



Managed IT & Cybersecurity for healthcare and mission-critical businesses

20+

Years in Business

97%

Customer Satisfaction

2x

Inc 5000 Honoree

WHAT WE DO



Cybersecurity & Compliance



Cloud & Infrastructure



Helpdesk & Onsite Support



Virtual CIO & Strategy



AI Integration & Automation



Data Backup & Disaster Recovery



100% Local Chicagoland Team



90-Day Satisfaction Guarantee

SAFEPOINT IT WEBINAR · JUNE 11

What Comes After Prompting: Agentic AI for Business Leaders

Most teams are still asking AI questions. The next wave is AI that takes action - and the leaders who understand it now will have a serious edge.

Presented by:

Dmitry Rudman | CTO & AI Automation Expert

Alek Pirkhalo | CEO & Technology Strategist

Kaden Evenson | Service Manager

Jared Bahr | Enterprise Engineer



PRACTICAL TAKEAWAYS

FOR BUSINESS LEADERS

What We Will Cover

The Rise of Multi-Agent AI Workflows

How Claude Fable 5 and OpenAI Codex are shaping the next phase of autonomous work.



The Risks Behind Agentic AI

Our team breaks down the ethical, security, and oversight concerns every organization should understand.



Real ROI, Right Now - Demos

Two different workflows where agentic AI is already paying off.

Final Q&A: The SafePoint Path Forward

Bring your questions about AI, automation, and security.

Claude Fable 5, Mythos 5, and the Codex Integration: The Next Step in Agentic AI Development

What Anthropic Announced

Claude Fable 5: Anthropic's first widely available Mythos-class model, built for advanced coding, research, vision, and long-running tasks.

Claude Mythos 5: Uses the same core model as Fable 5, but access is limited to trusted partners.

Why It Matters

More autonomous work: Fable 5 is built for longer, more complex tasks like coding, migrations, research, analysis, and vision work.

Built-in safeguards: Higher-risk prompts may route to Claude Opus 4.8, but Anthropic says most sessions stay within Fable 5.

Codex Integration

Two strengths: Codex is built for task execution and automation, while Claude Fable 5 focuses on deep reasoning, analysis, and knowledge work.

Better together: Teams can use Codex to act across tools and Claude to support research, planning, and decision-making.

The AI development workflow is shifting from single-model assistance to agentic, multi-model collaboration. Claude Fable 5 brings stronger autonomous reasoning, while Codex integration adds another layer of review and task delegation inside Claude Code.

Generative vs. Agentic

Both are AI. But they operate in fundamentally different ways - and confusing them leads to misplaced expectations and poor adoption decisions.

	Generative AI	Agentic AI
What it does	Creates content, answers questions	Takes actions across systems
You ask it to	Write, summarize, explain	Send, update, route, execute
Touches	A chat window	Email, tickets, calendars, workflows
Risk profile	Low - you review output	Higher - it acts on your behalf
Best starting point	Day one	After controls exist

❏ Most teams should be using both - generative AI to assist, agentic AI to automate. The key is knowing which is right for which task.

The Leverage vs. the Risk

When AI handles routine work like summaries, drafts, follow-ups, and data entry, small teams can spend more time on the work that actually needs human judgment.

Where the output comes from:

- Reports that build and send themselves
- Tickets sorted and routed the moment they arrive
- Follow-ups that go out without anyone remembering
- Hours every week reclaimed for the work that actually needs a human

Where the risk comes from:

- More tasks running means more places something can quietly go wrong
- Faster work means less time to catch a mistake before it spreads

Bottom line: more actions = more output, and more actions = more surface area for mistakes. The guardrails matter as much as the tools.



COMMON PITFALLS

The Mistakes People Are Making Right Now

Agentic AI adoption is moving fast - and most teams are making the same four mistakes. Knowing them in advance is half the battle.



No oversight — no monitoring, no way to catch problems, no way to undo. Build the guardrails before you build the speed.



Messy data in — bad data in means bad results out, just faster. Clean your inputs first.



Do it right from the start — "let's try AI somewhere" isn't a plan. Pick one task, one owner, one way to measure success.



No human in the loop — AI is a junior employee, not a senior one. Someone has to check the work, especially early on.

Customer Service & Daily Requests

Every business gets the same questions, requests, and follow-ups over and over. AI can handle the repetitive parts so your team can focus on the work that actually needs a human.

Where it shows up:

- Sorts and prioritizes incoming requests so nothing slips through
- Pulls the key points out of long email chains and conversations
- Drafts first-draft replies your team can review, edit, and send
- Turns every solved problem into a reference your team can reuse next time
- Spots patterns over time so leadership sees what keeps coming up

Why it works: Lots of similar tasks, clear rules, easy to check the work.

i Why it works: high volume, repetitive structure, clear success criteria.



The Hidden Risk: Prompt Injection

Agentic AI is dangerous when it can be influenced by the wrong input. Unlike a chatbot that just gives you a bad answer, an agentic system can take a bad action.

- Hidden instructions get planted in an email, document, or webpage - invisible to you, readable by AI.
- The agent reads them and runs them - forwarding data, changing settings, triggering workflows no one approved.
- Every major AI platform has been shown to be vulnerable. This isn't theoretical.

How to stay safe:

- Give every agent the smallest amount of access it needs
- Never connect it to systems it doesn't actually need
- Require human approval for anything that sends, modifies, or deletes data



Operations & Reporting

This demo shows how AI can replace manual spreadsheet reporting with a cleaner, easier-to-read HTML report. Instead of updating rows, formulas, and formatting by hand, you can make changes through a simple chat interface.

Where it shows up:

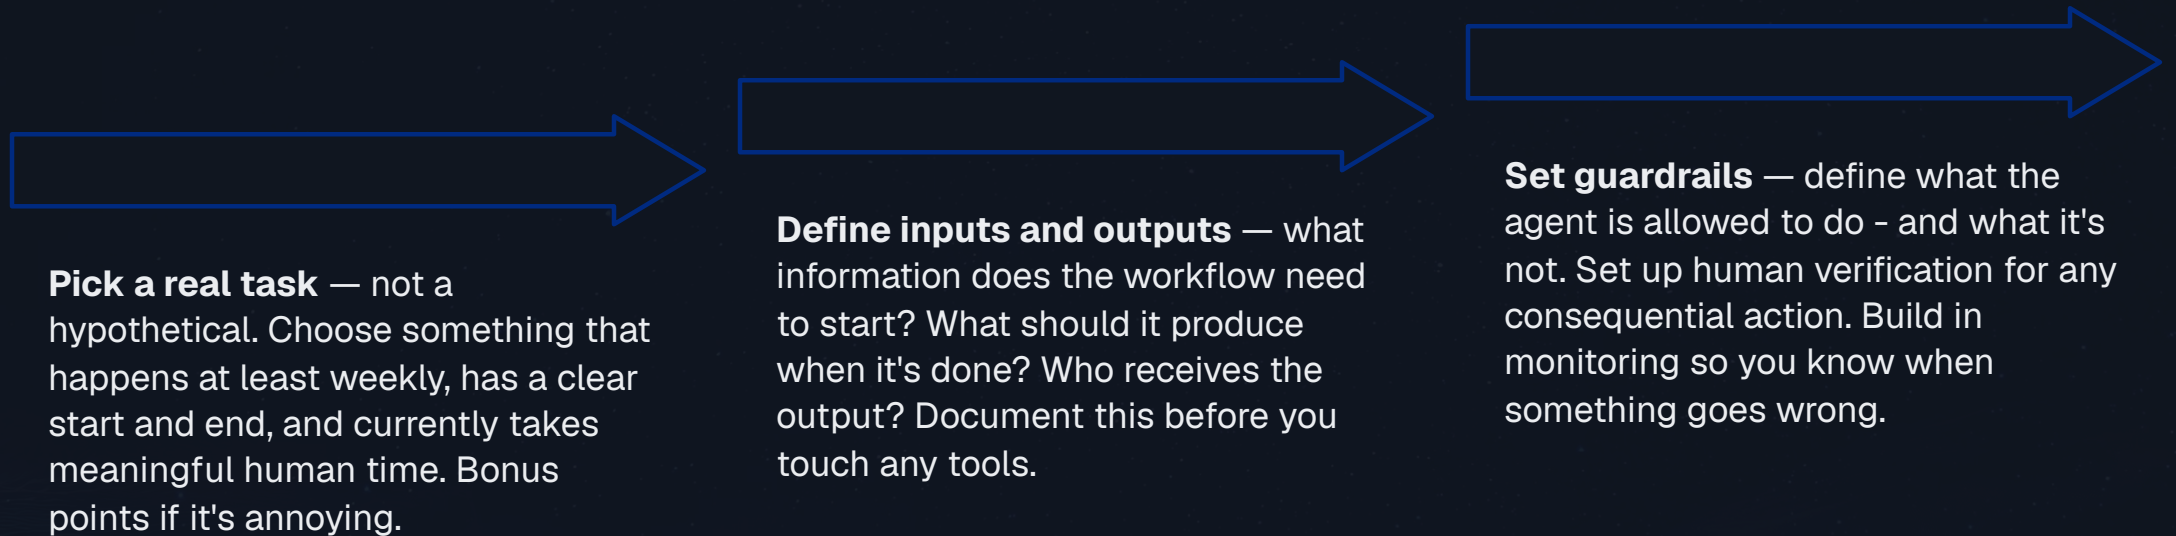
- Weekly and monthly reports generated in a polished format, without rebuilding the same spreadsheet every time
- Timesheet-style updates pulled into a report, instead of manually filling out every field
- Cleaner dashboards that are easier for leadership to read than a traditional Excel sheet
- Report edits made by asking Claude for changes, instead of adjusting cells, formulas, and layouts by hand
- Flexible report designs that can be updated for internal reviews, client updates, or leadership summaries

Why it works: Reporting follows a repeatable structure. Once AI knows the format, you can update, edit, and improve the report through chat instead of wrestling with the spreadsheet.



The Framework

The biggest mistake teams make is trying to automate everything at once. Start with one workflow. Do it right. Then expand. Here's how.



📌 Treat AI like a **junior colleague**: capable, fast, occasionally wrong, **always reviewed**.

Thank you for joining us.

We appreciate you taking time out of your day to join us and explore practical ways AI can support real business work, decision-making, and insight.

Don't forget to register for our next webinar in July -
The link is in your inbox along with your post webinar goodies.

LIVE WEBINAR • JULY 23, 2026

Using Claude as Your Business Agent

How small teams can use Claude to research, follow up, and take action across business tools.

DATE July 23, 2026 | **TIME** 3:00 PM CDT | **DURATION** 60 minutes

SAFEPOINT IT

The graphic features a central blue circle with the text 'AI' inside. Four lines radiate from this center to four circular icons: a magnifying glass for 'RESEARCH', a speech bubble for 'COMMUNICATION', a lightning bolt for 'ACTION & AUTOMATION', and a grid of squares for 'APPS & TOOLS'.

- We will be showing how Claude can work as a business AI agent across tools like CRM, email, billing, and research. We'll focus on practical use cases small teams can actually understand, control, and use safely.

Any Questions?

Bring them here. Ask about AI, automation, security, or where this could actually fit inside your business.

Still not sure where to start?

Schedule a practical AI conversation. We'll help you find the right opportunities, avoid the wrong ones, and move forward safely.

