Inteq Executive Briefing Q&A

Q1: We've invested heavily in Robotic Process Automation (RPA) over the past few years. Is this investment legacy and no longer has value?

It's a legitimate concern regarding sunk cost. However, RPA continues to deliver real value by significantly reducing manual effort on structured, rules-based repetitive tasks. It's not at all that RPA is legacy, or failed in some way, or is no longer useful. It's that RPA succeeded at what it was designed to do, and the organization has now reached the boundary of what RPA can achieve.

Your organization has probably automated everything rule-based automation can automate. The next tier of value - decision latency, exception handling, adaptive capacity - requires a different architectural approach. That is the handoff to AI agents. RPA remains the task-execution layer that AI agents orchestrate. Agentic AI builds-on, it does not replace RPA.

Q2: If this requires business sponsorship, how do we get business leaders to own an initiative they perceive as an IT project?

Cleary you understand the organizational dynamics of sponsorship. Leadership typically delegates anything with "AI" in the title to IT. Agentic AI is a new concept – part technology and part business process. Business leaders don't need to sponsor the AI technology platform, but they need to sponsor and own the process transformation that agents enable.

Frame it this way to business leadership: Who in the organization has the authority to decide that invoice approvals below $5,000 from established vendors in good standing no longer require human review? That's not an IT decision; that's a business decision. You can probably come up with many examples of similar knowledge and judgment-based business decisions in your business process workflows.

When you enumerate the specific decisions that must be made in the managment of ongoing workflows, it becomes clear that these are business decisions requiring business authority. IT sponsors the technology platform, and the business sponsors the operating model change.

Inteq Executive Briefing Q&A

Q3: How long does this upfront analysis take, and how do we prevent it from becoming an endless planning exercise?

Valid question for sure. Traditionally, analysis initiatives often expand seemingly indefinitely. Traditionally “Rigorous business process analysis” translated into months of workshops, documentation, and planning that delay any tangible result. The concern is legitimate - analysis paralysis is a real risk with traditional business process analysis and redesign. Agentic AI analysis and process redesign is scoped to specific target processes within the scope of an initiative.

For a well-defined industry-standard process like commercial invoice processing, procurement, customer onboarding, etc., the integrated analysis of decision logic, process structure, and data confidence can be completed, typically, in six to eight weeks with a dedicated cross-functional team. Non-standard novel processes could, but not always, take a few more weeks – but not months. The key discipline is time-boxing - define the process boundary, identify the primary decisions (typically four to eight per process), define the decision elements for each (logic, data, confidence, authority, adaptation), and validate with process owners.

The deliverable is a decision architecture document - not a comprehensive process redesign of the entire organization. Each process gets its own focused analysis sprint. The risk of analysis paralysis is managed by constraining scope to one process at a time (or several highly interdependent processes in parallel), requiring deliverable decision architecture documents at the end of each sprint.

Q4. Where should we start? If we were going to apply decision-flow design to one process as proof of concept, how would we choose it?

Select the first process based on four criteria:

- First, decision density. Choose a process with multiple decisions that currently require human judgment based on knowledge, experience and expertise. The more decisions, the more opportunity for AI agent-enabled improvement.
- Second, exception volume. Choose a process where exceptions are a significant percentage of total volume, because this is where the decision-flow model delivers its most dramatic advantage over task-flow.
- Third, data availability. Choose a process where the data required for decisions is largely digitized and accessible, even if data confidence hasn’t been formally assessed. You want the analysis to be about decision logic and confidence thresholds, not about data digitization.
- Fourth, executive sponsorship. Choose a process where both the technology leader and the business process owner are willing to co-sponsor the initiative and invest in the upfront analytical work. Also see Q2 regarding sponsorship.

Invoice processing, claims adjudication, customer onboarding, and procurement qualification are all strong candidates because they typically score high on all four criteria.

Q5: Can all of my task flows be used in the new decision-based processes? Can these new data sources (e.g., policy guides) be added to get the added value?

Part 1: Yes, your existing task flows are not discarded in the shift to decision-flow design; they are the analytical raw material from which the decision architecture is built. Every task flow contains embedded decisions, business rules, execution tasks, and data access patterns that are recoverable and directly reusable.

The transition involves decomposing each task flow to extract the decisions hidden within it such as the judgment calls, thresholds, and classifications that experienced staff apply tacitly - and formalizing those as explicit decision models with defined logic, data requirements, and authority boundaries. The execution tasks (data entry, record updates, notifications, system transactions) remain intact and continue to be performed by your existing automation or by agents after a decision determines they should occur.

The key distinction is that task flows are not reused as intact end-to-end sequences plugged into the new architecture. Instead, they are restructured around the decisions that create value, with tasks subordinated to those decisions rather than organized as fixed sequential steps. This means the organizational knowledge embedded in your current processes - the rules that your best people apply, the data they evaluate, the patterns they recognize are preserved and elevated into a form that agents can execute consistently, adaptively, and at scale. Nothing is thrown away. Everything is restructured to enable the autonomous, decision-driven operation that Agentic AI requires.

Part 2: Can new data sources (e.g., policy guides) be added to get the added value? Definitely, and this is one of the most powerful advantages of decision-flow architecture. Because each decision in the process has an explicit definition of what data it evaluates, new data sources such as policy guides, contract repositories, regulatory databases, historical exception patterns, market reference data are integrated as additional inputs to specific decisions rather than requiring process redesign.

For example, connecting a procurement policy guide to the approval routing decision enables agents to validate every transaction against current delegation authority in real time; connecting contract terms to the matching decision enables agents to resolve price variances that currently require days of human investigation. Each new data source directly expands what agents can evaluate and reason about, which increases data confidence, which in turn justifies expanding the agent's authority to act autonomously.

The compounding effect is significant because each data source that you integrate makes the process progressively more capable. The first integration might enable

agents to resolve a category of exceptions autonomously. The second embeds compliance into every decision point. The third allows agents to learn from organizational history. Over time, autonomous processing rates increase, cycle times compress, and human attention shifts from routine decisions to the genuinely complex situations where expert judgment is irreplaceable. This is the compounding capability advantage discussed in the briefing - and new data source integration is one of its primary drivers.

Q6: How do we quantify the difference between incremental efficiency and step-change improvement?

Finance-oriented executives need solid numbers before they'll fund and endorse a strategic shift. Your question tests whether a step-change claim can be translated into measurable business outcomes, or whether it remains a conceptual argument. The approach is to anchor on the three dimensions:

- Speed - measured as end-to-end cycle time, not task duration.
- Quality - measured as decision consistency and exception re-work rates, not data entry accuracy
- Resilience - measured as time-to-adapt when process conditions change. Weeks of developer reconfiguration versus real-time agent adjustment.

Identify agentic AI opportunities for step-change improvement, not just incremental improvement. For example, routine invoice end-to-end processing drops from 3-7 business days to minutes, or for example, a price variance scenario drops from 5-12 days to hours. These are order-of-magnitude compressions not 10–20% incremental improvements.

Consider running a decision-latency audit on two or three high-volume processes to quantify where the time is actually consumed in the current state process. If you are not able to identify step-change improvement opportunities from integrating AI agents into the workflow, then the business case to fund the initiative is more difficult. You will need to build the case based on incremental change further bolstered by less tangible value drivers such as improvement in the adaptably/agility and resilience of the process.

Q7: How do we know AI agents won't just create new kinds of errors that are harder to detect?

This is a legitimate concern about autonomous decision-making. And it's going to come from audit, compliance, or risk leadership. Expect the question and get ahead of it. The question tests whether the business and technical teams have considered the failure modes of the new model, not just the limitations of the old one. It's an excellent question because it demonstrates that stakeholders are engaged and are critical thinkers.

The answer has two parts. First, the decision-flow model explicitly addresses this through confidence thresholds and authority boundaries at every decision point. AI agents don't have unlimited authority to act, and the governance architecture defines exactly what happens when confidence is low or conditions are ambiguous.

Second, AI agent decisions are fully auditable in a way that human decisions typically are not. When a person approves an invoice, the organization knows the outcome but rarely documents the reasoning. When an agent makes the same decision, the logic, data inputs, confidence score, and decision path are all logged. The risk isn't that agents create undetectable errors; the risk is that organizations deploy agents without the governance structures that make errors detectable and correctable.

Q8: How should we budget for this? What's the ratio of technology investment to organizational change investment?

This is a CFO or budget-owner type of question that tests whether the investment argument can be translated into a practical financial model. Based on broad industry standards (understanding that enterprise Agentic AI is still in an emergent stage) most organizations budget approximately 80% of AI initiative funding for technology (licensing, infrastructure, development, integration) and 15–20% for everything else. The technology for supporting Agentic AI is stable and well understood. Accordingly, as Agentic AI rapidly transitions from a technology initiative to a mainstream business process initiative the weighting will invert to an allocation of 40-50% to business process analysis, decision logic definition, governance design, and operating model redesign; 30-40% to technology platform and agent development; and 15–20% to organizational change management, training, and capability building.

The rationale is straightforward. The analysis and organizational work are the primary determinants of whether agents deliver value. Organizations that under-invest in analysis and over-invest in technology end up with agents that automate poorly understood processes - which means they either fail or deliver incremental gains that don't justify the total investment.

Q9: Is a data governance program and process documentation a sufficient foundation for AI agent development and deployment?

Many organizations have invested in data quality programs and process documentation (often for regulatory compliance) prior to and independent of the rapid evolution and adoption of Agentic AI. Existing data governance and process documentation are valuable but insufficient for a specific reason: they weren't designed for Agentic AI.

Data governance programs typically assess data quality at the dataset or system level. Is the vendor master complete? Are financial records accurate? However, what is missing is the analysis of data confidence at the decision-point level: when an agent

encounters the next type of decision - does this invoice match the PO?. What is the confidence that the specific data fields required for that specific decision are reliable, complete, and current?

That's a much more granular assessment. Similarly, process documentation typically describes the base-case path (aka, the happy-path) task sequence and major exception branches. Decision-flow design requires documenting every decision within the process, the logic for each decision, the authority boundaries, and the adaptation rules. Think of it this way: your current process documentation describes what happens; decision-flow design describes what is decided, by whom (or what), based on what, and within what constraints.

* * *

Connect with me on LinkedIn:
linkedin.com/in/james-proctor-6a22329