Inteq Executive Briefing Series

Identifying High-Impact AI Agent Opportunities

Briefing Q&A

www.InteqGroup.com

Presentation, Video and Q&A are available at:

www.inteqgroup.com/assets-identifying-high-impact-ai-agent-opportunities

Thank you for attending "Identifying High-Impact AI Agent Opportunities" executive briefing. Eight questions came in during and after the session, and each deserved more depth than the live Q&A format allowed. I have written thorough, well-reasoned answers to all of them, and my goal throughout was to make every response immediately applicable to your own agentic AI initiatives.

A few notes on the document

- **On the questions themselves.** I have included each question as it was originally submitted. In a small number of cases, I made light edits for typos, brevity or readability - never the intent of the question.
- **On how the answers were produced.** In the spirit of the presentation itself, a note on transparency: each response began as my own draft, written from my professional judgment, experience, and expertise - much like the knowledge-work the presentation describes. I then ran each draft through a frontier LLM to pressure-test it for clarity, coherence, completeness, and factual accuracy. The result is a set of answers that combine human judgment with AI-assisted refinement - a working example of the collaboration the presentation argues for.
- **On numbering.** The Q&As are numbered for ease of reference. The numbering is not a ranking - every question received the same care.

Looking ahead

This document is meant to extend the conversation, not close it. If any of the answers raise follow-up questions, prompt a different angle, or apply unevenly to your specific environment, I would genuinely like to hear about it. The questions you have already submitted have shaped how I think about this material, and the next iteration of these ideas will be better for it.

I will also be holding additional executive briefings and presentations in the coming months. If you would like to be notified when those are scheduled, or if there is a topic in this space you would like to see explored in more depth, please let me know.

With my thanks again for your engagement,

James

James Proctor | Co-Founder & Managing Director
The Inteq Group, Inc.
People | Process | Technology | Agentic AI
Author, Mastering Business Chaos
Office: 800.719.4627
InteqGroup.com | INTEQ.EDU
LinkedIn: linkedin.com/in/james-proctor-6a22329

Q1: "How do you manage knowledge worker resistance when deploying AI agents?"

Knowledge worker resistance to AI agent deployment is real, predictable, and resolvable, but only when the change management approach addresses the actual concern rather than the surface concern. The surface concern is job security. The actual concern, typically voiced by the most experienced staff, is loss of professional identity and the devaluation of hard-won expertise. Change management strategies that address only the surface concern fail because they do not engage the deeper issue.

The strategies that work are grounded in three structural commitments.

First, frame the role transformation honestly and specifically. Generic reassurance ("AI will augment, not replace") is insufficient because it is not specific to the work the person actually does. Effective change management identifies, for each affected role, what the agent will handle, what the human will continue to handle, and how the human's role becomes more valuable as a result.

The accounts payable analyst whose routine matching work is handled by an agent does not lose their job. Their role transforms from operational throughput (reviewing 200 invoices per day) to expert oversight (reviewing the agent's exception escalations, monitoring decision patterns for drift, refining decision logic based on observed outcomes, and identifying process improvement opportunities). This is more valuable work, but only if the organization explicitly designs the new role and supports the transition.

Second, engage experienced staff as the source of agent decision logic, not as the subjects of agent replacement. The decision logic that agents execute exists today as tacit knowledge held by experienced operators. Externalizing that logic - capturing what experienced staff actually do at each decision point, the patterns they recognize, the judgment they apply - is essential to agent quality. Organizations that engage their most experienced people as the architects of the decision logic create two outcomes simultaneously: better agents, and a fundamentally different change dynamic. The experienced operator becomes the authority whose expertise is being formalized, codified, and elevated, not the worker whose skills are being made obsolete.

Third, invest in capability transition rather than communication campaigns. The most common failure mode in change management is treating communication as the primary intervention. Communication is necessary but insufficient. The real work is reskilling - giving experienced operators the analytical, oversight, and continuous-improvement skills that the new role requires. This also includes training in decision-flow analysis, agent governance and monitoring, exception pattern recognition, and process design. Organizations that invest in capability transition produce staff who are confident in their new roles because they have the skills to perform them. Organizations that rely on communication produce staff who hear the right words but feel unprepared for the actual transition.

What does not work, despite sounding good in theory: town hall reassurances without role-specific transformation plans; "AI champions" programs that treat enthusiasm as a substitute for capability; gamified training with no connection to the actual new role; vague commitments that "no one will lose their job" without specifying what each person's job will become. These approaches address the surface and miss the structure. The strongest predictor of successful knowledge worker transition is whether the new role has been explicitly designed, the capability gap has been honestly assessed, and the reskilling investment has been budgeted at the start of the initiative - not as a remediation after deployment.

Q2: "How do you formally model SME tacit knowledge and decision logic for AI agent design?"

Externalizing tacit SME decision logic into agent-usable artifacts is the single most consequential business analysis activity in any Agentic AI initiative. The work that experienced subject matter experts perform at each decision point, such as the conditions they evaluate, the patterns they recognize, the thresholds they apply, the judgment they exercise, is rarely written down. It exists as accumulated expertise in the operator's mind. Agents cannot inherit this expertise by observation; it must be deliberately decomposed, structured, validated, and translated into explicit artifacts that the agent can execute and the governance team can audit.

A model-driven analysis approach uses a small set of integrated visual artifacts that work together. The first is the decision table, which structures decision logic as a matrix of conditions and outcomes. For each decision point, the analyst captures the input conditions the SME evaluates, the threshold values that trigger different outcomes, and the resulting actions. Decision tables surface gaps and contradictions immediately. When SMEs review their own decision logic in tabular form, they often discover that the rules they thought were consistent actually vary based on context, or that two SMEs handling the same decision apply different thresholds. The decision table is also directly translatable into agent decision rules with minimal reinterpretation.

The second is the decision tree or decision-flow diagram, which maps the sequence of decisions and the conditions that determine which decision is made next. Where the decision table captures the logic of a single decision, the decision tree captures how decisions chain together to produce a process outcome. This artifact is essential for agent design because agents do not make isolated decisions, they reason through sequences of decisions where each decision's outcome conditions the next. The decision tree externalizes that sequencing logic.

The third is the state transition diagram, which captures how the state of a process instance changes as decisions are made. State transition modeling is particularly valuable for agent design because it forces explicit definition of every state the work item can occupy, every transition between states, and the conditions that govern each transition. Agents operate by reasoning about state and selecting appropriate transitions; without explicit state modeling, the agent's behavior becomes ambiguous.

The fourth is the information requirements model, which specifies what data each decision requires, where that data resides, and what confidence level the SME currently has in that data. This artifact directly informs the agent's data integration architecture and the data confidence assessments that determine agent autonomy boundaries. SMEs often assume access to data they actually obtain through informal channels or compensate for through experience; the information requirements model surfaces these dependencies explicitly.

The validation discipline is as important as the artifacts themselves. Three validation techniques produce reliable externalization.

- Walkthrough validation brings multiple SMEs together to review the same decision logic; disagreements among SMEs are not problems to suppress but signals that the decision logic was never as consistent as assumed.
- Historical replay validation applies the externalized logic to past decisions and compares the predicted outcomes to what actually happened; significant divergence indicates that the logic captured does not match the logic actually applied.
- Edge case validation deliberately tests the externalized logic against unusual situations - the cases where SME judgment most differs from rule-following; the agent will encounter these cases in production, and the logic must handle them or escalate appropriately.

The output of this work is an integrated set of artifacts: decision tables, decision trees, state transition diagrams, information requirements models that together specify the decision logic at sufficient detail to inform agent design and at sufficient transparency to support governance review.

This is not theoretical modeling; it is the foundational artifact set that distinguishes an agent that operates correctly from one that approximates SME behavior with hidden gaps. Organizations that produce these artifacts before agent development consistently deploy agents that work as intended. Organizations that skip the modeling and try to capture SME knowledge through informal interviews or "AI training data" approaches consistently produce agents whose behavior diverges from expectations in ways that are difficult to diagnose and correct..

Q3: "How is Agentic AI different from RPA and traditional automation?"

Traditional automation such as RPA, workflow engines, and scripted integrations executes pre-defined steps in a fixed sequence. It excels at high-volume, rule-based tasks where the logic can be fully codified in advance. RPA bots are essentially digital workers that follow scripts: when this happens, do that. They do not interpret, reason, or adapt. Agentic AI is fundamentally different. AI agents reason over conditions, interpret ambiguous inputs, make decisions within delegated authority, and adapt their behavior when conditions change. They handle the judgment-intensive, context-dependent work that RPA cannot touch.

The practical implication is that traditional automation discovery asks "What tasks can we codify into deterministic rules?" and identifies high-volume, repetitive, rule-based work. Agentic AI discovery asks "What decisions can be delegated to an intelligent system operating with defined authority?" and identifies judgment-intensive, multi-step reasoning work. These are different questions that produce different answers. An organization that applies automation thinking to agent selection will consistently choose the wrong processes - selecting work where RPA already performs well and missing the processes where agents create step-change value through decision compression, exception resolution, and adaptive reasoning.

A simplest test: if a process can be described as a fixed sequence of structured steps with deterministic rules, it is RPA territory. If a process requires judgment, contextual interpretation, or reasoning across multiple variables to produce an outcome, it is agent territory. Most enterprise processes contain both, which is why mature implementations layer agents on top of existing automation rather than replacing it. The agents handle the decisions; the automation continues to execute the tasks.

Q4: "Which business processes are best suited for AI agents?"

Decision density measures how many judgment calls, classifications, approvals, and contextual evaluations occur within the process. Processes with high decision density are where agents create value through decision compression and consistency. This is the single strongest predictor of agent ROI. Low-decision processes, even high-volume ones, are automation candidates, not agent candidates, because the value of agents lies in reasoning, not execution speed.

Exception volume measures what percentage of process instances deviate from the happy path and require human investigation, cross-functional coordination, or manual workaround. Exception-heavy processes are where agents deliver step-change improvement, because agents reason over exceptions contextually rather than routing them to human queues. A process with 35% exception rates often represents a larger agent opportunity than a process with 5% exception rates and ten times the volume because the cost of exception handling, not the cost of routine processing, is what agents collapse.

Data confidence readiness measures whether the data required for the process's key decisions is digitally accessible, reliable, complete, and current at the point where the decision is made. This is decision-point-level data confidence, not system-level data quality. Data confidence determines agent autonomy levels: a high-decision, high-exception process with poor data confidence will require extensive human-in-the-loop operation, limiting the value agents deliver. Data confidence readiness separates genuinely deployable opportunities from aspirational ones.

In practice, a senior process owner can estimate all three markers (high/medium/low) for a given process in a single conversation. Processes that score high on all three should advance to deeper analysis. Processes that score low on one or more markers are either deprioritized or routed to specific enablement work such as process redesign, decision logic externalization, or data quality remediation before agent deployment is considered.

Q5: "How do I decide which decisions an AI agent should make autonomously?"

Agent autonomy is not a binary choice between "the agent decides" and "the human decides." It is a five-tier classification applied to each decision point within a process: Fully Autonomous, Human-on-the-Loop, Human-in-the-Loop, Human-Initiated, and Human-Only. The right tier for each decision is determined by three factors: the cognitive complexity of the decision, the data confidence available at the point of decision, and the consequence of getting the decision wrong.

- **Fully Autonomous** decisions are those where the agent has explicit logic, reliable data, and bounded consequences - the decision can be made and executed without human involvement. For example, classifying incoming invoices by type, applying standard GL coding for established vendor patterns, and scheduling routine payments within defined cash management policy. The agent decides; humans review aggregate patterns periodically rather than individual decisions.
- **Human-on-the-Loop** decisions are made autonomously by the agent, but humans monitor in real time and can intervene if patterns deviate from expectations. The agent acts; humans observe. This tier is appropriate for decisions with moderate consequence where post-hoc correction is feasible. For example, exception resolution within defined tolerance thresholds.

- **Human-in-the-Loop** decisions require human confirmation before the agent executes. The agent proposes; the human approves or revises. This tier suits decisions where consequence is high, reversibility is limited, or organizational policy requires human accountability. For example, credit decisions above defined thresholds, exception approvals that exceed tolerance limits, and any decision that crosses a regulatory line where human accountability is non-negotiable.
- **Human-Initiated** decisions are made by humans but with agent support. The agent provides analysis, recommendations, and pre-assembled context; the human decides. This tier preserves human ownership of strategic and judgment-intensive decisions while leveraging agent capability to make the human more effective.
- **Human-Only** decisions are reserved for the agent entirely. These are decisions involving novel ethical questions, fundamental policy choices, situations of high political or relational sensitivity, or any context where the organization has determined human accountability cannot be delegated. The agent does not participate in these decisions, even in an advisory role.

The autonomy classification is performed during discovery, before any agent is built. It directly determines the agent's role, the human's role, and the governance design for the process. The most valuable agent opportunities are often at the Human-on-the-Loop or Human-in-the-Loop tiers, where the agent does the cognitive work and the human provides judgment on the highest-stakes decisions, not at the Fully Autonomous tier, which is reserved for the lowest-consequence, highest-confidence decisions.

Q6: "Why do AI agent initiatives fail?"

AI agent initiatives fail for predictable reasons. Five discovery anti-patterns derail most organizations' agent selection - and recognizing them is the first step toward replacing intuition-driven selection with structured, defensible analysis.

The five anti-patterns are:

- **Automation Bias:** only identifying tasks that traditional RPA can handle, missing the judgment-intensive, context-dependent tasks where agents create the most value. This selects processes where traditional automation already works, producing marginal improvement rather than step-change results. The step-change value of agents lives in analytical and judgment-based work, not in faster execution of deterministic tasks.
- **Volume Obsession:** prioritizing exclusively on transaction volume, selecting the highest-volume process regardless of decision complexity. High volume with low cognitive complexity is RPA territory. A 50,000-transaction process with simple rules may have lower agent ROI than a 2,000-transaction process with complex judgment at every step. Volume matters, but only as a multiplier on decision-density value, not as a primary criterion.
- **Technology Push:** starting from "we have an AI agent platform - where can we use it?" rather than from genuine business pain and decision-flow analysis. This selects processes that fit the technology rather than processes where agents create business value. It produces technically successful deployments that don't move operational metrics, which is the hardest kind of failure to acknowledge because the project itself "worked."

- **Perfect-Process Fallacy:** assuming the current process is optimal (or good enough) and agents should execute it as-is. This automates a flawed process at machine speed. Agents do not just execute flawed steps faster, they make decisions based on flawed logic, route work through unnecessary handoffs, and replicate dysfunction at scale. Process readiness must precede agent deployment, and the parallel-track approach (improve the process while developing the agent) is the resolution.

- **Scope Creep Optimism:** identifying opportunities too broadly ("automate all of accounts payable") without decomposing to specific, implementable decision points. Unscoped opportunities cannot be designed, built, or measured. Discovery must produce L3/L4 level opportunities specific enough to inform agent specification - individual decisions within processes, not entire functional areas.

If an organization's agent selection process exhibits any of these five patterns, it is likely selecting the wrong processes and will discover the misalignment after investment has been committed. Structured discovery methodology - grounded in decision-flow analysis, applied through the five discovery lenses, scored against the four-dimension assessment is the correct approach. The cost of structured discovery is measured in weeks. The cost of skipping it is measured in failed deployments and in the loss of stakeholder confidence.

Q7: "How do you measure the ROI of AI agents?"

AI agent ROI is measured incorrectly by most organizations because they apply automation-era metrics to a fundamentally different capability. Traditional automation ROI focuses on task speed such as processing an invoice faster or routing a request faster. This framing produces modest, incremental business cases that struggle to justify Agentic AI investment. The decision-flow economic case is fundamentally different because it measures value in three dimensions that traditional automation does not touch.

The first dimension is decision latency elimination. Decision latency is the elapsed time work items spend waiting for human judgment - sitting in queues, awaiting approval, awaiting interpretation, awaiting cross-functional coordination. In most enterprise processes, decision latency accounts for 70-90 percent of total cycle time. Compressing this latency from days to minutes is an order-of-magnitude improvement, not an incremental one.

For example, a routine invoice that previously moved through the process in 3-7 business days now completes in minutes when the agent makes the routine decisions autonomously; A price variance that previously required 5-12 days of cross-functional investigation completes in hours when the agent reasons over contract terms and historical patterns directly.

The second dimension is exception handling cost collapse. The fully-loaded cost of investigating, coordinating, and resolving exceptions, including senior analyst time, manager review, cross-functional coordination, vendor communication, rework cycles, etc. - is one of the largest hidden costs in enterprise operations. Agents that reason over exceptions contextually collapse this cost by resolving what human queues currently absorb. In exception-heavy processes (35 percent or more exception rates), this is often the largest single source of agent value, dwarfing the value created by accelerating the happy path.

The third dimension is human capital redeployment. The value created when skilled professionals shift from routine cognitive work, reviewing clear matches, applying standard policies, approving straightforward transactions, etc., to strategic work that only humans can do, including complex negotiation, novel exception resolution, process improvement, and relationship management. This is the most underestimated dimension of agent ROI because it does not appear as a line-item cost reduction; it appears as increased capacity for the work that creates competitive advantage.

The proper economic methodology is Agent Value Stream Mapping: map the current-state process step by step with elapsed time, queue time, exception cost, and human time at each decision point, then map the future-state process with the same lens. The difference is the agent ROI. Organizations that frame the case this way produce business cases that are both more accurate and more compelling than traditional automation business cases because they measure the dimensions where agents create their distinctive value, not the dimensions where automation already performs well.

Q8: "How do I identify the right business processes for AI agents?"

Identifying the right business processes for AI agents requires a structured three-stage methodology that moves from a broad portfolio of candidate processes to a prioritized shortlist of high-impact opportunities with clear next steps. Each stage produces a specific output and applies specific analytical instruments, transforming process selection from intuition into a defensible discipline.

Stage 1 is the Rapid Screen, executed in approximately one day. Apply three markers to your top 15-20 candidate processes: decision density (how many judgment calls, classifications, and contextual evaluations occur within the process), exception volume (what percentage of instances deviate from the happy path), and data confidence readiness (whether the data required for key decisions is digitally accessible, reliable, and current). Score each process as high, medium, or low on each marker. Processes that score high on all three advance. This eliminates 60-70 percent of candidates in a single exercise - focusing subsequent analytical effort on the genuinely viable opportunities.

Stage 2 is the Discovery Deep-Dive, executed in 1-2 weeks per shortlisted process. Apply the five discovery lenses to identify the specific decision points and activities where agents create value: Cognitive Task Analysis reveals where the process demands judgment versus rule-following; Decision Authority Mapping identifies every decision point and whether it can be delegated to an agent; Information Flow Analysis traces where information gets stuck or degraded; and the remaining lenses uncover patterns invisible to traditional process assessment. Classify each identified decision point by cognitive complexity and delegation potential using the five-tier autonomy taxonomy. This stage produces the opportunity register - a catalogue / pipeline of specific, scoped agent opportunities within each process, defined at the decision-point level rather than the process level.

Stage 3 is the Prioritized Portfolio, executed in approximately one week. Score each opportunity from the register across four dimensions: business value, technical feasibility, risk tolerance, and organizational readiness. Plot the opportunities on a value-feasibility matrix to identify Quick Wins (high value, high feasibility), Strategic Bets (high value, lower current feasibility but worth pursuing as enablement work proceeds), and Future Opportunities (lower priority for current investment). Sequence the prioritized opportunities by dependencies - some opportunities require shared infrastructure or capability assets that earlier deployments will produce. Present the prioritized portfolio to leadership for investment decision.

The methodology is deliberately structured for execution. Stage 1 takes a day and produces a defensible shortlist. Stage 2 takes 1-2 weeks per process and produces specific, scoped opportunities. Stage 3 takes a week and produces an investment-ready prioritized portfolio. Total elapsed time from broad portfolio to investment-ready opportunities: typically 4-6 weeks for an organization screening a meaningful set of candidates. This is dramatically faster than the typical organic discovery process, which often takes 3-6 months and produces less defensible results.

* * *

Related Posts and White Papers:

PR/FAQ Is Not a Specification

The Agentic AI Ontology Question

The Uncomfortable Truth About AI Agents

Data, Meaning, Reasoning – and Agentic AI

The Secret Sauce of Enterprise-Grade Agentic AI

Agentic AI - Breaking the Myth of the Iron Triangle

Why AI Agents Often Fail to Improve Business Processes

The Secret Sauce to Unlocking Enterprise Class Vibe Coding

Spec-Driven Development Starts with Model-Driven Analysis

* * *

Subscribe to my blog | Visit our Knowledge Hub

Visit my YouTube Channel | Connect with me on LinkedIn

Check out our business analysis Training Courses and Consulting Services

Contact us at info@inteqgroup.com