The AI Playbook

From Buzzword to Business Strategy

A practical guide to understanding, adopting, and governing AI

Rick Becker

April 2026

resume.beckerz.com

rick@beckerz.com

About This Paper

This paper was written for business leaders who would like a condensed overview of AI without wading through technical documentation or vendor marketing. It is not a product pitch. It is not a prediction about where AI is headed in ten years. It is a practical orientation: the history, the vocabulary, the opportunity, the risk, and a framework for getting started. I wrote this to support a busy executive who wants the TLDR version of what AI is and how to real decisions about AI in their organization right now.

The frameworks here, including the Crawl/Walk/Run adoption model and the risk governance approach, reflect the author's direct experience designing and deploying AI-assisted systems across insurance operations, enterprise platforms, and solo full-stack builds using AI development tools. The goal is to give leaders enough grounding to ask better questions, evaluate vendors more critically, and move forward with confidence rather than waiting for the landscape to "settle down." It won't.

Introduction

Artificial intelligence is not new. The concept of machines mimicking human intelligence has been around for over 70 years. What is new is the version of AI that has captured the attention of boardrooms, newsrooms, and dinner tables around the world. The technology that people are talking about today, generative AI, is the product of decades of incremental progress in computing, programming, machine learning, and deep learning. Each breakthrough built on the last. What makes this moment different is not the intelligence itself but the speed at which capabilities are advancing, the breadth of industries being affected, and the accessibility that puts frontier AI in the hands of anyone with a browser.

This white paper is designed for business leaders who want to understand AI without the hype, the jargon, or the fear. It covers what AI actually is, how to use it personally and professionally, where the real business opportunities are, what the risks look like, and how to adopt it in a way that is smart, governed, and aligned with your organizational strategy.

The goal is simple: by the time you finish reading, AI should feel understandable and actionable rather than abstract or overwhelming.

Section I: AI Is Not New. It Is Just Loud Now.

The history of artificial intelligence starts in 1950, when Alan Turing published a paper asking a deceptively simple question: "Can machines think?"[1] That question started decades of research and experimentation.

In the 1960s, researchers at MIT built ELIZA, widely considered the first chatbot.[2] ELIZA could simulate conversation by recognizing patterns in text and generating responses. It was rudimentary, but it demonstrated that machines could interact with humans in natural language.

By the 1980s, the field had moved to expert systems, which were AI programs designed to follow complex sets of rules to make decisions in specialized domains. There was enormous investment and enormous hype. But the technology could not deliver on its promises. Funding dried up, research stalled, and the field entered what is now called the "AI Winter."[3] This is worth remembering: the pattern of hype followed by disappointment is not new. It has happened before.

The 1990s brought a shift toward machine learning, where instead of programming explicit rules, researchers taught systems to learn from data. IBM's Deep Blue demonstrated this evolution dramatically in 1997 when it defeated world chess champion Garry Kasparov.[4] The machine was not following a script. It was analyzing positions and making decisions.

---

[1]Turing, A.M. (1950). "Computing Machinery and Intelligence." Mind, 59(236), 433-460.
[2]Weizenbaum, J. (1966). "ELIZA: A Computer Program for the Study of Natural Language Communication Between Man and Machine." Communications of the ACM, 9(1), 36-45.
[3]The AI Winter refers to periods of reduced funding and interest in AI research, primarily 1974-1980 and 1987-1993. See Crevier, D. (1993). AI: The Tumultuous History of the Search for Artificial Intelligence. Basic Books.
[4]IBM Deep Blue defeated World Chess Champion Garry Kasparov on May 11, 1997. See IBM Archives: https://www.ibm.com/history/deep-blue

The 2010s saw the rise of deep learning, a subset of machine learning that uses layered neural networks loosely inspired by the human brain. This is the era that gave us Siri, IBM Watson winning Jeopardy,[5] and Google DeepMind's AlphaGo defeating the world champion in a game that experts had said computers could not win.[6]

And then, in November 2022, OpenAI released ChatGPT.[7] That was the moment AI went from a technology industry conversation to a global cultural phenomenon.

What Enabled This Moment

Three things converged to create the generative AI moment we are living through today.

The Transformer Architecture. In 2017, a team at Google published a paper titled "Attention Is All You Need" that introduced the transformer, a fundamentally new way for AI to process language.

Before transformers, models read text sequentially, one word at a time. The transformer's self-attention mechanism allows a model to look at an entire passage simultaneously and understand the relationships between words, capturing context, meaning, and nuance. This single architectural breakthrough is the engine behind every major large language model: GPT, Claude, Gemini, Llama, and others. Without the transformer, none of what we are experiencing today would be possible.[8]

Compute and Data at Scale. A breakthrough on paper is not enough. You need the horsepower to build it. GPUs, originally designed for video game graphics, turned out to be perfectly suited for the parallel mathematical processing that AI training requires.

Cloud computing from AWS, Azure, and Google Cloud made that processing power accessible without owning a data center.[9] And the internet produced an unprecedented volume of text, images, code, and video that could serve as training data. The blueprint existed. Now the construction crew and raw materials did too.

Accessibility. GPT-3 existed in 2020 and it was remarkable. But it was locked behind an API that only developers could access. OpenAI's decision to wrap the technology in a free, simple chat interface changed everything. Suddenly, anyone could interact with frontier AI. No coding. No technical background. Just a browser and a question. That is what turned a technical achievement into a cultural moment.

---

[5]IBM Watson defeated Jeopardy! champions Ken Jennings and Brad Rutter in February 2011. Apple launched Siri in October 2011.

[6]Silver, D. et al. (2016). "Mastering the game of Go with deep neural networks and tree search." Nature, 529, 484-489. Google DeepMind.

[7]OpenAI launched ChatGPT on November 30, 2022. It reached 100 million users within two months, making it the fastest-growing consumer application in history at that time.

[8]Vaswani, A. et al. (2017). "Attention Is All You Need." Advances in Neural Information Processing Systems 30 (NIPS 2017). Google Brain / Google Research.

[9]NVIDIA GPU architecture for AI training. See Raina, R., Madhavan, A., and Ng, A.Y. (2009). "Large-scale Deep Unsupervised Learning using Graphics Processors." ICML 2009.

Section II: Demystifying the Jargon

With this cultural movement came a new language. The new terminology makes the technology feel more complicated than it needs to be. This section breaks down the key terms in plain English.

How It All Fits Together[10]

Think of AI as a set of nested layers, like Russian nesting dolls.

Artificial Intelligence is the broadest category. It refers to any machine that can perform tasks that typically require human intelligence. This concept has been around since the 1950s.

Machine Learning sits inside AI. These are systems that learn from data instead of being explicitly programmed. Rather than writing rules for every decision, you feed the system data and let it find patterns.

Deep Learning sits inside machine learning. It uses layered neural networks, loosely inspired by the human brain, to recognize increasingly complex patterns. This is what made image recognition, voice assistants, and language translation possible at scale.

Generative AI sits at the center. These are models that do not just analyze or classify. They create new content: text, images, code, audio. This is ChatGPT, Claude, Midjourney, Copilot. This is what everyone is talking about. But notice that it is built on top of 70 years of everything else.

The Vocabulary You Will Hear

Natural Language Processing (NLP) is how machines understand and generate human language. Every time you talk to Siri or type a question into an AI chatbot, NLP is what makes it work.

Computer Vision is how machines interpret images and video. It is what allows a self-driving car to recognize a stop sign or your phone to unlock with your face.

Training Data is the information a model learns from. The critical principle here is garbage in, garbage out. If the training data is biased, incomplete, or wrong, the model will be too.

A Model is the trained system itself. It takes your input, processes it through the patterns it has learned, and produces an output.

Tokens are how AI reads text. A token is roughly one word or piece of a word. When people say a model has a "128K context window," they are talking about how many tokens it can process at once.

Agentic AI is the next evolution. While generative AI creates content in response to a prompt, agentic AI goes further. It plans, makes decisions, uses tools, and takes action autonomously. Generative AI can write the email. Agentic AI can write the email, send it, check for a reply, and

[10] Definitions in this section are adapted from standard references in the field, including Russell, S. and Norvig, P. (2021). Artificial Intelligence: A Modern Approach, 4th Edition. Pearson; and Goodfellow, I., Bengio, Y., and Courville, A. (2016). Deep Learning. MIT Press.

schedule a follow-up meeting. This distinction becomes critical when we discuss the future of AI adoption later in this paper.

Narrow AI is what exists today. It is AI that excels at specific tasks but cannot generalize across domains the way a human can. Every AI system in production today, including the most advanced models, is narrow AI.

AGI (Artificial General Intelligence) is the theoretical concept of a machine that can reason across every domain the way a human can. It does not exist yet. Some researchers say it is five years away, some say fifty, some say never[11]. It is not what this paper is about. Everything here concerns the AI that exists right now and what you can do with it.

11

https://medium.com/@cognidownunder/agi-still-years-away-despite-tech-leaders-bold-promises-for-2026-146c9780af65

Section III: Making It Real

Before we talk about enterprise strategy, it is worth grounding this in the personal. AI is not a technology that only matters to IT departments. It is something you can use today, right now, to make your life easier.

Personal Use Cases

- **Writing and editing.** Draft emails, thank you notes, complaint letters, or anything you have been putting off. AI can generate a first draft in seconds, and you refine it.
- **Research and learning.** Ask AI to explain anything in plain language. Quantum physics, your kid's algebra homework, how to fix a leaky faucet. Tell it to explain like you are five, and it will.
- **Health and wellness.** Meal plans based on dietary restrictions, workout routines, symptom lookups. Not a replacement for your doctor, but an excellent starting point.
- **Travel.** Give AI your dates, your budget, and what you like to do. It will build you an itinerary.
- **Household.** Budgeting, recipes from whatever is in your fridge, scheduling, organizing.

Professional Use Cases

- **Drafting and summarization.** Write emails, summarize 40-page reports (TLDR), clean up meeting notes. This is where most professionals feel the immediate impact.
- **Data analysis.** Upload a spreadsheet and ask AI to find trends, outliers, or build a chart. It is not perfect, but it is a phenomenal starting point.
- **Brainstorming and strategy.** Use AI as a thought partner. Give it your business problem and ask it to poke holes in your plan. It is like having a smart colleague available around the clock who never gets tired of your questions.
- **Coding assistance.** Even if you are not a developer, AI can help you write formulas, automate tasks, and build simple tools.
- **Content creation.** Presentations, marketing copy, social media posts, internal communications.

Precautions: Before you Dive in

- **Make It Your Own** - AI is a starting point, not a final product. Ensure the output represents your thoughts and ideas. Where you use AI-generated content, note it appropriately.
- **Verify Everything** - AI makes mistakes. Check facts, review numbers, and never blindly trust the output. You are responsible for what you put your name on.
- **Think About Privacy** - Consider what you are putting into AI. Customer data, company financials, proprietary information. If you are using a public AI model, assume it is not private.

- **Follow Your Company's Guidelines** - Ensure you are in compliance with your organization's security policies and AI usage rules. If those rules do not exist yet, ask for them.

The bottom line: you do not need permission to start. Pick one thing from either list and try it today.

Section IV: Where AI Creates Business Value

Everything in the previous section is about individual productivity. This section is about what happens when you scale that across an entire organization.

As a starting point consider these categories of business value that AI can deliver today.

Staff Augmentation. This is AI as a force multiplier. Every employee in your company gets access to drafting, summarization, data analysis, and coding assistance. You are not replacing people. You are making your existing team dramatically more effective.

Software Integration. This is AI embedded in the tools your teams already use. Microsoft Copilot in Office. Salesforce Einstein in your CRM. You do not have to build anything. The AI comes to where the work already happens.

Content Generation. At scale, this means marketing teams producing campaigns in hours instead of weeks, internal comms writing company-wide updates faster, and training teams building course materials at a fraction of the previous cost.

Process Automation. This is where you start eliminating entire steps in a workflow. Invoice processing, contract review, employee onboarding paperwork. AI does not just help someone do the task faster. It does the task.

Forecasting and Analytics. Predictive modeling, demand planning, risk assessment. AI looks at your data and tells you what is likely to happen next, not just what happened last quarter.

Consultative and Decision Support. Feed AI your market data, your competitive landscape, your financial projections, and ask it to stress-test your strategy. It is not making the decision for you. It is giving you a better-informed decision.

Knowledge Management. Every company has institutional knowledge trapped in people's heads, in old documents, in email threads nobody can find. AI can make all of that searchable and accessible. It accelerates onboarding, reduces knowledge loss when people leave, and makes your organization's collective intelligence available to everyone. This should be the first enterprise AI deployment for good reason: low risk, high visibility, immediate ROI.

Customer Experience. Chatbots, personalization, self-service portals. The goal is not to replace your support team. It is to handle the routine stuff instantly so your people can focus on the complex, high-value interactions.

The ROI Framework: Where Do You Start?

The question every leader asks next is: where do I start? The answer is a simple prioritization framework based on two dimensions: business impact and implementation effort.

Start Here (High Impact, Low Effort): Knowledge management, content generation, and staff augmentation. These deliver significant value with tools that already exist. You do not need to build anything custom.

Strategic Bets (High Impact, High Effort): Process automation, forecasting and analytics, and software integration. These are worth doing, but they require time, budget, and cross-functional coordination. You get here after building confidence and capability with the easier wins.

Quick Wins (Low Impact, Low Effort): Customer experience chatbots and basic decision support tools. They will not transform your business overnight, but they build momentum and get people comfortable with AI in production.

Reconsider (Low Impact, High Effort): Custom-built AI solutions without a clear business case. If someone proposes building a proprietary model and cannot articulate the ROI in one sentence, push back. This quadrant is where money goes to die.

The key insight: start in the top-left quadrant, build capability and confidence, then move to the top-right. Do not skip to strategic bets before you have proven value with the easier wins.

Section V: The Risks Are Real

AI is powerful, and everything covered so far is real. But so are the risks. Adopting AI without understanding them is not innovation. It is recklessness. Here are ten risk areas every organization needs to understand.

Model Bias. AI learns from data, and data reflects human history, including our biases. If training data underrepresents certain groups, the AI will too. This has real consequences in hiring, lending, insurance underwriting, and anywhere else AI informs decisions about people.

Hallucinations. AI does not "know" things the way humans do. It predicts what comes next based on patterns. Sometimes it predicts wrong, confidently. It will cite studies that do not exist, invent statistics, and present fiction as fact. Verification is not optional.

Privacy and Data Leakage. When someone on your team pastes a confidential document into a free-tier AI product, that data may be used to train the model. It becomes part of a system that millions of people interact with. This is not a hypothetical risk. It has happened.

Security Risks. AI expands your attack surface in both directions. It can be used against you through AI-generated phishing that is indistinguishable from a real email. And your AI systems themselves can be targets through prompt injection, data poisoning, and model theft.

Lack of Guardrails. Without clear boundaries, AI can generate content that is off-brand, non-compliant, or harmful. If your teams are using AI without guidelines, you are one bad output away from a PR problem.

Governance Gaps. If nobody in your organization owns AI decisions, nobody is accountable when something goes wrong. And something will go wrong.

Shadow AI. This is the risk that keeps CISOs up at night. Your employees are already using AI tools you do not know about, on devices you do not control, with data you have not authorized. If you do not have visibility into this, you cannot manage it.

Vendor Concentration Risk. If your entire AI strategy depends on a single provider and they change their pricing, their terms, or their model quality, you are stuck. Diversification matters here just like it does in any other strategic investment.

Model Drift and Reliability. The model you tested and approved six months ago is not the same model running today. Providers update models constantly, and outputs can change, sometimes subtly, sometimes dramatically. You need ongoing monitoring, not one-time validation.

Lack of Board Oversight. Regulators, including the SEC, are increasingly clear that AI governance needs to be a board-level conversation. If your board cannot articulate how AI is being used in your organization and what controls are in place, that is a material governance gap.

Governing AI: The NIST Framework

The good news is that you do not have to build a governance framework from scratch. The National Institute of Standards and Technology (NIST) published the AI Risk Management Framework (AI RMF 1.0), which provides a proven, structured approach.[12] It is organized around four functions.

Map. Before you can manage risk, you have to know where it lives. Which teams are using AI? Which tools? What data is flowing in and out? Who is affected by AI-driven decisions? Most organizations skip this step and go straight to buying tools. That is a mistake. You cannot govern what you cannot see.

Measure. Once you have mapped where AI is in your organization, you need to quantify the risks. What is the error rate? What is the bias profile? How often does the model hallucinate? You need benchmarks, testing protocols, and metrics you track over time. Not a one-time assessment that goes in a drawer.

Manage. This is where you implement controls. Guardrails on what AI can and cannot do. Human-in-the-loop requirements for high-stakes decisions. Data classification rules that determine what can be sent to an AI and what cannot. Incident response plans for when something goes wrong.

Govern. This is the organizational layer. Policies. Roles. Accountability. Who is the AI owner? Who reviews outputs? Who reports to the board? This is not an IT function. It is a cross-functional responsibility that touches legal, compliance, risk, and the business.

The key message: governance is not something you bolt on at the end. It starts on day one, even if it is lightweight, and it evolves as your AI maturity grows.

Keeping It Safe: Enterprise AI Strategies

Understanding risk is one thing. Mitigating it is another. There are four practical approaches to keeping your data safe when using AI, and most mature organizations end up using a combination of them.

API with Enterprise Terms. Both Anthropic and OpenAI offer API access where your data is not used to train their models. You send prompts via API, get responses back, and the provider contractually agrees not to retain or train on your inputs. This is the most common enterprise path. The critical distinction is that this is different from someone on your team using the free consumer product, where data handling is much less favorable. The key is making sure you are on the right tier with the right terms of service.

Private or Self-Hosted Models. If you truly need data to never leave your environment, you can run open-weight models like Meta's Llama or Mistral on your own infrastructure or within a private cloud VPC. Nothing ever hits an external API. The tradeoff is that these models are generally less capable than frontier models from Anthropic or OpenAI, and you are taking on the infrastructure, maintenance, and GPU costs yourself.

[12]National Institute of Standards and Technology (2023). AI Risk Management Framework (AI RMF 1.0). NIST AI 100-1. https://www.nist.gov/artificial-intelligence/executive-order-safe-secure-and-trustworthy-artificial-intelligence

Cloud-Provider Hosted AI. AWS Bedrock, Azure OpenAI Service, and Google Vertex AI all let you use Claude, GPT, and Gemini within your existing cloud tenancy with data residency controls, encryption at rest and in transit, and contractual guarantees that your data stays within your account boundary. For a company already on one of these clouds, this is often the sweet spot: frontier model capability with enterprise-grade data controls.

The Hybrid Approach. This is what most sophisticated companies are doing. You classify your data into tiers (public, internal, confidential, restricted) and route accordingly. General knowledge questions and non-sensitive tasks go to a cloud-hosted API. Anything touching PII, protected health information, financials, or trade secrets either goes to a self-hosted model or gets redacted and anonymized before it hits the API. You build a gateway layer that enforces these rules automatically so individual employees do not have to make judgment calls about data sensitivity.

Regardless of which approach you choose, the enterprise checklist is the same: ensure your vendor agreement explicitly states no training on your data; have a data processing agreement that satisfies your compliance requirements (SOC 2, HIPAA if applicable); implement data loss prevention controls at the gateway level to catch sensitive data before it leaves; log all interactions for audit purposes; and document the architecture for your board governance reporting.

Section VI: The AI Onboarding Model

Everything in this paper has been building to this point. The history provides context. The terminology provides shared language. The use cases provide inspiration. The risks provide caution. Now, how do you actually do it?

The answer is a phased adoption model: Crawl, Walk, Run. Each phase builds on the previous one, and skipping ahead is how most AI initiatives fail.

This model also maps to the Prosci ADKAR change management framework.[13] ADKAR stands for Awareness (understanding why the change is needed), Desire (willingness to participate), Knowledge (understanding what to do), Ability (capability to execute), and Reinforcement (ensuring the change sticks). Successful AI adoption requires all five, in sequence.

Crawl: Foundation and Discovery

ADKAR focus: Awareness, Desire, Knowledge

Crawl is about building the foundation. It is tempting to skip this phase and jump to deploying tools, but that is asking people to execute before they understand or want to. Most failed AI initiatives fail here, not at the technology layer.

Governance comes first. Before anyone touches an AI tool, you need the rules in place. What tools are approved? What data can and cannot be shared with AI? What are the dos and don'ts? This is not bureaucracy. This is the foundation that makes everything else possible.

Enterprise-wide AI education. Not just IT, not just the data team. Everyone. What is AI? How do you use it? What should you avoid? This builds the Awareness and Desire that ADKAR requires before you ask anyone to change how they work.

Low-hanging fruit. Individual contributors start using AI for assistive tasks: knowledge lookup, drafting, summarization. At the enterprise level, you might deploy an AI-powered internal knowledge base. The goal is to start small, start safe, and start now. These early wins build confidence and demonstrate value.

Walk: Integration and Alignment

ADKAR focus: Knowledge, Ability

Walk is where AI moves from individual productivity to business impact. The shift is significant: you are no longer experimenting. You are deploying.

Consultative deployment. AI is applied to specific business problems, not broad experimentation. You are pointing the technology at defined challenges with measurable outcomes.

[13] Hiatt, J. (2006). ADKAR: A Model for Change in Business, Government and Our Community. Prosci Research.

Alignment to corporate strategy. Every AI use case ties explicitly to your organization's mission, goals, and targets. If an AI initiative does not connect to a strategic objective, it does not get prioritized. This is what separates productive AI adoption from expensive distraction.

Data pipeline architecture. AI is only as good as the data you feed it. In the Walk phase, you begin answering critical infrastructure questions: What data do we need? Where does it live? How do we structure it so AI can use it effectively? This infrastructure takes time to build, which is why you start it in Walk so it is ready for Run.

ROI measurement. Every use case needs to prove value before you scale it. This is a stage-gate approach: pilot, measure, decide. The ROI framework from Section IV becomes your prioritization tool here. It tells you what to tackle first and what can wait.

Run: AI as a Team Member

ADKAR focus: Ability, Reinforcement

This is where the model gets distinctive. In the Run phase, AI agents are not tools you use. They are team members you manage.

AI agents take on defined roles. Not vague "AI initiatives," but specific responsibilities with clear scope. "This agent handles first-pass contract review for standard vendor agreements." "This agent monitors daily sales data and flags anomalies." Clear, bounded, accountable.

SMART goals. Every employee in your organization has goals. Your AI agents should too.[14] Specific: what exactly is it doing? Measurable: how do you know if it is working? Achievable: is this within the model's capability? Relevant: does it connect to a business outcome? Time-bound: when do you evaluate? If you cannot write a SMART goal for an AI deployment, you are not ready to deploy it.

Human evaluation. A human reviews the agent's performance on a regular cadence, just like a quarterly performance review. Is it meeting its goals? Where is it falling short? What needs to change? If it is underperforming, you course-correct. If it is exceeding expectations, you expand its scope. Same process, same rigor, same accountability as any other team member.

Expanded data pipelines. As your data infrastructure matures through Walk and into Run, AI moves from task execution to business advisory. It is not just summarizing reports. It is telling you what the data means for your strategy: revenue trends, risk indicators, operational efficiency.

Consultative posture. This is the guardrail that keeps everything grounded. AI is helpful. AI is constructive. AI advises. Humans decide. That boundary does not go away in Run. If anything, it becomes more important as AI becomes more capable.

Run is not a finish line. It is the journey. As models improve, as your data matures, as your team becomes more sophisticated, the capabilities expand. You will always be in Run: continuously improving, continuously evaluating, continuously evolving.

[14]Doran, G. T. (1981). "There's a S.M.A.R.T. Way to Write Management's Goals and Objectives." Management Review, 70(11), 35-36.

The core insight is this: you do not need to invent a new way to manage AI. You already know how to set expectations, evaluate performance, and make personnel decisions. Apply that same discipline to AI, and you will be ahead of the vast majority of organizations trying to figure this out.

Conclusion

AI is not magic, and it is not a threat. It is a tool, and like every tool before it, its value depends on how you use it.

The organizations that will thrive are not the ones that adopt AI the fastest. They are the ones that adopt it the smartest: with clear governance, aligned strategy, measured outcomes, and the discipline to manage AI the way they manage any other critical business capability.

The playbook is straightforward. Understand the technology. Learn the language. Start using it personally. Identify where it creates business value. Manage the risks. Govern it properly. And adopt it in phases: Crawl, Walk, Run.

The only mistake is doing nothing.

About the Author

Rick Becker is a strategic leader with a career background as a CIO, COO, and senior executive in the insurance industry, with deep expertise spanning technology, product management, data and business intelligence, and distribution. He works at the intersection of domain expertise and emerging capabilities: designing systems, building platforms, and demonstrating what thoughtful, governed AI adoption looks like in practice.

Rick has built production AI systems solo using agentic development tools, including a fully functional insurance platform serving as a proof of concept for what a single non-technical subject matter expert can deliver without a traditional development team. He applies the same frameworks described in this paper in his own work every day.

He is based in the San Diego area and can be reached at rick@beckerz.com or at resume.beckerz.com.

Rick Becker

April 2026

resume.beckerz.com

rick@beckerz.com