Re-Engineering Your Business Using AI

Strategy, agentic transformation, and systems thinking

By Drake Wolfe



The Three Promises

Everything in this session rolls up to these core outcomes. By the end, you'll have practical frameworks you can implement starting Monday.

Practical AI Agent Strategy

Learn how to develop a clear, actionable AI agent strategy for your organization that delivers measurable results in 90 days.

Agentic Transformation

Understand how agentic AI fundamentally transforms processes and team structures, creating new roles and workflows.

Systems Thinking

Master techniques for building scalable, autonomous AI ecosystems that compound value over time through reuse and connection.

The Story: From Tools to Teamwork

AI is less about robots and more about how work flows through your organization. The evolution is simple but powerful:

Tools help individuals complete tasks faster.

Agents coordinate workflows and make decisions within guardrails.

Ecosystems learn from every interaction and continuously improve.

This progression changes everything about how teams operate, where handoffs happen, and what humans focus on day to day.





Part 1

Build a Practical AI Agent Strategy

Start With the Business Goal

Ask the Right Question

What outcome would a board member care about in 90 days? This isn't about technology. It's about business impact that shows up in your numbers.

Pick **one outcome per department**. Avoid boiling the ocean. Focus creates momentum.

Sales

Faster quote turnaround from days to hours

Support

Fewer escalations and faster resolution times

Finance

Shorter month-end close cycle

Process Inventory in 10 Questions

Answer these questions for any workflow you're considering. The answers will reveal where agents can deliver the most value.

- 1. What are the top recurring workflows?
- 2. Where are the handoffs between people or systems?
- 3. Where do delays and rework happen most often?
- 4. What rules or policies govern each step?
- 5. What inputs are missing or inconsistent?

- 1. What is the current cycle time and cost?
- 2. What user pain shows up in tickets or emails?
- 3. Which steps are document heavy?
- 4. Which steps are decision heavy but rule based?
- 5. What would "good" look like in one sentence?

The 1-Page Agent Strategy Canvas

This canvas is your north star for the pilot. Fill it out in one sitting with your team. Keep it simple and specific.

Target Workflow The specific process you'll transform first		
		04
		Inputs & Fields
What data the agent needs to do its work		
What data the agent needs to do its work		
06		
06		
06 Human Review Points		
Human Review Points Where humans stay in the loop		

Capability Menu

Choose 2 or 3 capabilities maximum for your first pilot. Starting small lets you learn fast and build confidence before expanding.



Summarize and Draft

Create initial versions of documents, emails, or reports



Classify and Route

Organize incoming requests and send them to the right team



Look Up and Validate

Check information against policies or databases



Schedule and Notify

Coordinate timing and send reminders



Generate Forms

Create checklists and structured templates



Create Quotes

Build proposals and financial summaries

Roles: Who Owns What

Clear ownership prevents drift and scope creep. Assign these roles before you start building.

Executive Sponsor

Sets the business goal and removes roadblocks when the team gets stuck

Process Owner

Defines what "done" looks like and accepts the final output

Builder

Glues tools together and configures the agent capabilities

QA Lead

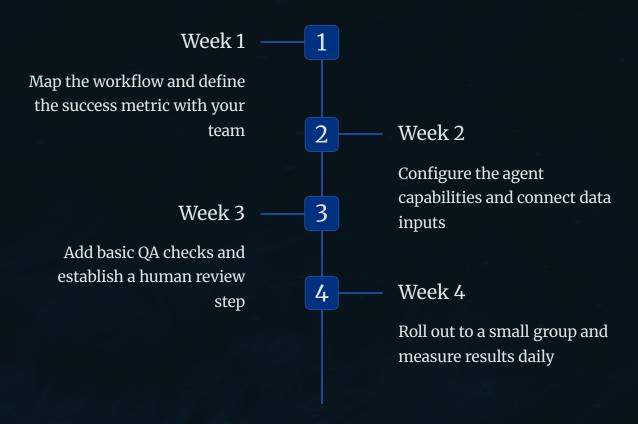
Defines acceptance checks and monitors quality metrics

Security Lead

Signs off on data access and ensures compliance requirements are met

Pilot Plan: 30 Days





Deliverables: Intake form, output template, success dashboard, rollback plan

Measure What Matters

Promise less. Deliver more. Publish the wins. Track metrics that your stakeholders already care about.







Speed

Cycle time reduction, time to first response, queue wait times

Quality

Error rate, rework frequency, customer satisfaction scores

Cost

Time saved per transaction, software usage reduction, payback month

Evidence: Independent Productivity Studies

Anchor your story with specific, credible results from peer-reviewed research and controlled experiments.

Customer Support

"Generative AI at Work" (Brynjolfsson, Li, Raymond)

Field experiment at Fortune 500 software firm: issues resolved per hour +14% average, +34% for novice agents, with better customer sentiment and retention.

NBER Working Paper w31161, 2023; QJE, 2025

Professional Writing

"Experimental Evidence on the Productivity Effects of Generative AI" (Noy & Zhang)

Mid-level professionals completed writing tasks \sim 40% faster with \sim 18% higher quality.

Science, 2023

Software Development

"The Impact of AI on Developer Productivity" (Peng et al.)

Controlled experiment with GitHub Copilot: developers finished a JS task **55.8% faster**.

arXiv:2302.06590, 2023

Consulting & Knowledge Work

"Navigating the Jagged Technological Frontier" (Dell'Acqua et al.)

BCG consultant field study: +12.2% more tasks, 25.1% faster, >40% higher quality on benchmarked work.

SSRN/Harvard, 2023; BCG 2024

Tie each stat to your pilot metric: speed, quality, or cost.



Evidence: Enterprise Examples

Klarna — AI Customer Service Assistant

Scale	Impact	Quality
2.3M conversations in first month	Equivalent work of ~700 FTEs	On-par CSAT with human agents
Handled two-thirds of all chats across 23 markets, 24/7, in 35+ languages	25% drop in repeat inquiries Customer errands resolved in <2 minutes vs 11 minutes previously	Sources: Klarna press release (Feb 27, 2024), OpenAI case study, Reuters/Business Insider coverage

Use as a directional benchmark for scale and metrics to track in your own pilots.

Illustrative Case Studies

Examples you can narrate as "if this were your workflow," without naming clients. These patterns show what's possible.

1

Construction Estimating

Build: Simple quote intake \rightarrow policy-grounded draft \rightarrow QA checklist \rightarrow scheduler suggestion

Outcome pattern: Turnaround from 3 days to same-day; win rate +9 points; rework -22%

2

Agency Orchestrator

Build: Orchestrator routes briefs \rightarrow drafts \rightarrow QA \rightarrow approval \rightarrow publish, with logs

Outcome pattern: Cycle time **-70%** on routine content; brand consistency up; fewer handoffs

3

Support Knowledge Assistant

Build: RAG answers with citations; human-in-the-loop for edge cases

Outcome pattern: First-contact resolution +12%; handle time -18%; CSAT up

E-commerce Listing Flow

Build: Tagging, descriptions in brand voice, alt text, marketplace formatting

Outcome pattern: Time-to-publish -60-70%; organic clicks +10-20%; clearer sizing reduces returns

Clearly label these as illustrative when presenting. Swap numbers if you prefer different targets.

Part 2

How Agentic AI Transforms Processes and Teams

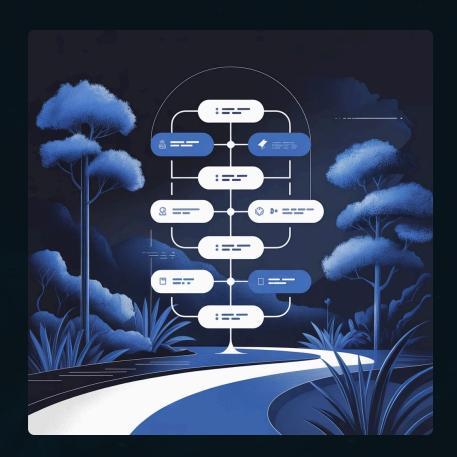


What "Agentic" Means in Plain English

An agent is software that can **decide the next step toward a goal**. It acts inside guardrails and asks a human when the situation is unclear.

Think of it as a teammate, not a toy. It doesn't just respond to commands—it makes decisions, checks its work, and hands off to humans at the right moments.

Guardrails keep it safe. Logs make it transparent. Feedback makes it better.





Before and After: A Generic Flow

1

Before

Manual intake → scattered docs → long back-and-forth → late handoff

Every step requires context switching, emails get lost, and nothing is tracked systematically.

2

After

Simple intake → agent drafts and checks → human approves → system logs

Clear handoffs, consistent outputs, automatic documentation, and metrics on every transaction.

Patterns by Department

Use this menu to pick one capability and one metric per team. Start where pain is highest.

Department	Typical Capability	Outcome Metric
Sales	Lead triage, discovery notes, proposal drafts	Response time, win rate
Marketing	Briefs, content variants, asset tagging	Time to publish, CTR
Operations	Scheduling, capacity checks, vendor comms	Cycle time, on-time rate
Finance	Invoice drafts, variance summaries	Close time, accuracy
HR	Interview kits, onboarding checklists	Time to hire, ramp time
Support	Knowledge-grounded replies	First contact resolution, CSAT

Team Structure With Agents

Humans

Set goals, handle exceptions, build relationships, make judgment calls on edge cases



Agents

Draft outputs, check quality, route requests, send reminders, log every transaction

Rules

Simple acceptance criteria for each output, clear escalation paths, documented policies

Roles get clearer. People do higher-value work. The system becomes more transparent.

Change Management in Two Moves



1 Champions

One person per team to own adoption, answer questions, and share wins

2 Micro Training

5-minute playbooks inside the tools—no long workshops or dense manuals

Announce the pilot like a product launch. Make it feel exciting, not mandatory.

The 30/60/90 Transformation Roadmap

Small wins that add up. Each phase builds on the last.

30 Days

One pilot working and measured with real users and real metrics

60 Days

Extend to an adjacent process in the same department or a different team

90 Days

Connect pilots, add dashboards, write the playbook for the next wave



Part 3

Systems Thinking for Scalable, Autonomous Ecosystems

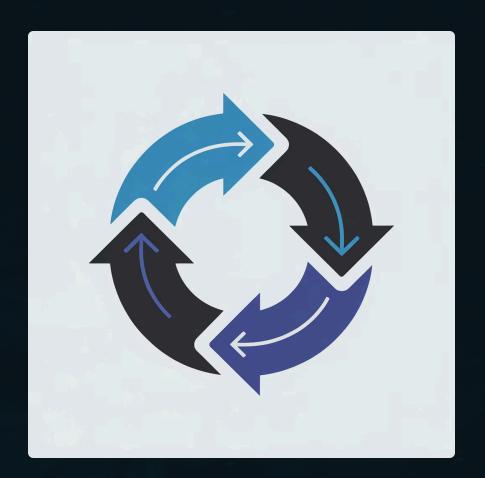
Think in Loops, Not Lines

The Simple Loop

Input → Action → Feedback → Improve

Each output teaches the next run. Bad outcomes trigger review. Good outcomes get reinforced.

This is how systems get smarter over time without constant manual tuning.



Guardrails - Safety First

Safety is designed in from day one. These four elements keep agents working inside boundaries.

Policies

What the agent is allowed to do—the scope of its decision-making authority

Checks

A short list the agent must pass before marking completion—quality gates

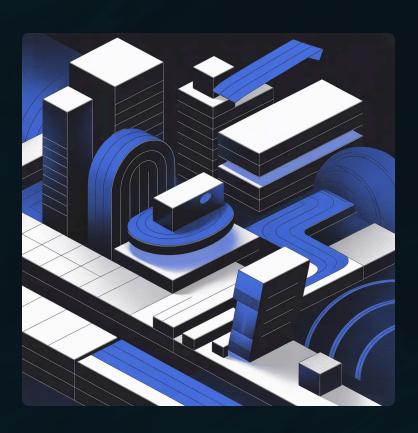
Escalation

When to ask a human—clear triggers for handing off uncertain situations

Logs

What happened and why—full audit trails for every decision and action

The Capability Bank



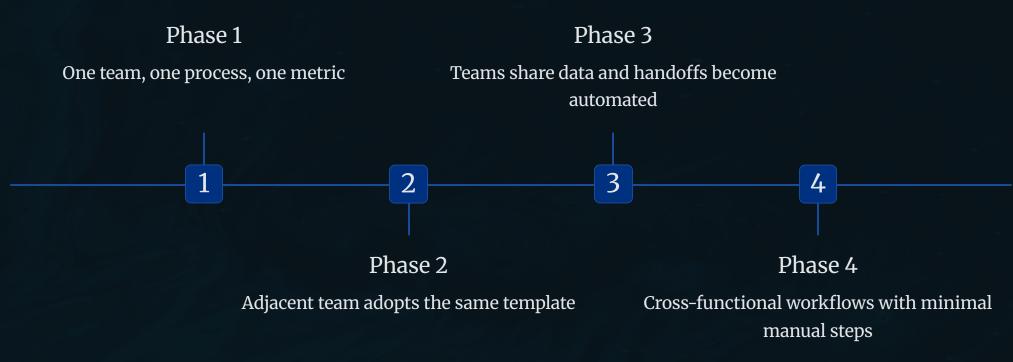
Every successful pilot becomes a reusable block: intake forms, checklists, templates, prompts, review rules.

Stack these blocks to speed up the next build. Your second pilot takes half the time. Your fifth takes a quarter.

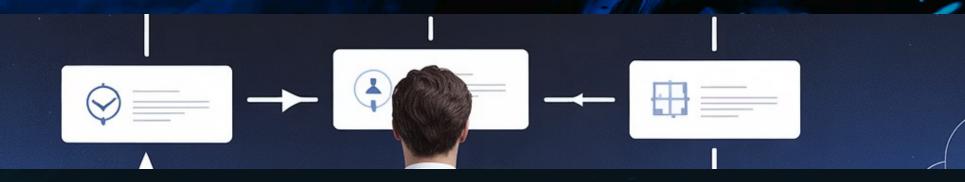
Reuse beats rebuild. Document what works and make it easy for other teams to copy.

Connect the Dots Over Time

Start with one team. Then connect teams around shared data and shared templates. This is how an ecosystem emerges without a big-bang project.



Crawl, walk, run. Keep the north star visible.



Quick Procurement Guide

Buy for speed. Keep your options open. Ask these questions before committing to any vendor.

1

Who owns the data and outputs?

Make sure you retain ownership and can use outputs freely

2

How do we export or migrate later?

Avoid vendor lock-in with clear exit paths and data portability

3

What access controls exist?

Ensure role-based permissions and security compliance

4

How is usage priced?

Understand cost per transaction, user, or volume to avoid surprises

5

What metrics and logs can we see?

Visibility into performance, errors, and audit trails is essential

Resource Map to Find Tools

Start with the capability you need, shortlist three options, trial with your pilot metric for one week.



Discovery

Product Hunt, There's An AI For That, G2, Capterra, StackShare browse by category and read real user reviews



Open Source & Demos

GitHub Awesome lists, Hugging Face Spaces—test capabilities before buying



Learning & Research

Papers With Code, vendor blogs with how-to guides—understand what's proven to work

Three Key Takeaways



Start with a Pilot

Implement agentic AI with a clear business goal and measurable outcomes. Focus on small, impactful wins to demonstrate value quickly.



Empower Agentic Workflows

Move beyond simple automation. Deploy AI agents that make decisions, check their own work, and collaborate effectively with human teams.



Scale with Loops & Guardrails

Design for continuous improvement using feedback loops and robust guardrails to ensure safe, transparent, and evolving autonomous systems.

Thank You

Thank you for attending our presentation on agentic AI. We hope you found the insights valuable and inspiring for your business transformation journey.

For further questions or discussions, please reach out to us:

Email: hello@goagentix.com

Connect with me on LinkedIn for more insights on AI solutions and business advice.

My debut #1 Best Selling book, The Final Business Model, breaks all of this down into great detail over 400 pages. You can find it on **Amazon**.

