BUILT TO BE TRUSTED

A Disciplined Approach to Agentic AI for HR

A WHITE PAPER ON AGENTIC AI METHODOLOGY

“Just got back from a conference where they talked a lot about Agentic AI. Seems like a no-brainer, and you could probably automate everything in a few weeks.”

-Senior leader, in conversation with the CHRO

Authored by:

- Gus Prestera, PhD MBA
- Meltem Bilgic
- Belen Bilgic-Schneider

June 2026

If only it were that easy. Agentic AI is real, it is powerful, and it is starting to reshape how HR and L&D functions deliver value. But the version sold from conference stages, the one that promises to “automate everything in a few weeks,” is an empty promise. Real, production-ready agentic AI takes deliberate design, disciplined testing, and humans who stay in the loop.

This white paper describes how the Prestera FX team approaches Agentic AI projects, walking through our five-phase methodology and the design principles that help our clients build workflows they can actually trust. Whether you are exploring your first pilot or scaling up your AI strategy, our hope is that this paper helps you set realistic expectations, ask better questions, and build something worth the investment.

Agentic AI Overview

Let's level set with a brief overview of what Agentic AI is, what it looks like, and why it matters.

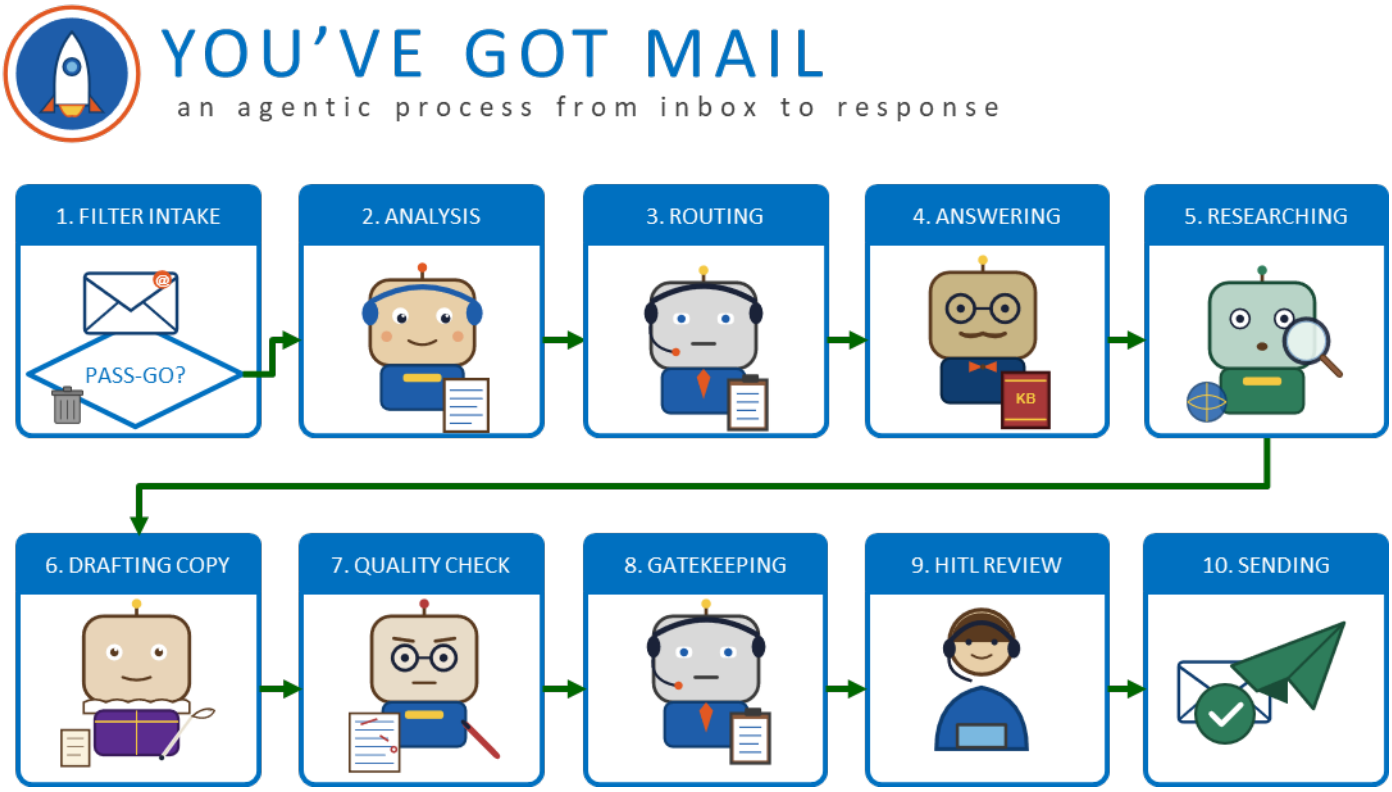

WHAT IS AGENTIC AI?

Agentic AI refers to autonomous systems that can be used to automate a process with the use of AI Agents that can execute tasks and make decisions without direct supervision.

Unlike Generative AI chatbots that you may be used to—such as ChatGPT—Agentic AI does not need to be prompted. It operates in the background and is usually triggered by an event that it monitors for on a constant basis.

WHAT'S IT LOOK LIKE?

Just about any routine process can be automated with Agentic AI. A simple example that we'll reference throughout this paper is an inbox management process. A customer sends an email to your company's general inbox, a series of agents analyze that message, interpret it, research how best to respond, compose a reply email, and evaluate the message. As you can see from the diagram above, the workflow can involve many specialized AI Agents as well as Human-in-the-Loop reviewers who make sure no errors slip through to the customer.

WHY DOES IT MATTER?

The year 2026 has been dubbed the Year of Agentic AI. This year, AI tools crossed over from being just *AI Assistants* helping individuals write emails, summarize meetings, and otherwise improve their personal productivity…to performing substantial work that requires creativity and judgment as *AI Agents,* working inside of automated workflows as part of *Agent Teams* performing complex tasks. Agentic AI helps organizations not only do things better, faster, and at larger scale, but also tackle challenges they never could before.

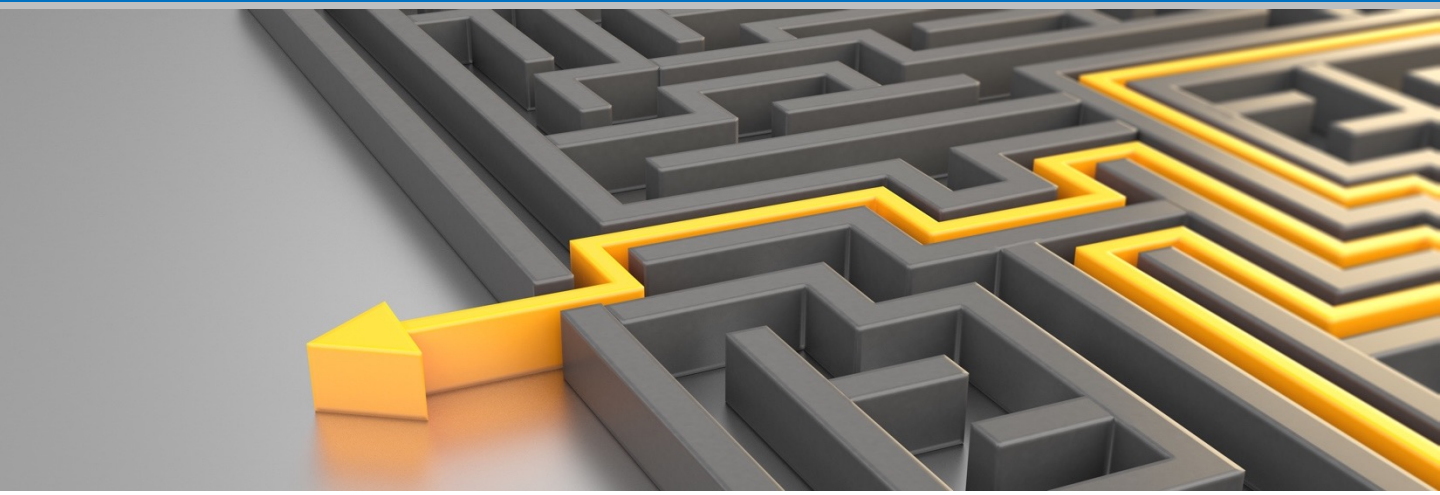

THE CHALLENGES

Three forces are converging in HR and L&D right now, and Agentic AI sits at the intersection of all three. Before you commit to a project, understanding the landscape matters, because the gap between what is possible and what is sustainable is wider than most senior leaders realize. The opportunity is real, but the failure rate is high. This paper is intended to help address this *discipline* gap between the 5% that succeed and the 95% that do not.

1 Productivity Pressure

HR and L&D teams are being asked to deliver more with less. Gartner projects HR budget growth will slow from 2.4% in 2025 to 0.7% in 2026, with only 29% of CFOs planning to increase HR budgets and 22% planning cuts (Gartner, 2026). Headcount growth expectations are collapsing from 6% in 2025 to just 2% in 2026, with Gartner describing this as "a structural pivot from labor expansion to optimization driven by automation and AI" (Gartner, 2026). The math is unforgiving: leaner teams, expanding mandates, and a hard ceiling on hours in the day.

2 Manual Process Drag

The processes most ripe for transformation are still mostly manual. Deloitte's research finds that HR staff spend up to 57% of their time on administrative tasks (Deloitte, 2024), and nearly one in five HR teams are still doing work manually that software could already be doing, even without AI (HiBob, 2026). Talent reviews, 360 administration, IDP creation, onboarding intake, and exit interview analysis all share the same pattern: high-volume, repeatable, and slow. These are also the processes where well-designed agentic workflows can deliver measurable value.

3 AI Hype vs. Practical Implementation

The third challenge is the noisiest one. Gartner predicts that over 40% of agentic AI projects will be canceled by the end of 2027 due to escalating costs, unclear business value, or inadequate risk controls (Gartner, 2025). IDC research, conducted with Lenovo, shows that 88% of AI proofs of concept never reach widescale deployment: for every 33 AI POCs launched, only four graduate to production (Schuman, 2025). MIT's Project NANDA found that just 5% of generative AI pilots achieve rapid revenue acceleration, while 95% see little to no measurable return on the bottom line (Estrada, 2025). Gartner has also called out "agent washing," where vendors rebrand existing chatbots and automation tools as agentic AI. Of the thousands of vendors making agentic claims, Gartner estimates only around 130 deliver genuine agentic capabilities (Gartner, 2025).

Where Agentic AI Fits In HR

Where could agentic AI deliver value for HR? Almost everywhere, but not all use cases are created equal. The six categories below show the lifecycle areas where teams are seeing meaningful results today.

TALENT ACQUISITION

- Resume screening
- Candidate screening and ranking
- Interview scheduling
- Reference checking
- Offer letter drafting

ONBOARDING

- Pre-boarding workflows ★
- New hire orientation Q&A
- 30/60/90 check-in drafting
- Compliance training tracking

PERFORMANCE & CAREER

- Performance appraisal drafting
- Individual Development Plans ★
- Coaching prep notes
- Career pathing recommendations

LEARNING & DEVELOPMENT

- Content curation and tagging
- Course recommendations
- Skills assessment scoring
- Training program administration

TALENT MANAGEMENT

- Talent review summaries
- Succession planning analysis
- HiPo identification
- 360 administration ★

HR OPERATIONS

- Policy Q&A and employee inquiries ★
- Benefits inquiries
- Offboarding workflows
- Exit interview analysis ★

The items marked with a star ★ are good ***first projects*** because they share four traits: clear inputs and outputs, repeatable patterns, defined success criteria, and relatively low downside if something goes wrong. Start there to build internal credibility, then take on higher-stakes processes once your team has earned the trust to do so.

OUR METHODOLOGY

Our methodology breaks an agentic AI project into five phases.

Each phase builds on the one before, and each has its own deliverables, success criteria, and stakeholder touchpoints.

1 PROCESS ENGINEERING

Design the right process before automating it.

2 AGENT DEVELOPMENT

Define each agent's role, capabilities, and guardrails.

3 WORKFLOW DEVELOPMENT

Wire the agents into an end-to-end process.

4 PROCESS VALIDATION

Pressure-test end-to-end process until ready to ship.

5 PROCESS SUSTAINMENT

Deploy, monitor, audit, govern, and continuously update and improve.

TESTING & ITERATING

Skip a phase,
and you increase the odds of
joining the 40% of projects that get canceled
before reaching production.

Follow them carefully, and you build
a workflow that earns trust.

THE SOLUTION ARCHITECTURE

1 PROCESS ENGINEERING

Design the right process before automating it.

Agentic AI gives you the ability to automate processes that were previously difficult or even impossible to automate. But it can also magnify the flaws in those processes. If you start with a bad process, automation will just scale the badness. So, before any AI work begins, we engineer the right process.

1A SELECT PROCESS

We help you identify processes that meet three criteria: they happen often, they matter to the business, and they are clunky to execute well. From that list, your first project should be a smaller, well-bounded one. Earlier wins build internal support, strengthen team confidence, and create momentum for higher-profile work later. Trust in the process has to be earned, and a clean win on a smaller project is worth more than a heroic effort on something that should have been a phase two priority.

1B DEFINE CURRENT STATE

We map the process as it actually runs today, not as it appears in your SOPs or process documentation. We quantify duration, time on task, quality, reliability, and other relevant metrics. Live mapping uncovers the workarounds, handoff failures, and rework loops that almost always exist in real-world processes.

1C DEFINE SUCCESS

What does great look like? If this process were running perfectly, what would the outcomes look like in terms of quality, speed, efficiency, and reliability? We set specific success criteria, with measurable metrics whenever possible, so we have something concrete to evaluate against later. Without this step, you are flying blind.

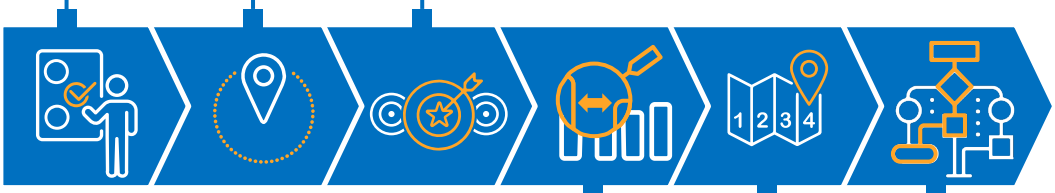

Product Requirements Document (PRD)

1D ANALYZE GAPS

We compare current state performance against your success criteria and identify what is working, what is not, and what is missing entirely. If new metrics need to be tracked, we start measuring them now to establish a baseline.

1E DESIGN FUTURE STATE

We design the future state process from the ground up, leveraging what works and changing what does not.

So far, this looks like any process engineering project. Now we weave AI Agents into the conversation.

1F ARCHITECT WORKFLOWS

Based on the future state process, we architect a blueprint defining what your agentic process will entail. The blueprint specifies how the work will be organized into one or more workflows, which agents will be involved, what capabilities they will need, what resources they will draw on, where handoffs will occur, what guardrails will apply, and where human oversight will sit.

The output of Phase 1 is your Product Requirements Document (PRD), a blueprint that we align on before any code is written.

WARNING: Avoid Common Pitfalls

Five ways agentic AI projects go off the rails before they even start:

#1

Picking the wrong starter project.
The bigger the ambition, the bigger the failure point. Start small, work out the kinks, let success breed success and earn you the right to go for bigger wins.

#2

Skipping the process engineering step.
Automating a bad process just makes it bad at scale and amplifies its negative effects. Optimize your process first, then make it agentic.

#3

No defined success criteria.
Gartner attributes the 40% projected cancellation rate of agentic AI projects to escalating costs, unclear business value, or inadequate risk controls (Gartner, 2025). Each of those is a measurement problem first.

#4

Short-changing the testing.
Let's face it, testing is boring, monotonous, and time-consuming. Who has time for that? Well…the companies that get AI right make the time to test each agent, each workflow, and each end-to-end process thoroughly before even considering going live.

#5

Assuming that humans aren't needed.
The bottom line is that AI can't work effectively unless humans are involved in designing, building, testing, governing, and updating it. Those humans need to have new skills and do their work differently, but they are still necessary, no matter what the hype may tell you.

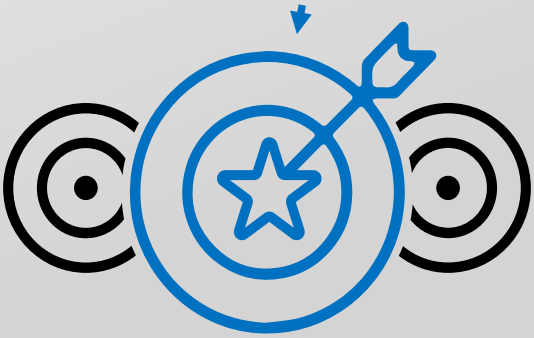

An Email Management Process

Now that we've optimized the process, we want to make it agentic, so we sketch out the process at a high level, starting to define what agents and workflow steps we'll need. This forms the basis for our blueprint.

AN EXAMPLE:
A PRD would spell out a high-level design of an agentic process. Here's an example using an email management process.

Emails arrive at the info@ inbox. Spam and other messages that do not warrant a response are filtered out.

Next, we analyze the tone, sentiment, and requests in the email.

Based on those things, we generate specific requests and route them to the other agents.

1. FILTER INTAKE

2. ANALYSIS

Agent
CINNAMON
Empathetic Listener

3. ROUTING

Agent
HOLLER
Traffic Controller

Three specialized agents research our internal knowledge base for the right answers, fill in knowledge gaps with external web research, and compose an on-brand draft response.

4. ANSWERING

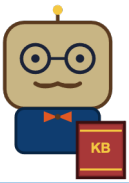

Agent
HATCH
Domain Expert

5. RESEARCHING

Agent
SCRATCH
Research Assistant

6. DRAFTING COPY

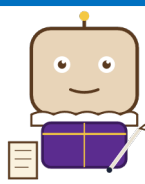

Agent
SHAKES
Email Copywriter

The draft is evaluated by an agent.

The Human-in-the-Loop (HITL) reviews the draft and (1) approves, (2) provides feedback, or (3) aborts the process.

7. QUALITY CHECK

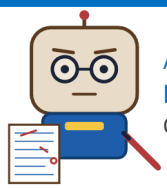

Agent
KANSA
QA Reviewer

9. HITL REVIEW

Human
APPROVER
HITL Reviewer

10. SENDING

DELIVERY
MS365 trigger

Based on that, the Gatekeeping agent approves the draft or routes feedback to the appropriate agent(s) to produce another draft (up to 3 loops).

8. GATEKEEPING

Agent
HOLLER
Traffic Controller

If there is HITL feedback, the Traffic Controller agent routes it back to the appropriate agent(s) to produce another draft and re-submit for HITL Review (up to 3 loops). Once a draft is approved, the Traffic Controller generates/formats the email and sends it to the person who submitted the inquiry.

THE SOLUTION ARCHITECTURE

2 AGENT DEVELOPMENT

Define each agent's role, capabilities, and guardrails.

Each agent in your workflow has a unique part to play. Define an agent's parameters with too much room for interpretation, and the agent will improvise in ways you do not want. Define them too rigidly, and the agent will not adapt. The sweet spot is somewhere in between, and finding it takes deliberate craft.

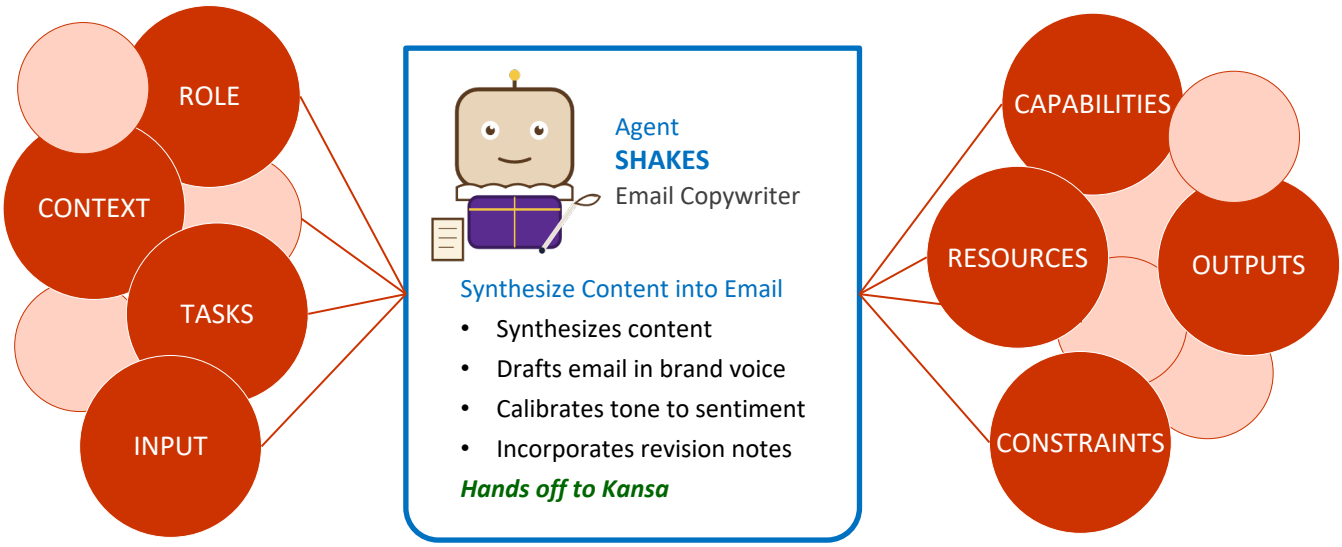

For each agent, we define eight parameters at a minimum. This page and the next page walk through each parameter.

ROLE

What this agent does, and just as importantly, what it does not do. Example: a summary agent that produces talent review summaries should never conduct its own employee evaluations or fabricate data points. Narrow roles improve quality, reliability and ease troubleshooting later.

CONTEXT

What the agent needs to understand about the broader goal, the process, the organization, and its place in the workflow. Context is the most undervalued lever in agentic design. We will return to this point in the Design Principles section.

TASKS

Specifically, what the agent does when activated. Tasks often include sub-tasks, branching logic, and decision criteria. We also decide which AI model powers each agent based on the task's complexity. Deep reasoning tasks may justify a more capable, slower model; simple lookups warrant a faster, more efficient one.

INPUTS

What information the agent needs from upstream agents or systems. The agent does not need every byte of upstream output. It needs the specific information required to do its job in a format it can use.

THE SOLUTION ARCHITECTURE

2 AGENT DEVELOPMENT (continued)

Define each agent's role, capabilities, and guardrails.

CAPABILITIES

The specialized knowledge and skills the agent draws on. This might include analytical methods, brand voice guidelines, feedback frameworks, routing rules, or coaching frameworks. We give each agent access to best practices for the things it needs to do well.

RESOURCES

Shared assets that multiple agents may call on at the point of need: company information, branding guidelines, SOPs, case studies, prior work products, and similar artifacts. These live in a Knowledge Base that agents can access using retrieval augmentation, a method that lets agents look something up when they need to. Each agent gets access only to the resources relevant to its role.

OUTPUTS

What the agent's output looks like, how it is formatted, and what criteria it must meet. Does the summary need to stay under 500 words? Does the chart need a specific layout? Does the response need to be plain text or HTML? We specify the requirements clearly and supply annotated examples of what bad, good, and great outputs look like, so there is no ambiguity.

CONSTRAINTS

Where the agent must stop, escalate, or check in with a human. We anticipate the ways the workflow can drift off-course and build in protections. There is also a setting called temperature that controls how predictable versus how creative the agent should be. We dial it toward predictability when the task is rules-based, and toward creativity when the task calls for genuine judgment. We rarely anticipate everything upfront, so guardrails get updated as we learn.

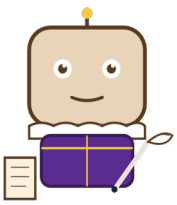

Once we flesh out all these details for each agent, we vibe code them into HTML and place that code into the system instructions of each agent, so that each agent is fully equipped to do its job.

Example of an agent system instructions form in the Cassidy AI harness

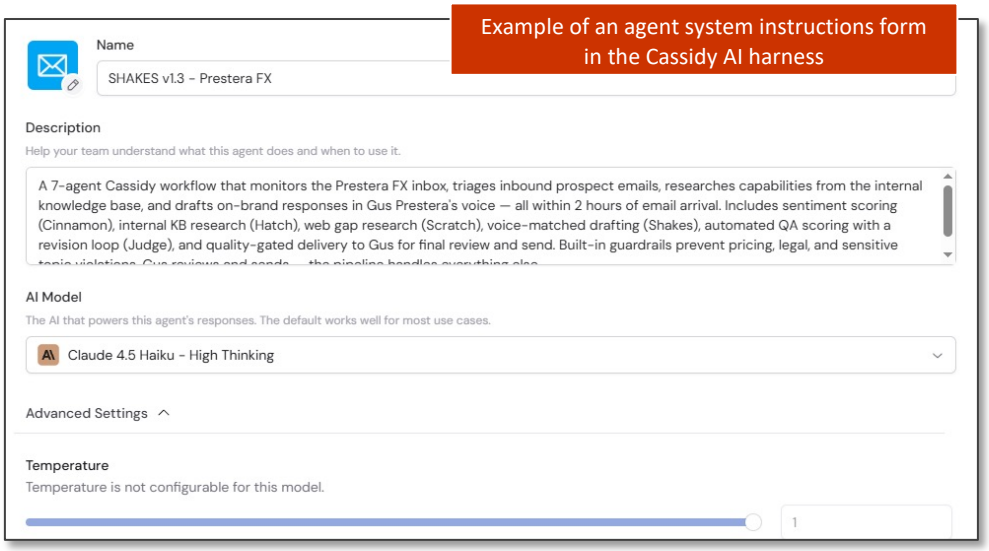

Why Specialist Agents Beat Generalists

It is tempting to build one big, all-purpose agent that can execute the whole workflow on its own.

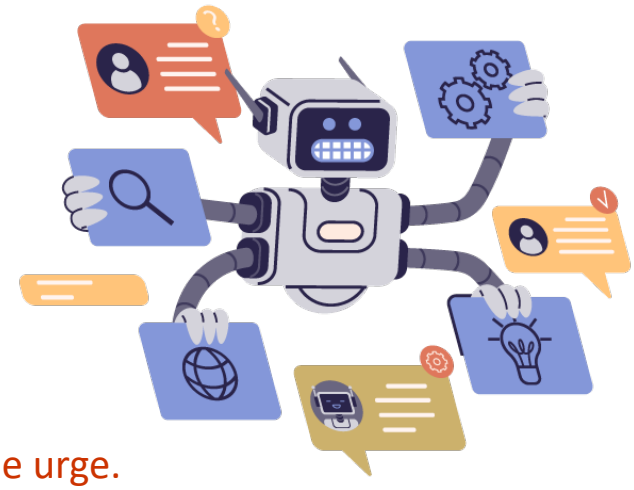

Resist the urge.

For most workflows, an AI Agent Team will be more effective than a single all-purpose AI Agent.

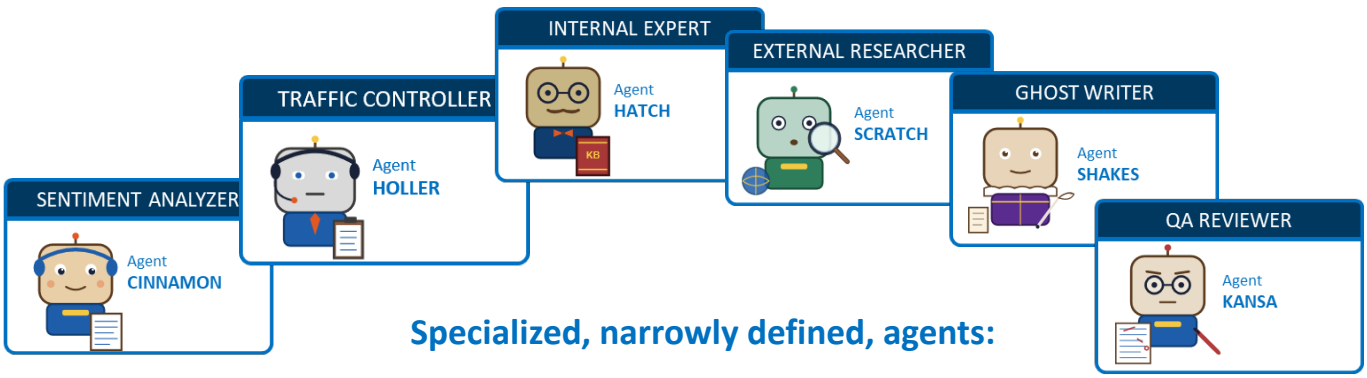

Specialized, narrowly defined, agents:

Perform their tasks accurately, reliably, and at scale.
All-purpose agents perform each task with a larger margin of error, which get amplified as the frequency and repetition increase.

Use their attention budget (i.e., "context window") more efficiently, which translates into higher quality outputs.

Are easier to troubleshoot: a routing error points to the orchestrator, an off-brand sentence points to the writing agent.

Consume less tokens, because you can match the right model to each agent and task, which helps you better manage AI usage costs.

The slight added complexity of more agents is more than offset by the gains in quality and reliability.

THE SOLUTION ARCHITECTURE

3 WORKFLOW DEVELOPMENT

Wire the agents into an end-to-end process.

With your agents built and tested individually, the next step is to wire them together into a working end-to-end process. This is where the workflow starts to come alive, and it is also where most of the unexpected complexity shows up.

WORKFLOW DESIGN DOCUMENT

We refine the PRD into a more detailed Workflow Design Document (WDD), which specifies:

- What triggers the workflow, how steps connect to one another, and how the workflow closes out
- The sequence of steps, including loops for review and iteration cycles
- Whether an orchestrator agent reviews each output and routes it forward
- Where and how the human-in-the-loop participates
- Which agent owns each step
- Which steps may be deterministic, because they are rules-based and predictable, requiring no agent creativity or judgment
- The step-level instructions and Knowledge Base resources each agent needs in the moment
- The handoff format between agents
- What data gets collected, used, and saved in the activity logs

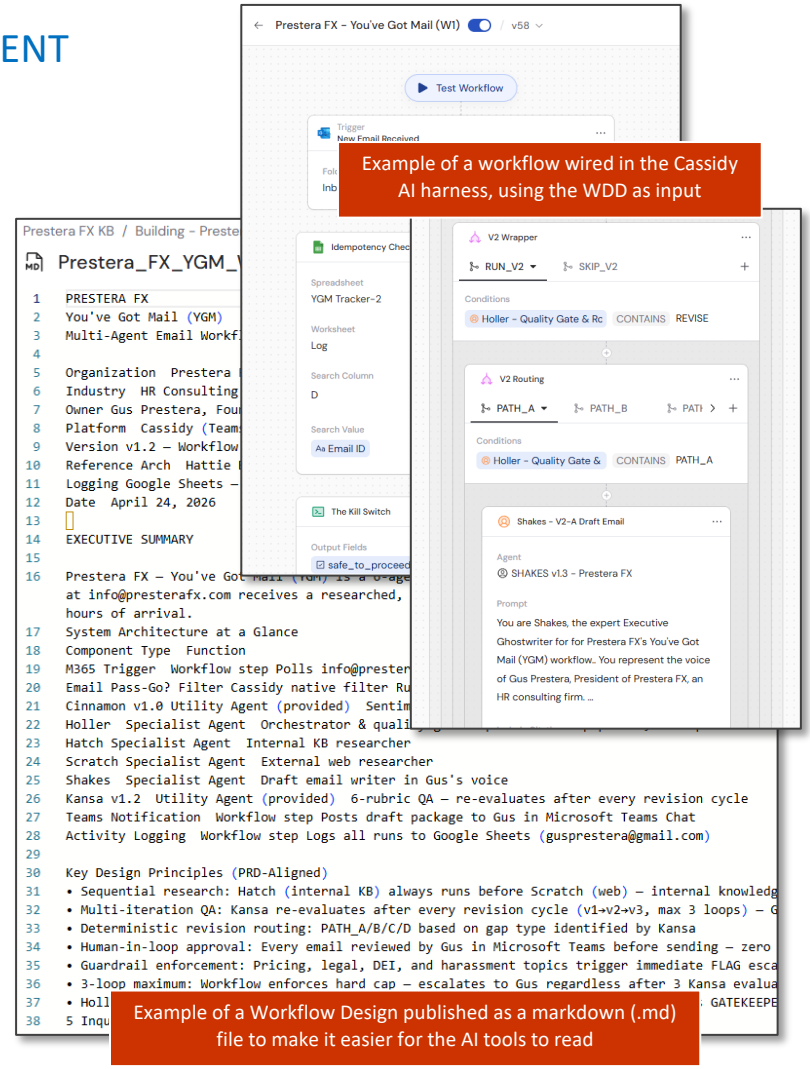

DATA DICTIONARY

The WDD includes a data dictionary that becomes the foundation for the build. The latest AI build environments can now read the PRD and WDD and generate much of the workflow code for you. From there, the workflow needs end-to-end review and testing.

CONNECTIONS

In order for the agentic process to function properly, workflows need to be connected to other systems. With the email management process, for example, our agents need to be able to read, write, and send emails as well as communicate with the HITL via Teams, so an MS365 connection is needed.

YOU'VE GOT MAIL

an agentic process from inbox to response

Example of a workflow design. While the PRD gives us a high-level blueprint, the WDD requires us to think through the details of each step in the process.

1. FILTER INTAKE

MS365 INBOX
15-min polling trigger

2. ANALYSIS

Agent
CINNAMON
Empathetic Listener

Sentiment & Intent Analysis

- Reads each email with care
- Identifies tone & sentiment
- Captures intent & urgency
- Documents context notes

Hands off to Holler

3. ROUTING

Agent
HOLLER
Traffic Controller

Routing & Coordination

- Receives sentiment data
- Routes to right specialist
- Manages team handoffs
- Tracks workflow status

Hands off to Hatch

4. ANSWERING

Agent
HATCH
Domain Expert

Senior Consultant Expertise

- Searches internal KB
- Applies domain expertise
- Identifies knowledge gaps
- Flags needed research

Hands off to Scratch

5. RESEARCHING

Agent
SCRATCH
Research Assistant

Web Research Specialist

- Researches web sources
- Verifies facts & citations
- Closes knowledge gaps
- Compiles research package

Hands off to Shakes

6. DRAFTING COPY

Agent
SHAKES
Email Copywriter

Synthesize Content into Email

- Synthesizes content
- Drafts email in brand voice
- Calibrates tone to sentiment
- Incorporates revision notes

Hands off to Kansa

7. QUALITY CHECK

Agent
KANSA
QA Reviewer

Evaluate Email Response

- Reviews accuracy & tone
- Scores quality on rubric
- Edits or returns for revision
- Logs feedback for learning

Returns to Holler

8. GATEKEEPING

Agent
HOLLER
Traffic Controller

Decide if ready for HITL review

- Reviews QA results
- Decides: ready or revise?

If more work is needed, routes to Hatch, Scratch, or Shakes

If ready, sends to Human

9. HITL REVIEW

Human
APPROVER
HITL Reviewer

Evaluates Email Response

- Reviews accuracy & tone
- Approves if response is ready to send
- If not, can abort or submit revision requests

Return to Holler

10. SENDING

DELIVERY
MS365 trigger

Response Email Delivered

- Holler formats email
- Sends from info@
- Logs conversation
- Tracks outcomes for continuous improvement

Process Ends

Max 3 quality check revision loops

Max 3 HITL revision loops

Test Early and Often

Testing, refining, and iterating are critical aspects of any development project, but it is especially important to test your agents and workflows early and often, because generative AI can sometimes work in unexpected ways.

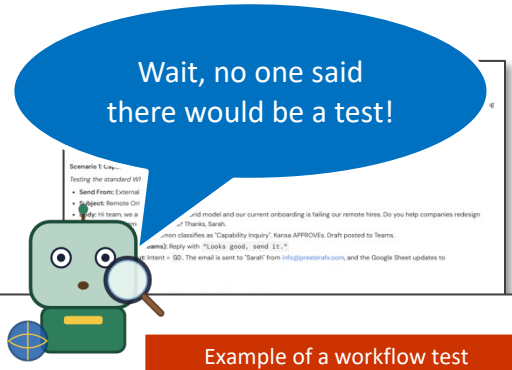

Scenario 2: Hot Lead

Testing W1 urgency flagging and the W2 "REVISE – PATH A" (Tone/Style).

- **Send From:** External Email
- **Subject:** Urgent: Leadership Development Program
- **Body:** Hi Gus, we urgently need help with a leadership development program for our mid-level managers. We have budget approved for Q3. Can we set up a call this week? Best, Marcus.
- **Expected W1 Output:** Cinnamon classifies as "Hot Lead" with "High" urgency. Kansa APPROVEs. Draft posted to Teams.
- **Your HITL Response (Teams):** Reply with `"Make it sound a bit more urgent and cut the second paragraph."`
- **Expected W2 Output:** Intent = `REVISE`. Holler routes to PATH A (Shakes only). A new draft is posted to Teams.

Example of a workflow test scenario we can use to test the steps and agents

As soon as we create an agent or a workflow, we start testing and iterating until that agent is consistently operating within our established parameters.

UTILITY AGENTS

Throughout the development process, we are working with a fleet of agents that we generally refer to as "utility" agents. These agents will not participate in the agentic workflows, but they can help us design, code, and test the agentic process.

ARCHITECT AGENT

Our architect agent, for example, helps us write the PRD, WDD, and system instructions, so it knows the agentic process well enough that it can also help us generate the quality standards, testing criteria, test scenarios, and test data…and can even help us conduct our testing.

RECURSIVE SELF-ASSESSMENT (RSI) LOOPS

RSI involves agents checking agents. When given tight testing criteria, rubrics, and parameters, Quality Assurance (QA) Agents can review the output of other agents, score them, and provide corrective feedback, so that the agent team can iterate its process and repeat the loop until it gets it right, or right enough.

This can happen live, in-process, to catch critical, time-sensitive issues, like making sure that outputs are accurate and complete. It can also happen post-process as part of routine audits conducted to gauge performance and uncover weak spots. Through these recursive loops, agent teams learn and get better.

TEST SCENARIOS

Test scenarios provide us with all the inputs needed to test an agent or workflow, helping us kick the tires on different use cases, edge cases, situations that trigger our guardrails, and ones that vary in complexity, difficulty, nuance, and impact. In this inbox management example we've been referencing, the test scenarios might include situations that involve casual inquiries (e.g., "Do you guys do __?"), hot leads ("I have an urgent request…."), out-of-bounds ("I need legal advice about…."), and so on.

THE SOLUTION ARCHITECTURE

4 PROCESS VALIDATION

Pressure-test end-to-end process until ready to ship.

There is a difference between a workflow that is operational and a workflow that is ready to ship. Operational means it runs. Ready to ship means it runs reliably, produces the experience you want, and earns your trust. The gap between those two states is where most agentic AI projects either pass or fail.

Before we deploy to a live production environment, we put the whole end-to-end process through its sea trials, looking for leaks and flaws. We test not only how each individual workflow and step performs but how the cumulative user experience comes together. Once we have run enough scenarios and fixed enough issues, we bring in beta testers who introduce a wider range of human variability.

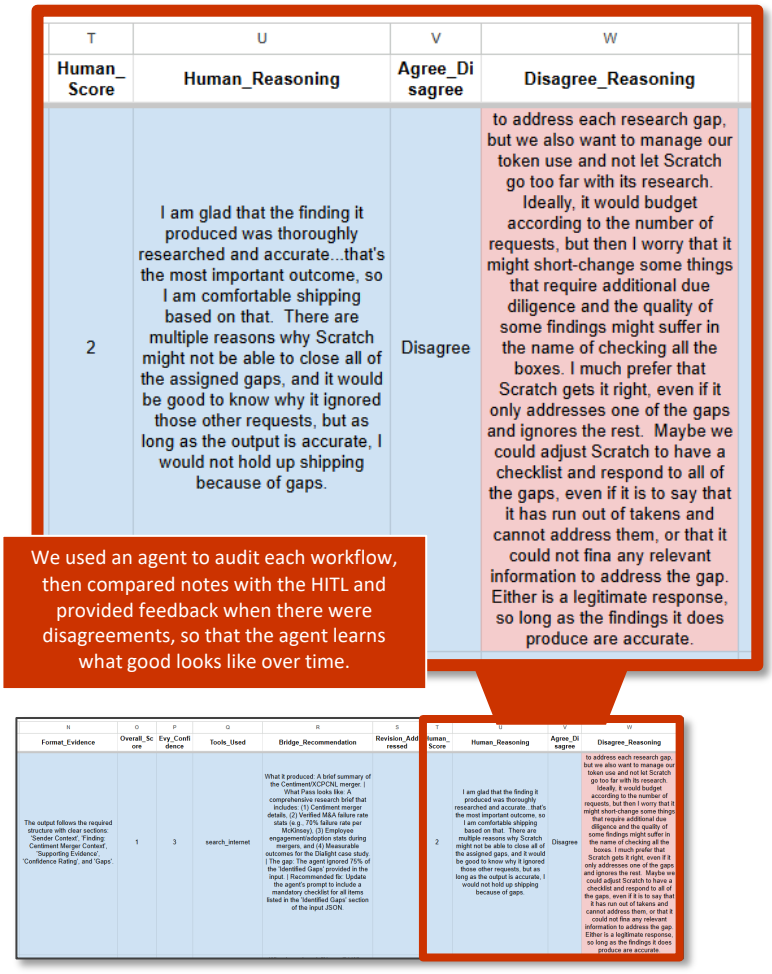

T	U	V	W
Human_ Score	Human_Reasoning	Agree_Di sagree	Disagree_Reasoning
2	I am glad that the finding it produced was thoroughly researched and accurate...that's the most important outcome, so I am comfortable shipping based on that. There are multiple reasons why Scratch might not be able to close all of the assigned gaps, and it would be good to know why it ignored those other requests, but as long as the output is accurate, I would not hold up shipping because of gaps.	Disagree	to address each research gap, but we also want to manage our token use and not let Scratch go too far with its research. Ideally, it would budget according to the number of requests, but then I worry that it might short-change some things that require additional due diligence and the quality of some findings might suffer in the name of checking all the boxes. I much prefer that Scratch gets it right, even if it only addresses one of the gaps and ignores the rest. Maybe we could adjust Scratch to have a checklist and respond to all of the gaps, even if it is to say that it has run out of takens and cannot address them, or that it could not fina any relevant information to address the gap. Either is a legitimate response, so long as the findings it does produce are accurate.

We used an agent to audit each workflow, then compared notes with the HITL and provided feedback when there were disagreements, so that the agent learns what good looks like over time.

EVALUATION CRITERIA

To keep things simple, we evaluate against a few criteria per step and a few for the workflow overall, then run lots of scenarios and score each one using a three-point scale:

3. **Ship As Is**. Performance is acceptable as-is.

2. **Ship with Edits**. The issue is minor and can be caught by human-in-the-loop controls.

1. **Do Not Ship**. Fix this issue before deploying.

SHIPPING TARGET

Our shipping target—the threshold we need to meet before releasing a workflow in a live environment—is at least **80%** of test runs scoring 3s, the remaining **20%** scoring 2s, and zero scoring 1s. At that point, we can confidently deploy, knowing that major deviations will not slip through and the smaller deviations will be caught by the HITL controls we have in place. It may take hundreds of test runs to reach that target. That is by design. The process needs to earn our trust.

COMMON PITFALL: Rush To Deploy

The agentic process must be made to earn the organization's trust before going live. Rushing the validation work is counter-productive and may erode trust, which can undermine your broader AI efforts. This has already been a painful lesson for many organizations that let their enthusiasm get the best of them.

"Just 5% of integrated AI pilots are extracting millions in value, while the vast majority remain stuck with no measurable P&L impact."

— MIT Project NANDA, *Challapally et. al (2025)'s The GenAI Divide.*

THE 5 PERCENT DIFFERENCE

The difference between the 5% that succeed and the **95% that fail** is almost always "flawed enterprise integration."

With the discipline to design the workflow correctly, pressure-test it, and force the Human and AI Agents to learn, adapt, and prove their trustworthiness, those integrations are much more likely to yield ROI for the organization.

"Most fail due to brittle workflows, lack of contextual learning, and misalignment with day-to-day operations" (Challapally, 2025).

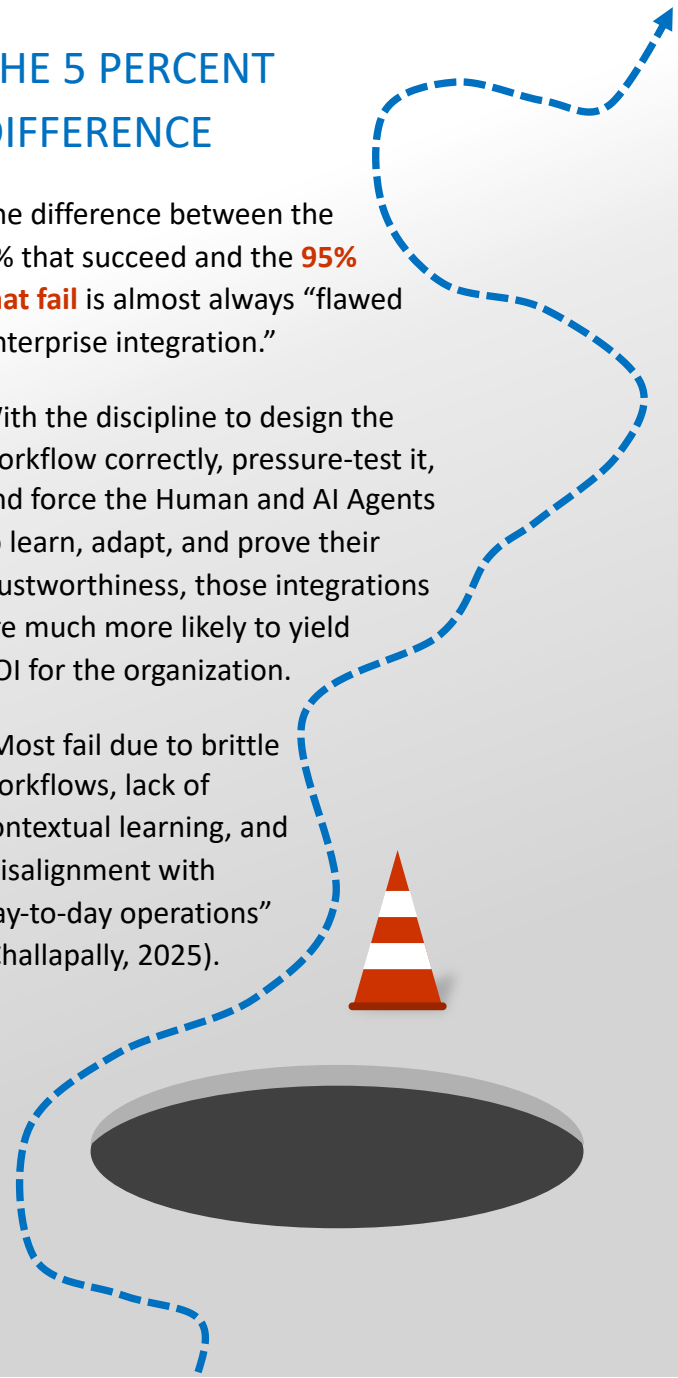

THE SOLUTION ARCHITECTURE

5 PROCESS SUSTAINMENT

Deploy, monitor, audit, govern, update, and continuously improve.

Deploying the workflow is the start of the next phase, not the end of the project. Live workflows degrade over time as the business changes, the systems they connect to evolve, and the use cases expand. Sustainment is how we keep workflows sharp.

After deployment, we continue to monitor and audit the workflow, run additional evaluations, and refine the workflow. The goal is to drive the 20% of 2-scoring outcomes down to a level that matches your tolerance for human review. If the use case allows it, we eventually phase out the in-the-loop controls and shift to on-the-loop controls instead.

The distinction matters:

SUSTAINMENT

Human
GOVERNANCE
Process Owner

Validate, Deploy, and Sustain

- Establish governance
- Test end-to-end process
- Decide when ready to ship
- Deploy to live environment
- Monitor and troubleshoot
- Update and maintain
- Audit and improve

Ongoing sustainment effort

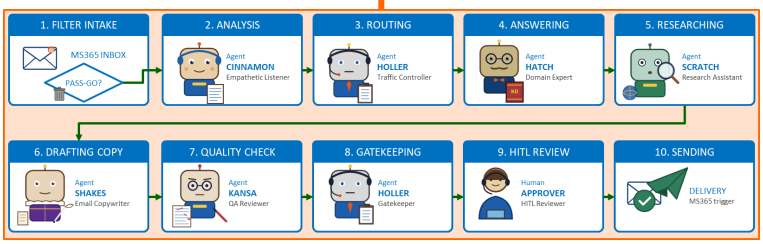

HITL:
HUMAN-IN-THE-LOOP

A human reviews specific outputs before they go out. This is the right approach when stakes are high and errors are costly. Examples include manager-facing 360 reports, executive talent review summaries, and external-facing communications.

HOTL:
HUMAN-ON-THE-LOOP

A human monitors performance trends, audits samples, and intervenes when patterns suggest something is drifting. This is the right approach when the workflow has earned enough trust to run unattended most of the time. Examples include policy Q&A responses, low-stakes employee inquiries, and routine status notifications.

HUMANS
NEVER OPTIONAL

Humans are essential to monitoring, troubleshooting, auditing, and updating the workflow as your business evolves. Even with processes that no longer require HITL or HOTL and are capable of self-improvement, there will still be a need for human governance of these systems. Only humans can take accountability, and accountability requires ongoing process ownership.

The Essential Human

If you haven't already, you will hear technology vendors hyping how many FTE hours of work, salaries, and budget dollars their products can save you. Perhaps they will even tell you that their AI agents and workflows can do the work of an entire department. Don't be afraid to call that out as the malarky you know it to be.

AI agents and workflows perform repeatable, well-structured tasks, and if properly designed and tested, AI can even perform those tasks better, faster, more consistently, and at greater frequency and volume than their human counterparts. They can increase the scale of your operations tenfold without breaking a sweat, though of course they may break your token budget.

It is also difficult to argue with the fact that many of those tasks are ones that currently consume a portion of your team's day-to-day workload, so offloading those tasks to AI will inevitably disrupt their jobs. But you have a choice about what kind of disruption that involves.

WHAT'S YOUR CHOICE?

Will you simply use AI to reduce headcount and lower your budget? Or will you work with your team to redesign their jobs, upskill them, and help them re-allocate their time towards higher value tasks that only humans can perform? Or will you do a combination of both, recognizing that not everyone on your team is willing and able to make this pivot with you?

While working through these difficult decisions, we encourage all of our clients to think carefully about what Nathaniel Whittemore (May 2026) calls the ***Human Premium***. In his podcast, the AI Daily Brief, Whittemore describes seven categories of value that evaporate when we remove humans from critical processes.

Consider the Talent Acquisition Process, for example. Ask where your company, the candidates, hiring managers, and stakeholders benefit most from that human premium and where they could benefit most from the scalability that agentic AI can provide.

When To Bring In a Partner

Expecting your busy HR team to invest time into building new agentic workflows can be a big ask. Here are some signs it is time to bring in outside help:

- Your process experts are too busy to learn and apply the advanced automation and AI skills required
- You are scaling beyond a pilot into production
- The workflow touches multiple HR systems or business units
- The stakes are high…compliance-sensitive processes, leader-facing reports, anything customer-visible
- You do not have the in-house engineering, design, or data capabilities to build and sustain it well
- Your team has the capacity to be smart buyers and orchestrators but not the bandwidth to build it themselves
- You want a thought partner who has done this before and can help you avoid the common traps

Let Prestera FX boost your AI deployment and adoption efforts

Our Services

Consulting Services
- AI Enablement Strategy
- Change Management
- Performance Consulting
- Group Facilitation

People Analytics
- AI-Generated Dashboards
- Employee Sentiment
- Program Evaluation
- Leader 360 Assessment

Program Design
- AI Mindset, Fluency, Skills
- Leadership Development
- Career Development
- New Hire Onboarding

Creative Services
- AI-Generated Content
- Agentic AI Workflows
- eLearning & Video
- Workshops & Virtual ILT

How can we help you get business results with AI?

Strategy

Getting your organization started:

- Tech Audit & Readiness Assessment
- AI Enablement Strategy Workshops
- AI Tool Selection & Architecture Design
- AI Roadmap & Change Management

Capacity

We can then supplement your team with AI specialists while also supporting your team's transition with:

- AI Mindset Education
- AI Skills & Fluency Training
- AI Tools & Resources
- AI Change Agent Training & Coaching

Content

For your programs, we can build entire AI-generated elearning courses or individual AI-generated:

- Graphics/ Animations
- Presentations
- Synthetic Videos
- Games & Sims
- Assessments
- Dashboards
- Portals

Agents

For specific use cases, we can design, build, test & deploy custom:

- Coaching Agents
- Interviewer Agents
- Simulation Agents
- Expert (RAG) Agents
- Researcher Agents
- Coordinator Agents
- Quality Control Agents

Workflows

For routine processes, we can design, build, test & deploy custom, automated, agentic workflows:

- Needs Analysis
- Course Creation
- Skill Assessments
- Program Evaluation
- Program Management

Let us accelerate your AI enablement, so your team can excel at what they do best!

DESIGN PRINCIPLES AT WORK

Across every agentic AI project, the same handful of design principles tend to separate the workflows that get shipped from the ones that get scrapped. Here are the ones we hold ourselves to.

1 **Start with Good Process**

Great agentic workflows can only be built from well-designed processes. We take time upfront to do the process engineering work, analyzing gaps between current and future states, and crafting a stronger process before adding AI. Starting with a bad process means you will end up scaling something bad at higher volume and higher speed. That is worse, not better.

2 **Provide Context**

Humans absorb and apply context effortlessly. AI models do not. We design system instructions, step-level instructions, and Knowledge Base resources to deliver the right amount of context at the right moment, so each agent has what it needs to perform. Most agentic AI failures trace back to context that was assumed but never supplied.

3 **Narrowly Defined Agents**

When agents are highly specialized, they perform their tasks reliably rather than occasionally. Specialization also makes troubleshooting easier. A routing error points to the orchestrator. An off-brand sentence points to the writing agent. Narrow agents also use their attention budget more efficiently, which translates into higher quality outputs across the workflow.

4 **Iterative Testing**

We do not wait for all the workflows to be built before we start testing. We test each agent in isolation, then test each workflow, then test the end-to-end process. We keep testing after deployment, until the workflow has earned the right to run with reduced human oversight. Testing is not a phase. It is a habit.

5 **Humans Essential**

No matter how well a workflow is designed, built, and tested, if it is important enough to deploy, it warrants ongoing human involvement. Humans monitor, troubleshoot, audit, and govern. Processes also need to evolve as the business and marketplace change, so process ownership is never a finished task. The agents you build today will need to be tuned, retrained, or retired tomorrow, and a human always needs to make those calls.

We welcome your perspective on other ways to make our Agentic AI development process more effective and impactful.

REFERENCES

- **Challapally, A., Pease, C., Raskar, R., & Chari, P. (July 2025).** *The GenAI Divide: State of AI in Business 2025*. MIT Project NANDA. Available at: https://www.artificialintelligence-news.com/wp-content/uploads/2025/08/ai_report_2025.pdf

- **Deloitte. (2024)**. *Modernizing HR: Design thinking and new technologies to help enhance employee experience*. https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consulting/servicenow-ffex-modernizing-hr.pdf

- **Gartner. (2025, June 25)**. Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027 [Press release]. https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

- **Gartner. (2026, February 10)**. Gartner research reveals CFOs' budget plans prioritize growth functions, tech, and AI in 2026 [Press release]. https://www.gartner.com/en/newsroom/press-releases/2026-02-10-gartner-research-reveals-cfos-budget-plans-prioritize-grotwth-functions-tech-and-ai-in-2026

- **HiBob. (2026, January 11)**. 23+ HR technology trends and statistics for 2025. https://www.hibob.com/blog/hr-tech-trends-statistics/

- **Schuman, E. (2025, March 25)**. 88% of AI pilots fail to reach production — but that's not all on IT. *CIO*. https://www.cio.com/article/3850763/88-of-ai-pilots-fail-to-reach-production-but-thats-not-all-on-it.html

- **Whittemore, Nathaniel (2026, May 10)**. The New Jobs AI Will Create. The AI Daily Brief (podcast). https://open.spotify.com/episode/0GljfKbZ6jMVUKtKr9kRbs?si=G76PGVA3QcCbf5snpqv3kA&nd=1&dlsi=a4e722f14ba447f4

Note that the AI Agents referenced in the Inbox Management Process example were, with permission, built or customized from agent templates created by Tyler Fisk and Sara Davison of AI Build Lab (https://aibuildlab.com/).