ORACLE
NetSuite + azdan

BUSINESS GUIDE

NetSuite Meets AI: Practical Lessons and Strategies

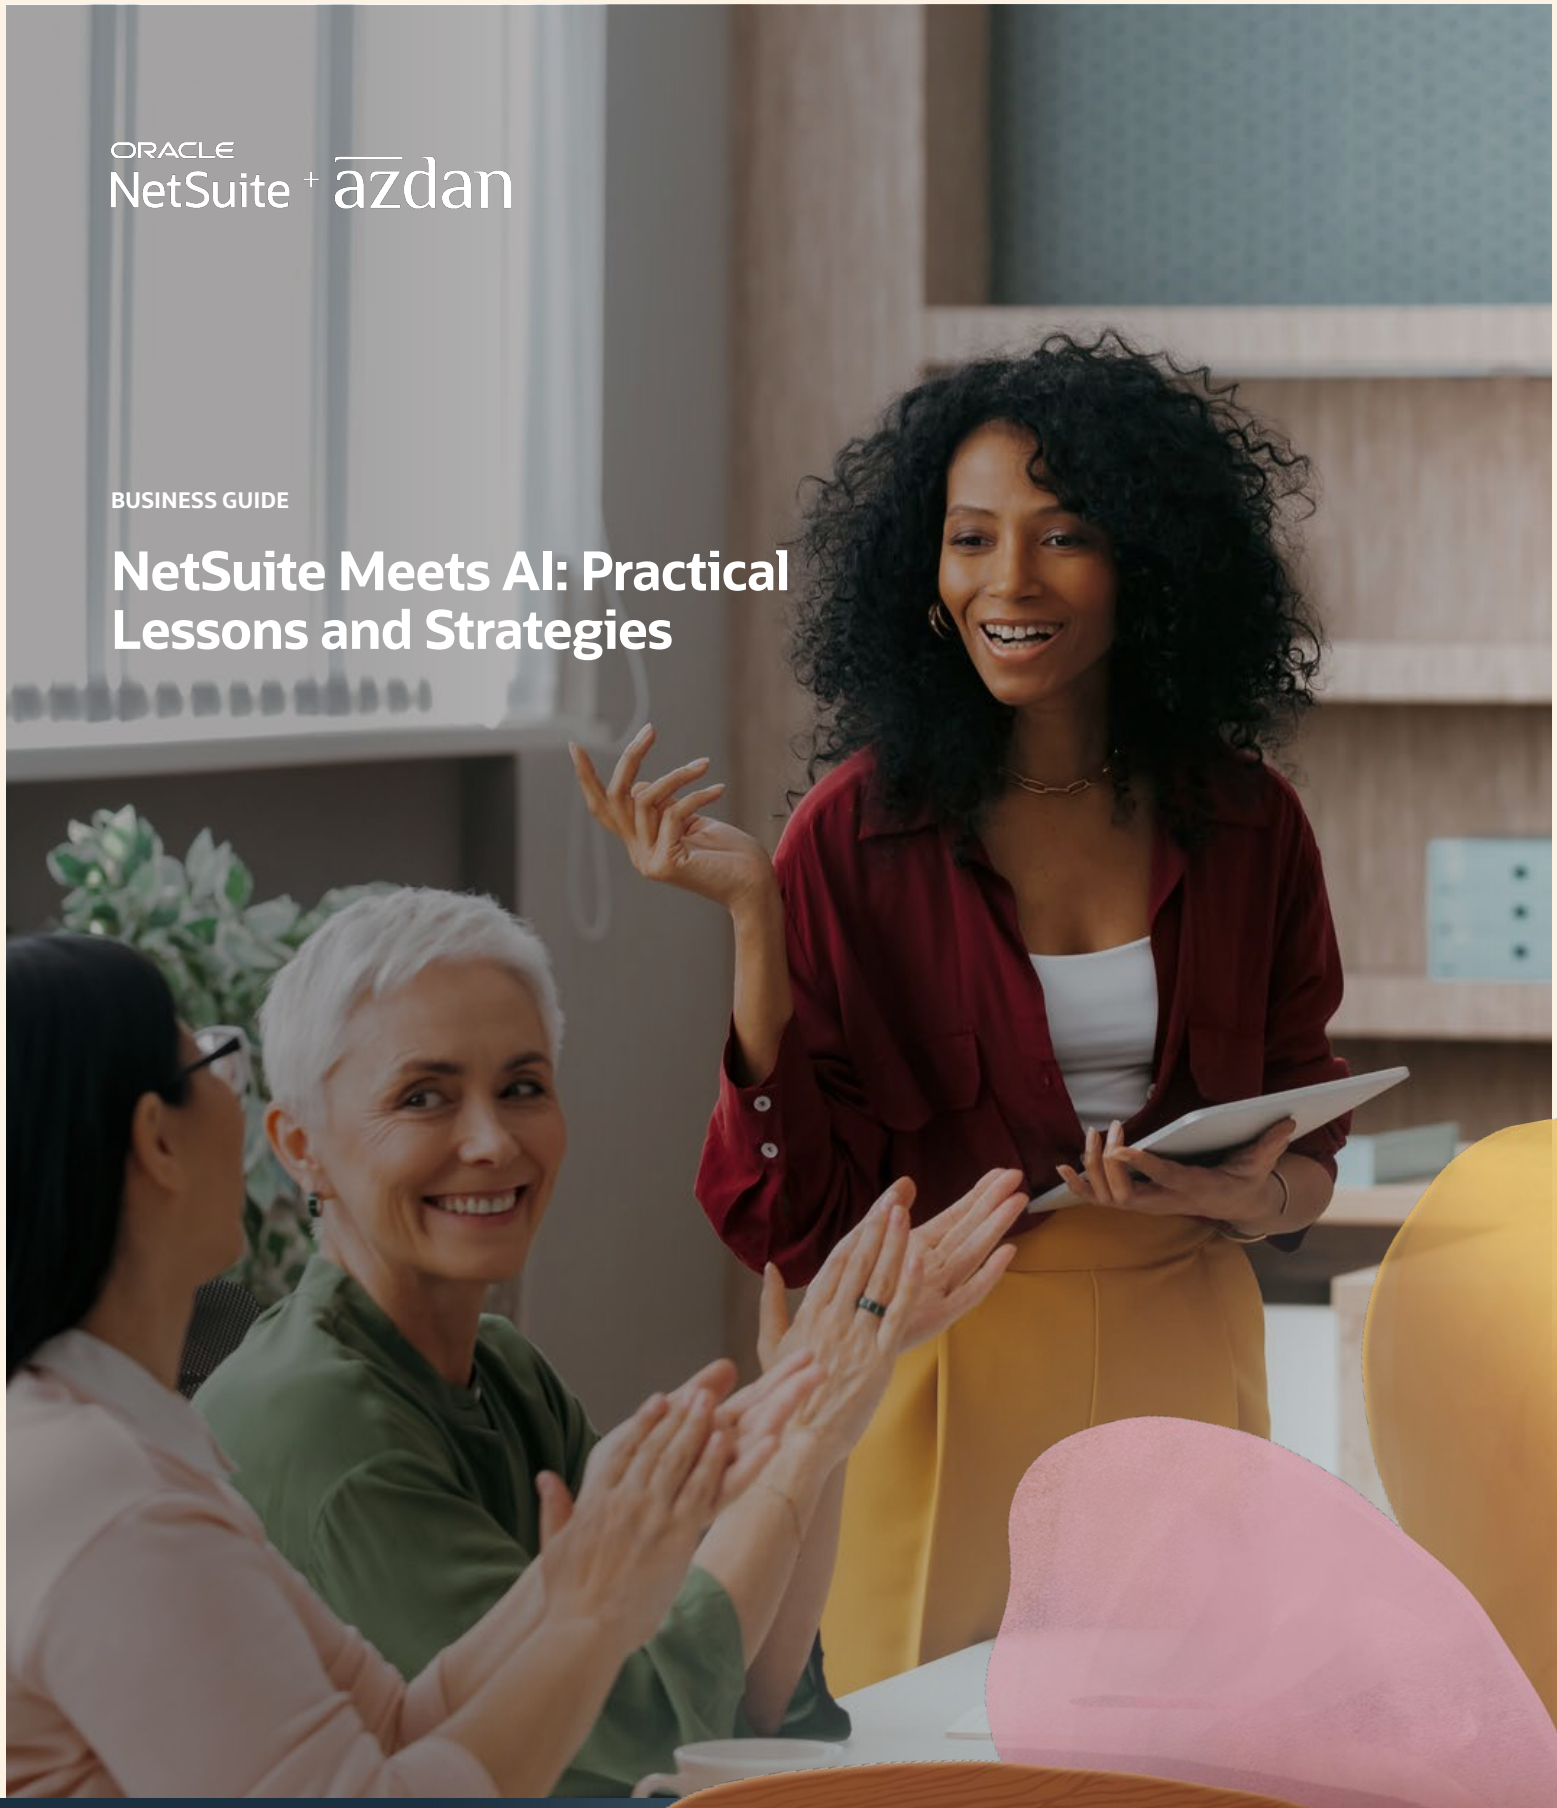

Learn more at netsuite.com

Grab a seat and enjoy.
Read Time: 16 minutes

NetSuite Meets AI: Practical Lessons and Strategies

Executive Summary

AI isn't coming to NetSuite—it's already here. The question isn't if it will reshape ERP work, but how to use it responsibly and profitably.

This guide explores what's actually delivering results today, what's being oversold, and what's ahead. We'll look at the rise of prompt-driven financial analysis, the reality of purpose-built AI agents, and the practical guardrails that keep everything trustworthy.

Because in the end, success with AI in NetSuite isn't about replacing people. It's about teaching the AI to think—the way the best analysts already do.

About the Author

Tim Dietrich is the founder of SuiteStep and a leading developer in the NetSuite ecosystem. With a software career that began in 1987, Tim brings decades of engineering experience to the ERP space. Since specializing in NetSuite in 2017, he has built a portfolio of high-leverage tools—including custom SuiteScript applications, developer-focused APIs, automation frameworks, and AI-enhanced financial analysis utilities. His work helps organizations modernize operations, streamline workflows, and fully capitalize on their NetSuite investment.

Tim Dietrich,
Founder,
SuiteStep

Table of Contents

Table of Contents

CHAPTER 1

The Current Landscape of NetSuite AI

Oracle's AI Connector has changed how AI interacts with ERP systems. Instead of exporting data or exposing internal reports to public tools, the Connector gives AI platforms such as ChatGPT and Claude a secure, governed way to access NetSuite data and reports in real time.

That capability—safe, permission-based access—turns AI from a novelty into a dependable extension of your finance team. The Connector functions like a controlled API, not a backdoor feed. Every request is logged, scoped, and governed by the same role-based permissions already enforced inside NetSuite.

But even with that secure bridge, results depend on how you ask the question.

CHAPTER 2

The Role of Prompt Engineering

Prompt engineering is the craft of giving an AI unambiguous, structured instructions that shape how it thinks. It's less about creative phrasing and more about defining the context, boundaries, and expectations the model should operate within.

In the NetSuite world, that means something very specific: structuring financial logic in a way that helps the AI interpret information the same way an accountant or analyst would. When you do it well, the AI stops guessing and starts reasoning.

I believe we're already seeing the emergence of a new role: the NetSuite AI Prompt Engineer. This is someone who understands both the data model and the dialogue model. I also believe this won't remain a niche skill for long. It'll become a requirement across both IT and finance roles.

CHAPTER 3

What's Working Now: Prompt-Driven Financial Analysis

If there's one area where NetSuite customers are already seeing results—and consistently—it's in prompt-driven financial analysis. This is where structured ERP data meets structured reasoning.

For example, consider this basic prompt:
"Show me AP problems."

Now consider this engineered prompt:
"You are a senior accountant performing a month-end AP review.

Retrieve all AP transactions from the past 30 days.

Flag any invoices that meet one or more of these conditions:

- Amount > $50K and no PO reference
- Duplicate vendor/invoice pair
- Payment terms differ from vendor master

For each finding, list the transaction ID, vendor, amount, issue, and recommended action.

Summarize findings in a risk-ordered table with an executive summary suitable for board reporting."

That's the difference between a casual chat and a controlled process. The engineered prompt is consistent, explainable, and repeatable.

Prompt-driven analysis aligns perfectly with how finance professionals think: structured, contextual, and rule-based. It brings reasoning to data that's already well-organized.

Across the projects that I've worked on, teams using prompt-driven analysis have:

- Eliminated hours of manual data gathering.
- Caught exceptions that manual reviews routinely missed.
- Produced executive-ready narratives in minutes, not hours.

I'm seeing NetSuite prompt engineering follow the same learning curve that SuiteQL did: first curiosity, then proficiency, then mastery.

CHAPTER 4

Data Quality: Garbage In/Governance Out

Most AI initiatives struggle because their data isn't ready. NetSuite's greatest strength is that its data is already ready.

NetSuite enforces relationships, balances, and auditability by design. For example:

- Every transaction ties to a subsidiary and account.
- Vendors and customers follow normalized structures.
- Approvals and scripts ensure traceability.

NetSuite doesn't just store data. It enforces discipline.

"Garbage in, garbage out" still applies, but NetSuite users start from a far stronger foundation than most.

CHAPTER 5

Trust, But Verify

In finance, trust is both earned and verified. And when it comes to AI, there's no exception. Even with high-quality data, every conclusion needs validation.

The most successful NetSuite AI workflows include verification loops and confidence scoring. Verification loops instruct the AI to reconcile totals or rerun checks before finalizing output. Confidence scores show how sure the AI is about its findings.

When used correctly, AI should present its results and its certainty. Verification doesn't have to be a separate process—it can be built directly into your prompt.

Here's an example:

"For each anomaly you report, calculate a confidence score (0–100%) based on how many independent data points support the conclusion. Include a brief explanation of why you assigned that confidence level. Then, double-check the total variance and reconcile it against the GL summary before reporting."

Using advanced techniques such as prompt-chaining, you can make the model evaluate its own reasoning and output a confidence metric automatically. That's how you move from experimental AI to enterprise-grade AI.

CHAPTER 6

What's Just Hype: The "Agent" Problem

The term "AI agent" is now being used for everything from chatbots to macros. But true agents are rare—and valuable.

A real AI agent:

- Understands ERP context and structure.
- Applies defined business rules.
- Reconciles before reporting.
- Explains its reasoning.
- Operates within governance and audit trails.

If it can't do these things, it's not an agent.

My advice is to avoid AI solutions that:

- Try to replicate NetSuite's native reports.
- Extract and store data outside of NetSuite.
- Create vendor lock-in through proprietary formats.
- Promise "fully autonomous" financial decisions with no oversight.

If your "AI agent" can't explain what it did or where it got its data, it's not an agent. I see it more as a liability.

CHAPTER 7

What's Actually Possible: Purpose-Built AI Agents

True AI agents are the next logical step toward systems that don't just analyze data, but act on it.

The most effective ones I've seen aren't generic tools. They're custom, purpose-built systems that run outside of NetSuite but operate tightly alongside it. They orchestrate processes by pulling together:

- Large language models (LLMs)
- Live NetSuite data and reports
- Business-specific rules and workflows

Think of them as intelligent, high-speed assistants that review all the relevant data and deliver the recommendations you need.

This outside-but-connected model keeps NetSuite's governance intact while enabling more sophisticated logic. It minimizes risk and preserves data integrity.

These agents work by combining:

- AI-based financial analysis (predictive reasoning and anomaly detection).
- Company-specific logic (policy thresholds, approval flows).
- Governance (reconciliation gates and logging).

For example, consider this AI agent that's been designed to assist with accounts receivable collections:

- Runs daily outside of NetSuite and obtains data via secure API calls.
- Analyzes aging and payment patterns.
- Flags high-risk accounts and suggests next actions.
- Drafts communication templates for review.
- Logs reasoning and confidence scores for audit.

That's the model for a real agent that's in production today. It's delivering real ROI with a governed, tested design. And what makes it work is the HITM ("Human in the Middle") approach. The agent doesn't act on its own, at least not yet. It evaluates the data, makes recommendations, and relies on a human to review and decide what happens next.

CHAPTER 8

The Road to Autonomy

If agents work so well with a human in the loop, what does the path forward look like? How do we get from recommendation engines to true autonomy?

This transition isn't about racing to autonomy. It's about establishing consistent, verifiable reliability first.

Organizations that are deploying agents responsibly are moving through a clear maturity path:

- AI assists – human decides
- AI recommends with confidence scores – human approves
- AI executes under review – human monitors
- Narrow autonomy on low-risk tasks – eventually

My recommendation for finance teams is simple. Stay in stages 1 or 2 for now. And then as trust, reliability, and verification strategies strengthen, gradually move forward.

This progression shouldn't be measured in speed, but in confidence. Let AI be brilliant at the financial analysis—and let your team be brilliant at the judgment, oversight, and decisions that actually move the business forward.

Here are a few things to put in place as you move your agents toward greater autonomy:

- Gather at least six months of validated accuracy metrics.
- Document and implement clear exception-handling logic.
- Establish formal audit and approval checkpoints.

Put proven governance and oversight workflows in place.

CHAPTER 9

Common Pitfalls—and How to Avoid Them

After helping multiple companies implement NetSuite AI, one pattern has become clear to me: most challenges have very little to do with the models themselves—and everything to do with strategy, structure, and mindset.

Here are some of the most common mistakes I've seen, and some advice on how to avoid them:

- Starting too big. Don't automate everything at once. Prove value with a single report or workflow first.
- Skipping governance. If you don't build governance in early, you won't add it later. Treat it as part of the foundation, not an afterthought.
- Expecting magic. AI amplifies whatever process you already have—good or bad. Clean up the workflow before you automate it.
- Ignoring the learning curve. Prompt engineering is a skill. Invest in it. Train your team through education and hands-on practice.
- Undefined success metrics. Set targets upfront: hours saved, errors caught, confidence scores, and cycle-time reductions. If you don't define success, you can't measure improvement.

The teams that I've seen succeed with AI are the ones that started small, built governance early, invested in skills, and took the time to measure what matters to them. If you get these fundamentals right, everything else becomes dramatically easier.

CHAPTER 10

Responsible Adoption and Governance

I want to take a few minutes to talk about AI adoption and governance, because it's extremely important. As I mentioned earlier, many companies either skip it entirely or treat it as an afterthought.

Governance needs to be considered from the very beginning. The teams that succeed aren't the ones with the most advanced models—they're the ones with the strongest foundation of security, transparency, and disciplined process.

Security Comes First

Strong security eliminates most risks long before AI scales. Here are a few things that I see as non-negotiable:

- Never copy NetSuite data into ChatGPT's web interface.
- Keep all API interactions inside Oracle's ecosystem, where access controls apply.
- Use role-based permissions, audit trails, and the same security posture you expect from any financial system.

Treat Prompts Like Software

Your prompts are part of your system architecture. Treat them with the same discipline you apply to code and scripts.

- Store prompts in Git, SharePoint, or a shared drive with controlled access.
- Track changes, comments, and approvals.
- Document who tested what, how it was tested, and when.

Build for Transparency

Your prompts should make the AI's reasoning visible—not hide it.

- Instruct the AI to show what data it queried, what logic or rules it applied, and its confidence score.
- Ask it to report missing information or additional data that would have improved the analysis.
- Require it to list any assumptions it made.

This turns AI into an explainable, audit-ready system rather than a black box.

Shape Your Company's Culture

Governance only works when a company's culture supports it.

- In every demo, show the prompt behind the result.
- Encourage feedback and iteration from your team.
- Celebrate transparent, repeatable outcomes—not just speed.
- Trust grows when people understand how the answer was produced.

And finally, consider this: Governance isn't bureaucracy. It's how you make AI predictable enough to trust.

CHAPTER 11

Getting Started: A Practical Roadmap

A lot of companies I've worked with fall into one of two categories: they're not sure how to begin adopting AI, or they've already started and aren't sure what to do next. The good news is that progress doesn't require a massive plan—just a deliberate sequence of small, verified wins.

Most teams underestimate how important pacing is. Success with AI comes from stacking those small, validated steps—not big leaps. Here's a streamlined roadmap that gets you moving without overcomplicating the process.

Phase 1: Experiment Safely

Start with one targeted financial analysis prompt.

Example:

"You are a financial analyst reviewing this month's income statement. Identify unusual expense trends compared to the prior three months and summarize the findings."

Goal: Produce a clear, explainable result.

Success Metric: The AI surfaces at least one pattern the team can verify manually.

Phase 2: Operationalize the Basics

Once the first prompt is reliable, expand it slightly and make it repeatable.

Key Steps:

- Test across a few different reporting periods
- Refine prompts for edge cases
- Add basic verification steps
- Document what works
- Review results with stakeholders

Phase 3: Extend and Automate Carefully

Build an AI agent based on what's proven. Don't automate anything untested.

Key Steps:

- Identify one or two repeatable workflows
- Define the data flow and review checkpoints
- Prototype a small orchestration
- Add human-approval steps
- Document prompts and logic
- Assess readiness for limited, controlled automation

My general advice is this: Start small. Move deliberately. Build trust one workflow at a time.

CHAPTER 12

What's Next?

Here's what I see coming.

The next phase of NetSuite AI won't be driven by hype or headlines—it will be driven by discipline.

We're early enough to set standards, but late enough to have proven patterns.

Expect to see:

- Deeper prompt integration into native NetSuite workflows.
- Reusable agent frameworks for finance and operations.
- New roles—NetSuite AI Prompt Engineer, AI Analyst, and Agent Architect—that combine aspects of finance and IT with AI fluency.

What I've learned working with finance teams across industries is that AI succeeds not through innovation alone, but through discipline.

NetSuite users already understand the value of process—and that's exactly what makes them poised to lead this next wave.

Conclusion

NetSuite AI is already delivering value—quietly, practically, and responsibly. The biggest wins today come from prompt-driven analysis built on clean, governed data.

The next wave—purpose-built agents—will extend that value through structured autonomy and explainable intelligence.

We're still early—and that's the best place to be.

The prompts and agents we're building now will set the standard for what comes next.

Start small, move deliberately, and scale what works.

The hype will fade, but the practical wins will compound.

If you want to implement AI in a way that produces results you can trust, build systems that earn that trust. Do it one prompt, one agent, and one workflow at a time.

ORACLE
NetSuite

www.netsuite.com
Infonetsuite_WW@oracle.com
877-638-7848

facebook.com/netsuite

twitter.com/netsuite

linkedin.com/company/netsuite

instagram.com/oraclenetsuite