

The Agentification Playbook: How to Prioritize 50+ AI Agent Opportunities

By Rahul Jindal

Every enterprise function now has a list of 30-100 possible AI agent use cases. Ticket triage. Recruiting screening. Pay equity analysis. Meeting summarization. Workforce planning. Contract review. The list grows weekly. The question is no longer "what can AI do?" It is "what should we build first?"

Why ROI Rankings Fail

The default approach: estimate the ROI of each agent, rank by expected value, build from the top. This sounds rigorous. In practice it produces garbage rankings because the ROI estimates require conversion assumptions that are essentially made up.

How many hours does a recruiting screening agent save? Depends on whether you count the recruiter's time, the hiring manager's time, or the candidate's time. At what hourly rate? Fully loaded or base? What about quality improvement? How do you dollarize "fewer bad hires"?

Whoever controls the conversion assumptions controls the ranking. The exercise becomes political, not analytical.

“The most valuable IP your organization produces isn't its data. It's its accumulated patterns of judgment.”

The Agentification Prioritization Framework

The framework solves this by keeping impact dimensions separate rather than collapsing them into dollars. Four impact dimensions:

1. **Efficiency:** Cost and toil reduction. Doing the same work with fewer resources. The CFO cares about this one.
2. **Quality / Risk:** Error reduction, compliance, consistency. Legal and audit care about this one.
3. **Speed / Throughput:** Faster cycle times, higher capacity. Business leaders waiting on your function care about this one.
4. **Strategic Elevation:** Freeing a role to operate at a higher altitude. Enabling work that was not previously possible. The hardest to measure and often the biggest prize.

Each agent is also scored on three feasibility dimensions: technical complexity, data readiness, and change management. Impact times feasibility produces the priority score. Multiplication, not addition, so a zero in feasibility means zero delivered value.

The Strategic Weight Profiles

Here is where it gets powerful. The weights on the four impact dimensions are adjustable. Four preset profiles:

- **Balanced:** Equal emphasis. Surfaces broadly strong agents. Good starting point.
- **Cost Takeout:** Heavy on efficiency and speed. Under budget pressure, this is your profile. Ticket triage and scheduling rise. Strategic tools fall.
- **Risk-First:** Heavy on quality and compliance. Post-audit finding, this is your profile. Pay equity scanners rise. Chatbots fall.
- **Transformation:** Heavy on strategic elevation. Post-cost-takeout maturity, this is your profile. People Business Partner tools rise. Simple automation falls.

You do not walk into the leadership meeting with a pre-ranked list. You walk in with the framework. You show how rankings shift under different profiles. Leadership picks the posture. The ranking follows from their choice.

The weights are the strategic conversation. The ranking is the consequence.

What We Learned from 83 Real Agent Opportunities

When this framework was applied to a real HR portfolio of 83 agent opportunities at a 200,000-person organization:

- 4 of 8 Tier 1 agents held their position under all four profiles. These are "no-regrets" bets.
- Pure automation plays consistently scored high on feasibility but narrow on impact. They dominate Cost Takeout but drop under Transformation.
- Advisory augmentation agents (tools for People Business Partners, compensation specialists) had the highest strategic elevation scores but lowest feasibility. The biggest prizes are the hardest to ship.
- Some product areas consistently ranked low regardless of profile. Not every area is an equal agentification target.

“Data tells an agent what happened. Decision traces tell it why: the soft constraints, the exceptions, the judgment calls that actually run a business.”


Getting Started

1. List every AI agent your function could build. Be exhaustive. 50+ is normal.
2. Score each on the four impact dimensions (1-3 scale) and three feasibility dimensions (1-3 scale).
3. Calculate priority scores under balanced weights first.
4. Run the same scores against Cost Takeout, Risk-First, and Transformation profiles.
5. Identify agents stable across all profiles (no-regrets bets) and agents that shift dramatically (strategic choice signals).
6. Present the framework to leadership. Let them choose the posture.

Related Insights


Operating Model 8 min

The Operating Model Every Enterprise Services Org Is Missing




Framework 7 min

Reimagine vs Automate: Why Most AI Initiatives Pave the Cowpath



Research 8 min

Pilot Purgatory: Why 75% of AI Pilots Never Reach Production



Before you prioritize agents, measure your absorption speed

An agent portfolio only delivers value if the organization can absorb it.

Take the OMI Assessment

Shared privately. Please do not redistribute.