added capacity or as reduced headcount.

3.1 Accountability for AI-Assisted Decisions

When accountability is unclear, problems caused by AI-assisted decisions can be blamed on the tool, the data, or whoever happened to be involved. Clear responsibility supports safer use.

Develop leaders' ability to establish minimum–viable accountability for verification and outcomes before problems occur.

4.1 Culture of Responsible AI Adoption

AI adoption culture is shaped by what leaders model, what behaviors are rewarded, whether concerns are heard, and whether workflows and training evolve alongside the technology.

Help leaders assess AI adoption culture and focus leadership attention on changes that make responsible use easier to sustain. Cover indicators of a healthy culture such as whether concerns about AI outputs are raised openly or whether people disclose AI use.

AI-Augmented Cybersecurity Awareness

AI-Augmented Cybersecurity Awareness is a 2–4 hour Specialized Credential for professionals who want to build the practical judgment needed to recognize, evaluate, and respond to the cyber and information–security risks that AI tools introduce into everyday work, both threats coming in and confidential data going out. AI has made social engineering more convincing, security tooling more automated, and the line between a genuine signal and a fabricated one harder to see at a glance. At the same time, the most common daily incident is not a sophisticated attack but an employee pasting client or company data into a public AI tool without realizing the risk. Professionals who treat AI-generated security alerts, communications, or media as automatically trustworthy, or automatically suspect, and who are not alert to their own outbound data risk, are poorly equipped to protect themselves, their teams, and their organizations.

1.1 AI-Enabled Cyber Threats

AI makes convincing cyber threats easier and faster to produce, changing what suspicious activity can look and sound like.

Help learners update their assumptions about common cyber threats in an AI-augmented environment. For example, AI makes convincing phishing, voice cloning, deepfakes, and automated attack activities.

2.1 AI-Generated Security Signals

AI increasingly generates legitimate security alerts while also helping attackers create convincing imitations. Neither automatic trust nor automatic suspicion is a safe default.

Build learners' ability to treat AI-generated security signals as inputs that require verification before action. For example, this might include an AI-generated phishing email impersonating IT support or a deepfake voice request for an urgent wire transfer.

3.1 Verification Practices in an AI-Augmented Environment

Many AI-enabled attacks succeed because no verification step exists at the moment it matters. A simple practice established in advance can reduce that risk.

Develop learners' ability to use a practical verification routine for identity, requests, or content before acting. For example, a verification practice might be confirming an urgent payment request through a separate, pre-established phone call rather than replying to the email or message that made the request.

4.1 Readiness for AI-Related Security Risks

AI-related security risk includes both attacks coming into the organization and sensitive information leaving through unsafe AI use. Effective readiness must address both.

Help learners assess organizational readiness across inbound threats and outbound data exposure rather than focusing on attack awareness alone.

AI for Customer Insights

AI for Customer Insights is a 2–4 hour Specialized Credential for professionals who want to use artificial intelligence to deepen customer understanding, accelerate sense-making, and improve the quality and speed of customer insight. The credential focuses on using generative, analytical, and emerging agentic AI to work with large volumes of structured and unstructured customer evidence; surface patterns, themes, sentiment, behaviors, and meaningful customer differences; strengthen research and feedback workflows; and generate hypotheses that can be tested with real customers. AI can dramatically reduce time-to-insight, but it can also create convincing patterns that are incomplete, biased, poorly grounded, or disconnected from actual customer experience.

1.1 Insight Questions and AI Fit

AI can produce summaries, personas, predictions, and recommendations quickly, but speed does not guarantee useful insight. Poorly framed questions, weak source data, or using AI for a task it is not suited to can create confident output without improving customer understanding.

Develop learners' ability to begin with a clear customer-insight question and determine where AI can add value. Introduce the differences among generative, analytical, predictive, and emerging agentic uses of AI at a practical level; consider the evidence available, the decision the insight should inform, and where human or direct customer research is still required.

2.1 AI-Assisted Synthesis and Pattern Detection

Customer evidence is increasingly spread across surveys, interviews, support interactions, reviews, call transcripts, behavioral data, digital interactions, images, and other sources. Manual analysis can miss patterns or take too long, while AI can surface themes at scale but can also amplify noisy, unrepresentative, or contradictory signals.

Develop learners' ability to use AI to synthesize structured and unstructured customer evidence and surface themes, sentiment, behaviors, friction, emerging needs, and meaningful differences among customer groups. Explore triangulation across sources, contradiction and outlier detection, basic clustering or segmentation support, and the importance of tracing an AI-generated insight back to the underlying evidence.

3.1 AI-Assisted Research and Hypothesis Generation

AI can help create interview guides, adapt follow-up questions, conduct conversational feedback, summarize interviews, identify research gaps, and simulate possible customer responses. These capabilities can expand the speed and reach of research, but poorly designed use can lead questions, flatten nuance, or mistake synthetic output for actual customer voice.

Develop learners' ability to use AI as a research partner across the customer-insight workflow. Explore AI-assisted research design, conversational or adaptive feedback, interview and qualitative synthesis, hypothesis generation, and synthetic customers or simulated feedback as tools for exploration—not validation. Reinforce when and how to return to real customers to test important assumptions.

4.1 Validation, Governance, and Human Judgment

AI-generated customer insights can sound authoritative even when they are based on incomplete data, hidden assumptions, poor representation, model bias, or inappropriate use of sensitive customer information. The risk increases when AI-generated summaries or recommendations move directly into decisions without validation.

Develop learners' ability to evaluate whether an AI-generated customer insight is sufficiently grounded and appropriate for action. Explore source provenance, representativeness, hallucination and inference risk, privacy and consent, bias and inclusion, confidence and uncertainty, human review, and proportional validation based on the consequence of the decision. Introduce the Tomorrow-Ready discipline of using AI to accelerate learning without surrendering accountability for the judgment.

AI for Product Discovery

AI for Product Discovery is a 2–4 hour Specialized Credential for professionals who want to develop their ability to apply artificial intelligence tools to accelerate and enrich product discovery work. AI can help teams surface customer problems faster, expand the idea space, and generate early-stage discovery artifacts — but only when practitioners can critically assess what AI produces and know where human judgment remains essential. This credential builds the practical capability to combine AI-generated insight with disciplined human analysis across the full discovery cycle.

1.1 Framing AI for Discovery

AI can accelerate discovery work, but it can also produce plausible answers to poorly framed questions. If the problem, hypothesis, source evidence, or constraints are unclear, the tool may increase speed without increasing learning.

Develop learners' ability to frame the discovery question before using AI and determine the role AI should play. Explore context quality, source evidence, assumptions, and where AI is useful for synthesis, challenge, generation, or prototyping versus where direct customer or technical evidence is required.

2.1 AI-Assisted Discovery Synthesis

AI can surface patterns across interviews, surveys, support data, research notes, and usage evidence much faster than manual review. The same capability can also flatten contradictions, reinforce the team's framing, or miss customer realities that are absent from the data.

Develop learners' ability to use AI to surface themes, contradictions, blind spots, risks, alternative explanations, and hypotheses from customer and product evidence. Reinforce traceability to source evidence and deliberate prompting that asks AI to challenge, not only summarize, the team's current thinking.

3.1 AI-Assisted Prototypes and Experiments

Discovery artifacts and prototypes are useful when they help test a specific product assumption. AI can create interview guides, prototype variants, test scenarios, and experiment ideas quickly, but generic inputs tend to produce generic artifacts that are disconnected from the real uncertainty.

Develop learners' ability to use AI to accelerate build-to-learn cycles. Explore rapid prototype generation, variant creation, discovery

artifacts, test design, fidelity choices, evidence thresholds, and how AI can help explore multiple solution directions before commitment.

4.1 Responsible Use of AI in Discovery

AI tools can reproduce bias, miss underrepresented customers, invent detail, and make synthetic users or polished summaries feel like real evidence. Responsible AI-assisted discovery therefore requires human judgment throughout the process to challenge outputs and determine where additional customer or technical validation is needed.

Build learners' critical awareness of where AI creates risk in discovery practice. Explore source provenance, representativeness, hallucination and inference risk, synthetic users as hypothesis tools rather than substitutes for research, privacy, consent, and appropriate human review. Develop learners' ability to validate AI-assisted discovery by applying appropriate human judgement and oversight before AI influences product decisions.

AI for Product Strategy

AI for Product Strategy is a 2–4 hour Specialized Credential for professionals who want to develop their ability to apply artificial intelligence tools to sharpen strategic product thinking. AI can accelerate market analysis, surface competitive signals, and support the generation and prioritization of product initiatives; but translating AI outputs into sound strategic decisions requires human judgment, contextual knowledge, and critical thinking that no tool can replace. This credential builds the practical capability to use AI as a strategic thinking partner while remaining clear-eyed about its limits.

1.1 AI-Assisted Market and Competitive Insight

AI tools can process large volumes of market and competitor data to surface patterns that inform strategic direction, but that analysis reflects the data it was given, not necessarily the full strategic picture.

Develop learners' ability to use AI for market intelligence as a starting point for strategic thinking. Explore synthesis across market and customer evidence, competitor moves, technology shifts, and internal performance, and distinguish useful signals from unsupported inference. Developing the critical human judgement needed to assess what AI analysis reveals, what it misses, and how to interpret and validate it.

2.1 AI-Assisted Strategic Options

AI can generate strategic product initiatives quickly, but volume is not the same as strategic quality. Its greater value is helping teams explore genuinely different options, challenge assumptions, and see how alternatives perform under different market or business scenarios.

Develop learners' ability to use AI as a strategic thinking partner to generate genuinely different product options, surface assumptions and risks, develop counterarguments, and explore scenarios rather than merely generate feature ideas. Develop critical human judgement to assess relevance to a defined customer problem, feasibility within organizational constraints, and alignment with business goals.

3.1 AI-Assisted Prioritization

Good prioritization decisions leverage both AI's analytical power and human wisdom about context, ethics, and organizational reality. Learning to intentionally design where AI informs and where human judgment decides is what enables better strategic choices.

Build learners' ability to design a decision process where AI enhances human judgment without replacing it. Help them understand that a good decision process clarifies what criteria matter, how AI will inform each criterion, where human judgment must step in, and how to know when to trust or question AI analysis. Evaluate options against customer outcomes, business impact, evidence, economics, feasibility, risk, and coherence from product vision through strategy to roadmap.

4.1 Communicating AI-Informed Product Strategy

A product strategy is a set of deliberate choices that must be communicable to stakeholders who will fund and act on it. AI can help synthesize evidence, tailor communication, and monitor new signals, but strategic coherence, accountability, and the decision to change direction remain human responsibilities.

Develop learners' ability to communicate an AI-informed strategy with clear separation between evidence, inference, and human choice, and to identify signals that should trigger strategic reassessment. Include key elements such as target market, strategic goals, success metrics, and AI-sourced benchmarks. Use AI to accelerate alternative explanations and stakeholder-specific communication without surrendering accountability.

AI for Product Metrics

AI for Product Metrics is a 2–4 hour Specialized Credential for professionals who want to develop their ability to apply artificial intelligence tools to identify, analyze, and act on product metrics. AI can help teams surface relevant metrics faster, detect patterns in existing data, and generate recommendations. But measurement decisions that are driven by AI without critical human oversight can produce metrics that optimize for the wrong things or mask the limitations of the underlying data. This credential builds the practical capability to use AI to enhance measurement practice while maintaining the judgment needed to act on data responsibly.

1.1 AI-Assisted Metrics Review

Many product teams inherit metric sets designed for different questions or product lifecycle stages. AI can accelerate the review of these metrics by surfacing relevance gaps, coverage issues, and proxy measures.

Develop learners' ability to use AI analysis as a diagnostic lens on existing measurement practices. Position AI as a thinking partner that can challenge an inherited metric set, not only summarise it.

2.1 AI-Assisted Metric Proposals

AI tools can generate metric suggestions quickly, but suggestions vary widely in quality. Some are well-aligned to outcomes, others are proxies or vanity measures. A proxy that a team is measured on will eventually be optimized at the expense of the outcome it stood for.

Develop learners' ability to treat AI-generated metric proposals as raw material, not finished recommendations, by testing each against a stated business or product goal before adopting it. Explore the difference between outcome, proxy, and vanity measures, and the behavior a metric starts to drive once a team is measured on it.

3.1 AI-Informed Product Decisions

AI tools can identify patterns and flag anomalies in metric data much faster than manual analysis, but the decision itself, what the evidence means and what action is warranted, remains a human responsibility.

Equip learners to complete the loop from AI-generated analysis to an accountable human decision. Explore how AI analysis can sharpen understanding of customer behavior, what it can and cannot establish about cause, and how recording the metrics and reasoning behind a decision lets it be revisited once the outcome is known.

4.1 AI-Assisted Metrics Reporting Quality

AI-generated reports and dashboards can appear more rigorous than they are, carrying the look of analytical authority even when the underlying data or metrics are flawed.

Build learners' ability to be a responsible final check before AI-informed measurement reaches decision-makers. Explore how chart choice, aggregation, and framing can overstate what the data supports, how bias enters through the data that happened to be available, and why measurement should run through the entire product or service lifecycle. Explore data provenance, definitions, missingness, representativeness, privacy, hallucination, model or analysis bias, and the level of human review appropriate to the consequence of the decision.

AI for Product Planning

AI for Product Planning is a 2–4 hour Specialized Credential for professionals who want to develop their ability to apply artificial intelligence tools to key product planning activities including roadmap creation, backlog prioritization, and release planning. AI can accelerate planning work, surface trade-offs that manual analysis might miss, and generate structured planning artifacts; but effective planning still requires human judgment about organizational context, stakeholder needs, and the consequences of sequencing decisions that no tool can fully anticipate. This credential builds the practical capability to integrate AI into planning practice while maintaining the accountability that sound planning demands.

1.1 AI in Product Planning

Product planning decisions are as much about organizational dynamics and stakeholder trust as they are about data. AI can improve planning efficiency, but plans built on AI alone risk being analytically plausible yet organizationally unworkable.

Establish where AI genuinely adds value across roadmap creation, backlog prioritization, and release planning, and where human judgment cannot be automated. Explore how AI-generated data supports replanning as conditions change, so planning is treated as something revisited rather than produced once.

2.1 AI-Assisted Planning Scenarios

AI can generate roadmap structures quickly, but formats that are technically coherent may not fit the organization's context or the stakeholders who will use the roadmap to decide.

Build learners' ability to treat AI-generated roadmaps as drafts, not finished artifacts, and to judge fit, trade-offs, and the adjustments needed before it is used. Explore multiple roadmap formats such as outcome-oriented and now-next-later views, how the choice of format changes what a roadmap communicates and what it hides, and how to use AI to test strategic coherence and surface alternatives rather than only to produce a layout.

3.1 AI-Supported Backlog

AI can apply prioritization frameworks, cluster requests, and draft epics or stories at scale, but the inputs to AI are judgment calls. Accepting the output without scrutiny can make uncertain estimates look objective.

Develop learners' ability to use AI to accelerate work decomposition and prioritization while preserving human accountability for value, scope, sequencing, estimates, and cross-functional feasibility. Explore how a framework converts judgment into scores, why its inputs need scrutiny, and how to use AI to check whether the order it produced actually serves the product strategy rather than only to produce it.

4.1 AI-Generated Release Plan Risk

AI-generated release plans can appear comprehensive and data-driven while reflecting assumptions, like team capacity or dependency risk, that don't hold in the specific organizational context. Assumptions and biases in AI-planning decisions negatively affect organizational priorities and stakeholder trust.

Build learners' ability to critically review and adapt AI-generated plans before they become commitments. Explore feasibility, dependency risk, change triggers, version rationale, and when a plan change requires explicit renegotiation.

AI for Stakeholder Management

AI for Stakeholder Management is a 2–4 hour Specialized Credential for professionals who want to develop their ability to apply artificial intelligence tools to identify stakeholders, understand their needs and influence, and design and deliver communication that reaches them effectively. AI can accelerate stakeholder mapping, surface patterns in stakeholder behavior, and support the creation of tailored communication at a scale and level of personalization that manual approaches alone cannot achieve. But stakeholder relationships are fundamentally human, built on trust, empathy, and the quality of real interaction, and no tool can substitute for the judgment and relational skill that effective stakeholder engagement requires. The ethical boundaries of AI use and the irreplaceable role of human skills are woven throughout all four learning outcomes.

1.1 AI-Assisted Stakeholder Identification

AI can help identify stakeholder groups, relationships, and informal influence by analyzing available data. Because AI-generated maps reflect the data provided, missing or biased data can create important gaps.

Develop learners' ability to use AI to identify stakeholder groups, influence, dependencies, decision relevance, and engagement needs while validating the map against human knowledge and respecting data boundaries. Explore the sources an AI can draw on, techniques for surfacing informal influence, and why a generated map understates the stakeholders least represented in the data available.

2.1 AI-Assisted Stakeholder Insights

AI can surface patterns in stakeholder needs, sentiment, relationships, and influence more quickly than manual analysis. The quality of these insights depends on the data, methods, and assumptions used.

Develop learners' ability to assess AI-generated stakeholder insights critically and distinguish useful patterns from artifacts of data or modeling choices. Explore techniques such as sentiment analysis, influence or network mapping, and persona generation, what each can and cannot establish about a stakeholder, and how to judge validity and reliability before acting on an insight.

3.1 AI-Assisted Stakeholder Communication Plans

AI can accelerate message planning, tailoring, channel selection, and timing for different stakeholder groups. Human review remains necessary to ensure the communication fits the relationship, context, and desired tone.

Develop learners' ability to use AI to accelerate stakeholder communication while retaining human judgment about context, relationships, and tone. Explore different approaches, such as a communication matrix, an engagement strategy, or a persona-based plan; how content, channel, timing, and degree of tailoring should differ by group; and when relationship history and organizational politics should override a plan that is analytically sound but relationally wrong.

4.1 AI-Assisted Messaging and Human Judgment

AI-assisted messages may be accurate but still lack appropriate tone, authenticity, or relational judgment. Stakeholder work also relies on human capabilities such as empathy, listening, and negotiation, while AI use raises ethical questions about profiling, transparency, and consent.

Build learners' ability to evaluate AI-generated messaging, recognize the limits of automation, and apply appropriate ethical and human judgment. Explore how to judge a message for effectiveness, clarity, accuracy, and tone; where human skills cannot be substituted; and the fairness, transparency, and consent questions raised by profiling and targeted communication.

AI-Enabled Change

AI-Enabled Change is a 2–4 hour Specialized Credential for practitioners who want to support organizations through the human side of AI adoption.

Note: This credential focuses on the human and organizational dimensions of AI adoption. It does not address technical AI implementation, AI tool selection, or how to use AI in change management practice.

1.1 What Makes AI-Driven Change Different

AI adoption can change both tasks and how people understand the value of their expertise. This creates identity and role disruption alongside the practical demands of technology adoption.

Build learners' understanding of what makes AI adoption a distinctive organizational change challenge. For example, unlike most technology rollouts, which reach a stable end-state once implemented, AI systems keep changing after go-live. The tool itself gets updated, retrained, or improved on an ongoing basis. This means practitioners can't treat AI adoption as a one-time transition with a clear finish line.

2.1 The Human Experience of AI Adoption

AI adoption can raise concerns about professional identity, future expertise, fairness, transparency, and job security. Different stakeholder groups may experience these questions in very different ways.

Develop learners' ability to recognize how AI adoption affects identity, role meaning, and psychological safety across stakeholder groups. For example, swapping one piece of software for another rarely makes someone question whether their expertise still matters, but AI can, when it starts doing part of what a person's professional identity was built on, like judgment, analysis, or craft.

3.1 Engaging Through AI Transitions

With AI, the answers themselves are often still being worked out what the tool can reliably do. Because the technology keeps evolving, engagement can't be a one-time occurrence.

Build the ability to design engagement that continues alongside AI's evolution. Cover why revisiting communication, policies, and boundaries is what keeps engagement credible. Cover the importance of continuous learning throughout AI transitions such as what the tool can reliably do, which decisions it will handle versus a person, and what guardrails or policies will govern the use of AI.

4.1 Managing Uncertainty and Risk in AI-Enabled Change

AI-enabled change introduces uncertainties you can't fully predict or control. Building resilience into the change requires anticipating what you don't know and designing systems that can adapt when things go wrong.

Build the ability to recognize that AI-enabled change is inherently more uncertain than traditional change. Cover specific risks such as AI behaves unpredictably, unintended consequences emerge, safeguards fail, algorithmic bias, automation creating new problems, or AI overriding human judgment inappropriately. Introduce resilience practices such as monitoring and feedback loops to surface problems early, contingency plans for when AI doesn't work as expected, maintaining human oversight, and designing for reversibility.

Data, AI, and Decision Accountability

Data, AI, and Decision Accountability is a 2–4 hour Specialized Credential for professionals who want to develop their ability to use data and AI-enabled recommendations responsibly in organizational decisions.

1.1 Role and Limits of Data and AI

Data and AI can improve speed, pattern recognition, consistency, and scale, but they do not replace human judgment, context, values, or accountability.

Build learners' understanding of how data and AI can support decisions while leaving responsibility with human decision-makers.

2.1 Evidence Quality, Bias and Accountability

Data- or AI-supported decisions can be weakened by incomplete data, historical bias, model limitations, proxy variables, automation bias, unclear ownership, or poor transparency.

Develop learners' ability to identify where decision inputs need additional scrutiny and where accountability risks remain.

3.1 Traceable Decision Rationale

When data or AI informs a decision, people need to know what evidence was used, what limitations were considered, who decided, and why the action was chosen.

Develop learners' ability to document decision rationale in a way that supports trust, auditability, learning, and responsible action.

4.1 Human Oversight in Practice

Human oversight is easy to establish but hard to keep effective. A reviewer handling two hundred recommendations a day is not exercising judgment on any of them; a reviewer shown the recommendation before forming their own view is anchored to it; a reviewer measured on throughput will approve. These are the consequences of how the loop was built.

Develop learners' ability to judge whether human judgement in an AI-augmented decision process is real, and to identify what would restore it. Cover the observable signals of a degraded loop and the levers that restore it.

SPECIAL THANKS

ICA would like to give special thanks to the contributors:

Kerri Sutey · Paulo Dias · Christina Hartikainen · Tricia
Broderick · Mike Lyons · Greg Pfister · Gabriel Opoku