IronRoot
Risk Consultants

Real Business Examples: AI Done Right (and Wrong)

What works, what fails, and what you can learn from both.

AI adoption is accelerating -- but results vary dramatically based on how it's implemented. These real-world examples (details generalized) show the difference between thoughtful adoption and costly missteps. The lessons apply across industries.

AI DONE RIGHT
Industry: Financial Services

Customer Service Team Uses AI to Draft, Not Decide

SITUATION

A mid-sized financial firm introduced an AI writing assistant to help customer service reps draft email responses to account inquiries.

WHAT WORKED:

-- AI output was always reviewed and edited by a human before sending.

-- The tool was only used for internal drafts -- never customer-facing automation.

-- Team received a 30-minute training session on what the tool can and cannot do.

-- Response time improved 40% without any compliance incidents.

KEY LESSON

Keep humans in the loop for any customer-facing or regulated communication. AI accelerates -- it doesn't replace judgment.

AI DONE WRONG
Industry: Professional Services

Vendor Contract Reviewed Exclusively by AI

SITUATION

A growing firm used a general-purpose AI chatbot to review a vendor contract and summarize key terms and risks before signing.

WHAT WENT WRONG:

-- The AI missed an auto-renewal clause that locked the company into a 3-year term.

-- A liability cap that was below industry standard went unnoticed.

-- No attorney or legal professional reviewed the summary before execution.

KEY LESSON

AI can help you read faster -- not smarter. High-stakes documents still require qualified human review.

AI DONE RIGHT
Industry: Technology / SaaS

IT Team Builds an AI-Assisted Vulnerability Triage Process

SITUATION

A SaaS company's small security team was drowning in scanner findings. They implemented an AI tool to help categorize and prioritize vulnerability data.

WHAT WORKED:

-- AI was used to cluster similar findings and suggest severity tiers -- not make final calls.

-- The team validated a sample of AI-suggested priorities each sprint.

-- Remediation time dropped significantly while analyst time shifted to high-severity items.

-- The process was documented for audit purposes from day one.

KEY LESSON

AI excels at pattern recognition in large data sets. Use it to triage -- then let experienced staff make the decisions.

AI DONE WRONG
Industry: Accounting / Tax

Employee Pastes Client Data into a Free AI Tool

SITUATION

An employee at an accounting firm pasted a client's full financial summary into a free, consumer-grade AI chatbot to help draft an analysis memo.

WHAT WENT WRONG:

-- The consumer AI tool's terms of service allowed use of inputs to improve the model.

-- Client PII and financial data were potentially retained by a third-party vendor.

-- No AI use policy existed -- the employee had no guidance to follow.

-- The firm faced a difficult disclosure conversation with the client.

KEY LESSON

Free AI tools are not enterprise tools. Without an AI Use Policy and approved tool list, your team will fill the vacuum -- with whatever's available.

AI DONE RIGHT
Industry: Healthcare

HR Team Uses AI for Job Description Drafts with Bias Review Step

SITUATION

A healthcare organization used AI to speed up drafting job postings, which previously took HR staff significant time to write from scratch.

WHAT WORKED:

-- Every AI-generated draft was reviewed by an HR professional before posting.

-- The team added a step to specifically check for gendered or exclusionary language.

-- Output quality was consistent, and time-to-post dropped by over 50%.

-- The process was included in HR policy documentation.

KEY LESSON

Structured review steps catch what AI misses. Build the governance into the workflow, not as an afterthought.

AI DONE WRONG
Industry: Retail / E-commerce

AI Chatbot Deployed Without a Fallback or Escalation Path

SITUATION

A retailer deployed an AI customer service chatbot without a clear handoff process to human agents for complex or sensitive issues.

WHAT WENT WRONG:

-- Customers with nuanced complaints received irrelevant or looping AI responses.

-- Frustrated customers escalated to social media, amplifying negative sentiment.

-- The AI was not trained on current return policy -- giving outdated answers.

-- No monitoring or feedback loop was in place to catch these failures.

KEY LESSON

Every customer-facing AI deployment needs a human escalation path, a feedback loop, and someone responsible for keeping it current.

The Difference Is Governance, Not the Tool.

IronRoot Risk Consultants helps organizations build the guardrails that make AI adoption defensible and audit-ready.

chris@ironrootrisk.com | ironrootrisk.com