
QUALITY-FIRST AI:
A Roadmap for AI Transformation
Dr. Nicole Radziwill &
T. Scott Clendaniel
Based on a chapter of the forthcoming book The AI-Native Leader (2026)
1

QUALITY-FIRST AI: A Roadmap for AI Transformation

Leaders can't just distribute new Artificial Intelligence (AI) tools and licenses, run a few training sessions, and hope for the best. Becoming an AI-native company requires all of us in an organization to fundamentally change how we work together, building our ability to see opportunities that were previously invisible while staying alert to risks that could create quality problems later.

These opportunities exist at every level of the Gen AI Value Pyramid: individual improvements, collective intelligence, transformation and growth, and visionary innovation. This chapter focuses on how to do it: building a solid **foundation** (ideally a quality management system), **activating** the workforce (through five key habits and practices), and establishing feedback loops that increase maturity and lead to **self-renewal**.

Figure 1: The Gen AI Value Pyramid (adapted from McLees, Radziwill & Satell, 2024)

THE AI QUALITY GAP

In the rush to AI-enable our workforces and workplaces, many organizations are unknowingly pushing serious problems into the future. We call this the **AI Quality Gap**: the difference between the assumed quality suggested by AI demos, benchmarks, and vendor promises and the sustained, trustworthy quality organizations actually need to experience once AI is deployed to drive real business processes at scale. It is the mismatch between how good AI looks on the surface and the actual reliability, consistency, and business outcomes it delivers in practice.

This gap manifests in several interconnected ways.

- **The objectives gap.** Too many organizations are using AI as a "hammer in search of a nail." The "AI-Native" approach is supposed to prioritize the application of AI techniques once clear objectives are established, not a replacement for having those objectives clearly defined.
- **The scalability gap.** Organizations rarely have a clear sense of how large their AI initiatives need to grow. They either build prototypes that have little chance of being scaled or over-engineer a process that works fine for individual employees but fails upon any attempt to scale. Having an established staged-gate approach can limit these issues.
- **The content gap.** AI-generated code and content can look impressively productive at first glance, but beauty is sometimes only skin deep. Without adequate review, code can introduce bugs, instability, and what many practitioners now call "AI slop" - output that increases technical debt and erodes long-term quality.
- **The delivery gap**. There is a significant qualitative difference between a working AI prototype and a production-ready system. Getting from "80% quality" to "99% quality" is disproportionately hard with AI. Organizations consistently underestimate this last mile.
- **The model performance gap.** Models that perform beautifully on clean benchmark data often struggle with messy, real-world scenarios. Research has shown that large language models that look excellent in evaluations can fail on a majority of real tasks without careful specialization and tuning. Similarly, a solution built assuming one model may need to be carefully monitored because there is no guarantee performance will improve as new models are released.
- **The financial lifecycle gap.** Quality means meeting stated and implied needs, and sometimes those needs only emerge well after deployment. A company may build an impressive AI product only to discover it is merely a wrapper around functionality that a general-purpose assistant has recently subsumed. Or a team builds, develops, and deploys a solution without modeling lifecycle costs, creating a business model that can't sustain itself. Or, an organization fails to consider externalities, like: what happens when the cost of AI services inevitably increases? What happens when regulatory requirements change?

- **The adoption gap.** User expectations of AI, shaped by marketing and dazzling demos, frequently exceed the actual experience (especially for AI capabilities embedded in traditional enterprise software). When reality falls short of the promise, trust erodes (sometimes permanently). More significantly, people just lose interest when they have to struggle to get an AI capability to work.
- **The human and organizational gap.** A widening divide is emerging between teams getting extraordinary results from AI and teams that tried it but have little to show. This divide is driven by differences in skills, clarity of use cases, and strength of collaboration, not tools. Whenever a technology changes the power dynamics between people, relationships shift and this has a direct impact on output and outcomes.
- **The human nature gap.** Perhaps the most insidious: in the early years of workplace AI (2023 to early 2026), pulse surveys reported that nearly *everyone* was excited about AI. But *of course* they were. Job security often depends on supporting your company's mandates, and *it was too politically risky to say anything else*. Many executive leaders are still not getting an accurate read on whether people are truly transforming their work, or just trying to shift the spotlight to some other person or department.
- **The perception gap.** Even when managers want to address employee concerns regarding AI, they focus on the potential loss of jobs. Deloitte's recent research (Deloitte 2025) indicates that employees fear loss of personal autonomy in their jobs almost as much as job loss. Yet this fear is frequently ignored.

Each of these gaps represents a risk to financial performance. Undetected quality problems compound over time, manifesting as rework, customer attrition, compliance gaps, and eroded margins. When leaders recognize that the AI Quality Gap is not a technical problem but a systemic one, the case for a quality-first approach becomes clear.

QUALITY-FIRST AI

Quality-First AI is a deliberate choice any leader can make. It means grounding every step of AI adoption and enablement in proven quality management principles rather than chasing technology for its own sake. It starts by adopting a quality-first philosophy, and continues by establishing a clear understanding of organizational context: who you serve, how work flows, and what outcomes matter. Using foundational frameworks like the Baldrige Excellence Framework or ISO 9001 can help your company ensure that AI initiatives are rooted in the reality of how the business actually operates.

The Philosophy

Quality-first leaders think in **systems**. They understand **variation** in their data and processes. They build **knowledge** through theory and testing rather than blind automation. And they account for the **psychology** of how people interact with intelligent systems, because leveraging new tools to accomplish

new things is ultimately a human challenge. These are the four elements of W. Edwards Deming's System of Profound Knowledge, covered comprehensively in Watson (2018). Rather than automating isolated workflows or adding intelligence to tasks that may or may not need it, Quality-First AI targets constraints that actually limit performance:

- It **demands data integrity** where decisions depend on it, recognizing that machine learning models inherit the flaws of their training data.
- It **insists on validation** at every stage of the AI lifecycle, from data preparation through deployment, verifying quality before passing work downstream.
- It **relies heavily on "own process ownership"** (自工程完結, or "ji kotei kanketsu") for agentic AI, where it is especially critical because compounding errors across chained steps can silently (and catastrophically) degrade outcomes.

The executive case is straightforward: technology alone does not produce results, *collective learning* does. Organizations that embed AI with quality systems as their guardrails avoid the bottleneck paradox, where speeding up one step simply overwhelms the next. They build trust, reduce costly failures, and create the feedback loops between humans and machines that drive sustainable, compounding improvement rather than surface-level efficiency gains.

When Other Priorities Should Lead

A Quality-First approach to AI should not, however, be a universal mandate. As much as the authors believe in quality as a foundational principle for AI, sometimes, it makes more business sense to focus *less* on quality (for example, when you'll lose an opportunity if you're not first to market). But this is a decision that must be taken carefully.

Most successful AI deployments require various *tradeoffs*. You might adopt a quality-first position for your flagship customer experience and cost-first for background automation. The real strategic skill is being able to differentiate which priority fits which scenario.

- **Prioritize speed** when the value of output degrades rapidly with time. In real-time fraud detection, content moderation at scale, or live support triage, a 90%-accurate answer in 200 milliseconds can be far more valuable than a 99%-accurate answer in 10 seconds. When the cost of waiting exceeds the cost of expected errors, move fast.
- **Prioritize cost reduction** for high-volume, low-stakes use cases. Auto-generating product tags, bulk-categorizing support tickets, or summarizing internal meeting notes are situations where "good enough" at a fraction of the cost is a rational tradeoff. Remember: quality is about meeting needs, and if the customer needs lower cost and accepts the risk of occasional failures, that *is* quality.

- **Prioritize scale** when broad coverage matters more than depth. Building a multilingual knowledge base or providing basic AI assistance across thousands of small businesses can create more total value through reach than by providing a perfect customer experience for just a few.
- **Prioritize exploration** when you are still discovering what to build. In early prototyping or testing whether an AI use case has product-market fit, imposing rigorous quality discipline before validating the concept is premature optimization.
- **Prioritize user agency** when building tools for experts. A developer using an AI code assistant or a data scientist doing exploratory analysis may prefer several rough options over one curated answer. They want leverage, not curation.
- **Prioritize resilience** when failures could be catastrophic. In safety-critical domains like aviation or medical devices, the priority is ensuring the system *never* produces a dangerous output. The goal is limiting downside risk rather than maximizing the upside.

The quality-first leader's job is to know which priority fits where, and to make those tradeoffs *explicit* rather than accidental.

TRANSFORMING YOUR ORGANIZATION

Above all, leadership must get everyone on the same page. Making sure people explicitly understand the work environment, its customers, its goals, and its business concepts in consistent ways provides the essential context an AI assistant needs to understand an organization. Intentionally building this context stops people from inadvertently pulling in different directions. When humans and AI are aligned around leadership intent, strategy, and customer needs, organizations stop wasting time and money on unproductive work.

Quality Management Systems (QMS) as the Foundation

A quality management system (QMS) provides an excellent foundation for AI because it already defines how the organization works: its processes, roles, controls, metrics, risks, and improvement loops. This gives AI assistants grounded, governed context instead of a random pile of documents to make sense of. AI performs better when its context layer includes semantics, approved sources, operating rules, and observability, rather than just the raw files and folders (or “digital exhaust”) produced during operations.

A QMS organizes the business around consistent procedures, documented information, audits, corrective actions, and performance measurement. These are *exactly* the signals an organization-wide AI assistant needs to answer questions reliably. A QMS also creates a common language for decision-making, enabling an AI assistant to distinguish policy from practice, exceptions from standards, and symptoms from root causes.

Among the available frameworks, the **Baldrige Excellence Framework (BEF)** administered by the U.S. National Institute of Standards and Technology (NIST) stands out as a particularly solid foundation for AI context (NIST, 2026). Curated since 1988 by hundreds of industry experts across multiple generations of the workforce, Baldrige is a book of probing questions that leaders can use to define the *why, who, what, and how* of a business. Most companies ignore this step, and rely on everyone to individually create their own mental models of the organization… hoping, for the most part, that those mental models are similar.

The seven categories described in the Baldrige Excellence Framework (Leadership, Strategy, Customers, Measurement and Knowledge Management, Workforce, Operations, and Results) help leaders define the organization as a system of interconnected parts. That matters because AI needs to understand how people and processes are expected to work together: priorities, decision rights, customer expectations, workflows, and success metrics in a structured way that typical knowledge repositories rarely provide. BEF is also non-prescriptive, so it works across industries and operating models while still surfacing the key context AI needs to respond consistently and responsibly. Notably, NIST has explicitly connected Baldrige to AI-relevant issues including strategy, knowledge management, innovation, risk, and responsible use.

Activating the Workforce: The Lighthouse Model

Once you have a quality framework that enables AI assistants to recognize your company as a system with interdependent components, make sure people are developing the habits that will empower them to see opportunities and act on them collaboratively, with full awareness of risk. The Lighthouse Model uses that foundation as the basis for reaching the pinnacle of quality: customer-centric innovation, meaningful business results, and a self-renewing workforce that adapts to new tools and capabilities quickly enough.

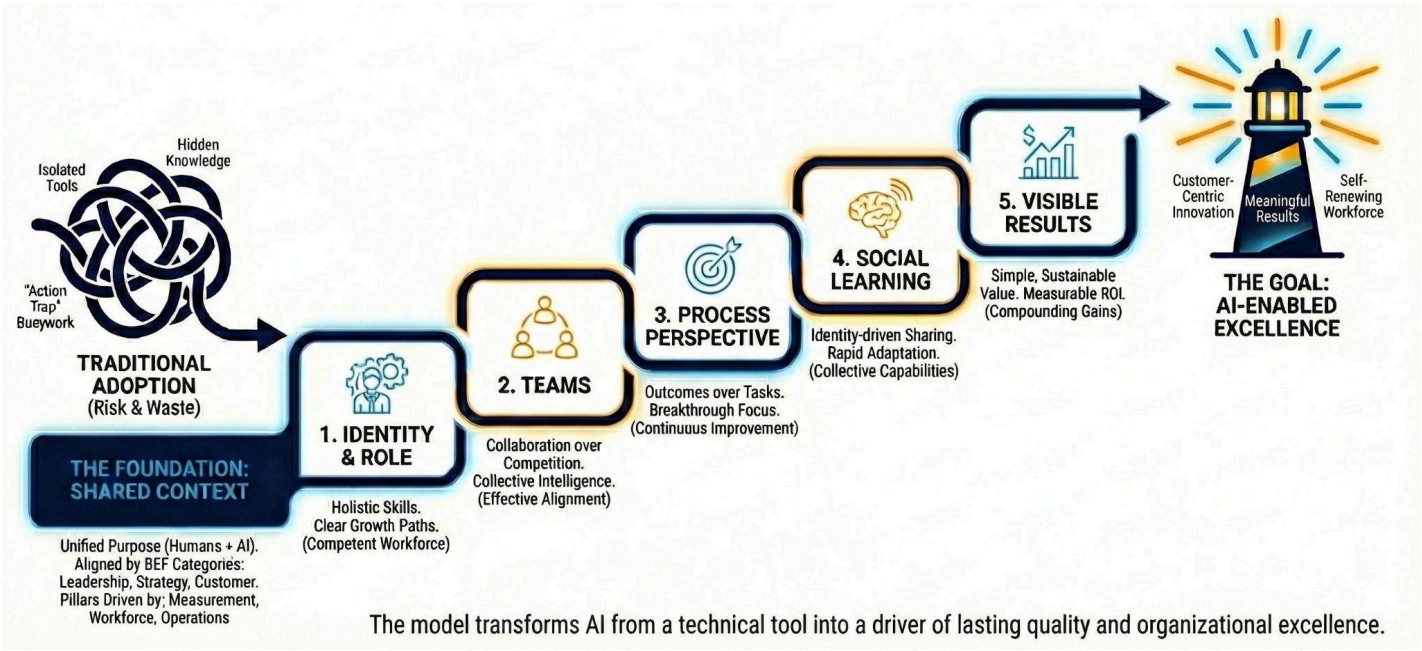

Figure 2: The 5 habits & practices that align the workforce around AI transformation.

The 5 Pillars

- **Pillar 1: Identity and Role.** Not everyone is a natural experimenter, and not everyone will gravitate to AI tools and assistants. Leaders should more broadly develop their team's skills around native competencies and aptitudes so employees know how to *think* with AI rather than trying to train everyone to become an expert in specific tools that change constantly. When people develop identities that capture what they're good at with AI, the team knows exactly who to ask for support when a challenge arises, preventing missed opportunities. (We recommend using Team-X AI Personas: Experimenter, Shapeshifter, Optimizer, Stabilizer, Observer, Integrator, and Steward. You can find out more at https://team-x.ai.)

- **Pillar 2: Teams.** Work on unblocking the human barriers to fast, frictionless application of AI. If we treat AI as a solo sport, a few super-users gain disproportionate influence while everyone else falls behind. By focusing on team-level collaboration, everyone can grow into the role their team needs, and the whole group learns (and gets results) much faster.

- **Pillar 3: Process Perspective.** Focus on the goals and outcomes you need to achieve *together* over tasks and actions people do *independently*. Primarily, this requires leaders to systematically evaluate, prioritize, govern, and measure every AI initiative they run.[1] This also means avoiding the "Action Trap" of implementing AI just because someone *said* to use AI, and instead keeping attention on the actual business and financial results. This opens the door for breakthrough improvements rather than indiscriminate productivity gains that create new bottlenecks.

- **Pillar 4: Social Learning.** Cultivate the habit of connecting people who *have* problems with people who know how to *solve* them using emerging technologies. AI moves too fast for rigid training programs and committee-driven governance; active feedback loops are needed to keep people connected to new solutions. The goal, then, is building a culture where colleagues quickly and naturally share wins, which creates a self-renewing and adaptable workforce. But this won't happen on its own... It's up to leaders to create the time, space, and incentives for these exchanges to happen naturally.

- **Pillar 5: Visible Results.** Focus on simple solutions that bring sustainable results, and provide a platform for everyone's contributions to be seen. Executives need to clearly see the value AI brings in real time, employees need recognition for collaborative solutions, and the organization needs to quickly distinguish what is working from what is not. This is another capability that depends on healthy and active feedback loops.

[1] Nautilus AI Academy offers interactive sessions to help leaders build their AI Roadmap by applying the process perspective to their own initiatives. Contact Scott.Clendaniel@nautilus-ai.com for more information or to enroll in live events or Maven courses.

The Activating Forces

Across our client work at Team-X, informed by working with neurodivergent and cognitively diverse teams, we found that these five pillars are activated by three categories of **action** that turn AI adoption from haphazard to habitual.

- **Building shared understanding** is the foundation. When people develop common language and mental models around AI, they stop reinventing wheels in isolation and start building on each other's progress.
- **Managing friction** is about removing damaging barriers while keeping protective ones. Standards and regulations, for example, can prevent reckless action. The goal is to unblock collaboration and streamline how AI gets applied to real outcomes—not to eliminate all resistance.
- **Managing cognitive load** is what makes adoption sustainable. When people have clear roles and can see how their work benefits the collective, they don't have to carry everything in their heads. AI changes fast; if every decision requires deep analysis, people freeze. Reducing cognitive load helps people act with confidence.

These forces interact dynamically. Sufficient shared understanding reduces cognitive load. Too much drives analysis paralysis. Lower friction accelerates learning. Higher friction reduces risk. When all three are managed together, the five pillars stop being aspirational and start producing compounding results.

EVALUATING YOUR PROGRESS

Measuring AI adoption and enablement is vital for translating technological investments into tangible business value. While tracking tool use and employee proficiency is a first step, it remains difficult for companies to accurately calculate ROI, pinpoint leadership gaps, and ensure AI genuinely enhances productivity beyond the scale of the individual. While counting active software licenses is straightforward, quantifying intangible benefits like improved creativity, better decision-making, or cognitive offloading is highly complex. Furthermore, the prevalence of "shadow AI" (unsanctioned, untracked usage by employees) obscures true adoption rates. Lacking standardized KPIs and navigating rapid technological shifts, many leaders are left relying on intuition rather than concrete data to guide their workforce enablement and scaling strategies. There’s a way around this.

Assessing Your Readiness

Use the "What's Your Organization's AI IQ?" assessment[2] from Team-X AI to find out where your organization stands with respect to the 5 habits and practices described by the pillars in the lighthouse model. In 10 questions, it quickly evaluates your status in each of the 5 areas: identity and role clarity, team dynamics, goal-oriented process thinking, social learning, and results visibility. This will measure the extent to which your teams are building the habits and practices you need for deep AI enablement.

Each question is scored from 0 to 3, yielding a maximum of 6 points per section and 30 points overall. In categories where only 0-2 points are reported, focus on building capabilities. When 4-6 points are reported in a category, talk with the teams to find out winning practices that can be replicated across the organization to enhance maturity.

Estimating Maturity Levels

The assessment roughly maps to five standard maturity levels.

- At **Level 1 (Initial)**, AI adoption is chaotic and driven by individual heroics, with results invisible to the broader organization.
- At **Level 2 (Managed)**, some teams have basic practices in place, but they are reactive rather than planned.
- At **Level 3 (Defined)**, the organization has standard processes for all five areas, with regular reviews focused on outcomes.
- At **Level 4 (Quantitatively Managed)**, practices are measured and data-driven, with leadership able to quantify ROI.
- At **Level 5 (Optimizing)**, AI enablement is a core organizational capability that continuously evolves, with peer learning deeply embedded in culture and results visibility driving real-time strategic decisions.

Note that organizations *rarely* jump levels. The key transitions are:

- Establish *any* regular practice to move from Level 1 to 2,
- Make those practices consistent and habitual to move from Level 2 to Level 3,
- Commit to meaningful measurement to go from Level 3 to Level 4, and
- Create self-renewing feedback loops to leap to Level 5.

An individual can take the assessment for a quick self-read or it can be deployed organization-wide (with 75% or greater participation) to give leaders a clear picture of where teams are ahead or behind.

[2] Team-X AI (https://team-x.ai) administers AI IQ on behalf of organizations as a survey or in "radar" mode where the workforce is continuously and adaptively sampled to determine how AI maturity is changing and get ROI signals that help leaders make informed decisions about resourcing.

What's your Organization's AI IQ? *From https://team-x.ai*

Question	0	1	2	3
Identity & Role Clarity				
Do you have a clear sense of what you're uniquely good at in an AI-enabled workflow, and do the people you work with know how to leverage those strengths?	No to both	Unclear / partial	Clear but not shared	Yes to both
In the past year, have you had a conversation with a colleague or leader about how AI might change your role and how that could benefit your work?	No	Briefly mentioned	Substantive discussion	Ongoing alignment
Team Relationships & Communication				
In the past quarter, has your team had at least one structured conversation focused on improving how you work together with AI?	Never	Once	2-3 times	Ongoing practice
When there is friction or confusion on your team, how quickly do you address it together?	We don't address it	Eventually	Within a week	Same day
Process Perspective – Why Over How				
When introducing new AI capabilities, does your team focus primarily on goals (e.g., improving client satisfaction, reducing errors) over technology (e.g., applying a new tool)?	Rarely goal-focused	Sometimes	Usually	Always goal-focused
In the past quarter, has your team discussed whether your AI-enabled tasks and workflows are helping you reach your goals faster or just creating new bottlenecks?	No	Informal discussion	Formal review once	Regular practice
Social Learning				
In the past month, have you shared something useful you learned about working with AI, or learned from someone else's experience?	Neither	Yes to one, informally	Yes to both, informally	Yes, through a formal process
Does your organization actively encourage AI learning and knowledge sharing by providing time and resources for collaboration?	No incentive	Occasional encouragement	Frequent encouragement	Rewards and/or time allocated
Visible Results				
When you or your team achieves something meaningful with AI, is there a clear way to share that success and be recognized for it?	No mechanism	Informal only	Clear but underused	Active and valued
Do you believe the leaders of your organization can *see* what you and your team contribute with AI and how it benefits the organization?	Definitely no	Probably no	Probably yes	Definitely yes

Rate each item from 0 to 3 using the descriptions provided. Total your score across all 10 questions (max: 30).

CONCLUSION

Quality-First AI is not about slowing down. It's about committing to the **long game** by making the organizational infrastructure for AI (the shared context, human habits and practices, and disciplined feedback loops) a priority. This allows the technology to deliver on its promise sustainably. The organizations that will lead in the AI era will be the ones that adopt wisely, embedding AI within systems designed to learn, adapt, and improve.

The quality gap is real, but it is closable... as long as you, the leader, is willing to commit AI transformation as a marathon where you're building organizational excellence over time - not just rolling out a new technology or training program in an isolated way.

REFERENCES

Deloitte. "2025 Global Human Capital Trends: Turning Tensions into Triumphs - Helping Leaders Transform Uncertainty into Opportunity." 2025.

McLees, Todd, Nicole Radziwill, and Greg Satell. “How to Create Value Systematically with Gen AI.” 2024. Harvard Business Review. December 6, 2024. https://hbr.org/2024/12/how-to-create-value-systematically-with-gen-ai

National Institute of Standards and Technology (NIST). Baldrige Excellence Framework (Business/Nonprofit): Proven leadership and management practices for high performance. Available from https://www.nist.gov/baldrige/publications/baldrige-excellence-framework/businessnonprofit, 2026.

Watson, Gregory H. The Theory and Practice of Profound Knowledge: An Inquiry into Quality and Strategy Management. Oklahoma State University, 2018.

ABOUT THE AUTHORS

Dr. Nicole M. Radziwill

Dr. Nicole Radziwill is Co-Founder and CTAIO of team-x.ai, helping leaders flourish with AI at scale. A former SVP, Chief Data Scientist, and tenured Associate Professor of Data Science and Production Systems with over 5000 students in data science and applied machine learning, Nicole is an Academician in the International Academy of Quality (IAQ), and a Fellow of the American Society for Quality (ASQ). She has a PhD in Quality Systems from Indiana State University and is the author of data science and statistics textbooks used in 30+ universities and 12+ countries. Previous works include Connected, Intelligent, Automated: The Definitive Guide to Digital Transformation and Quality 4.0 (ASQ Quality Press, 2020), Data, Strategy, Culture & Power (2024) and Machine Learning Fundamentals (ASQ Quality Press, 2026).

Driving Value with AI since 1997!

T. Scott Clendaniel

T. Scott Clendaniel, Chief Learning Officer of Nautilus-AI Academy, is one of the world's top AI educators and experts in Artificial Intelligence Impact and ROI. An advocate of "AI access for all" spanning four decades, Scott's methods have catalyzed business results for over 3.1 million students in 30 countries. His clients have included Wharton School of Business, NVIDIA, Google, Gartner, and the International Olympic Committee. He has been a consultant, speaker, and educator for Booz Allen, Gartner, the University of Pennsylvania, Wharton, and Johns Hopkins. He finds the patterns others miss and builds proprietary techniques that find unique value for his clients.

Driving Value with AI since 1986!

ACKNOWLEDGEMENTS

This work is strongly informed by work with Nicole's collaborators Todd McLees and Greg Satell, the content in *Connected, Intelligent, Automated: The Definitive Guide to Digital Transformation & Quality 4.0* (ASQ Quality Press, 2020) and our collective experiences advising the senior and executive clients we supported as independent consultants and for Nautilus AI, Gartner, Inc., Ultranauts, Inc., Team-X AI and clients of those organizations.