


STOP BUILDING HR AGENTS

Choosing the Right AI Architecture for People Functions.

A diagnostic framework to cut through vendor hype, control computational costs, and preserve procedural justice in HR AI deployment.




Choosing the Right AI Architecture for People Functions. A diagnostic framework to cut through vendor hype, control computational costs, and preserve procedural justice in HR AI deployment.

THE INSTITUTIONAL ISOMORPHISM TRAP

AGENTIC NODE DIAGRAM



LINEAR FLOWCHART



WARNING: Conflating architectures causes 300% – 500% budget overruns for identical operational outcomes.

DEFINING THE COMPUTATIONAL STACK

1. AI TASK (STATELESS)




One request, one response. No memory between interactions.

Cost: \$0.001 – \$0.01 / execution

Use Case: Summarizing a single CV.

2. WORKFLOW (PREDETERMINED)



Cost: \$0.05 – \$0.30 / execution

Chained sequence of steps. Humans design the path; AI powers specific nodes.

Cost: \$0.05 – \$0.30 / execution

Use Case: 6-step new hire onboarding sequence.

3. AGENT (AUTONOMOUS)



Cost: \$0.50 – \$15.00+ / execution

System receives a goal, dynamically selects tools, and autonomously determines the path.

Use Case: Ensure 30-day onboarding completion for all hires.

THE CORE SHIFT IN CONTROL AND COST

WORKFLOWS

Practitioners Own the Logic

Processing 50,000 Candidates: ~\$150,000


AGENTS

AI Owns the Logic

Processing 50,000 Candidates: ~\$600,000

Autonomy carries a 3x to 10x computational tax over predetermined sequencing.

The Hidden Operational Tax of Autonomy



THE ARCHITECTURAL DIAGNOSTIC PLAYBOOK



1. Complexity

Is the task genuinely complex, or just complicated?



2. Economics

Does the value improvement justify the 3x-10x cost multiplier?



3. Capability

Can current models reliably execute the subjective nuances?



4. Risk Profile

What is the cost of a delayed error detection?

Deploy agents only when all four pillars align. Otherwise, optimize a workflow.


Diagnostic 1: Is the Task Genuinely Complex?



EVIDENCE: SALESFORCE

Transitioned global customer success planning from static models to agentic scenario planning.

Achieved 40% faster cycle times and 30% improved executive insights, justifying the 6x computational cost increase.



COMPLEX (Agent)

Example: Workforce Scenario Planning

If the optimal next step depends on unpredictable intermediate results, an Agent is justified.

Diagnostic 2: Does the Value Justify the Cost?

Break-Even Analysis



EVIDENCE: UNILEVER NATURAL EXPERIMENT


THE SETUP: Built a dynamic, agentic assessment for graduate recruitment (\$8.00 per candidate).

THE PIVOT: Attempted to scale to high-volume retail hiring, but costs were prohibitive. Redesign as a structured workflow (\$0.80 per candidate).

THE OUTCOME: Workflow reduced costs by 90% while predictive validity dropped by only 3%.

Diagnostic 3: Can AI Reliably Handle the Nuance?

System Component Pipeline



EVIDENCE: JOHNSON & JOHNSON

THE SETUP: Piloted agentic assessment for VP-level roles.

THE FAILURE: The autonomous agent systematically undervalued candidates from startup backgrounds due to training data biases.

THE FIX: Pivoted to Hybrid Workflow. AI handles structured data; human experts handle contextual interpretation of career transitions.

Diagnostic 4: What is the Cost of Error?



EVIDENCE: IBM

Segmented AI deployment based strictly on this matrix.

Compensation & promotion utilize strict human-in-the-loop workflows; internal logistics utilize autonomous agents.

The Architecture Selection Matrix

DIMENSION	Pure Workflow	Hybrid Architecture	Autonomous Agent
Logic Type	Deterministic 	Segmented 	Dynamic 
Cost Profile	\$ (\$0.05 - \$0.30)	\$\$ (\$0.20 - \$1.00)	\$\$\$ (\$3.00 - \$15.00+)
Auditability	Perfect step-tracing 	Clear escalation logs 	Complex / Opaque 
Primary Use Cases	Onboarding, Benefits, Basic Reporting	Candidate Screening, Scenario Modeling	Complex Org Design, Strategic Talent Matching

Visibility + Consistency + Correctability = High Procedural Trust



BLACK BOX (Agent)

System optimizes for aggregate outcomes.
Individual cannot self-correct.




GLASS BOX (Workflow)

System logic is visible. Individual understands
exactly why an outcome occurred.

INSIGHT

Unilever candidate focus groups revealed a 40% higher satisfaction score with Workflow transparency over initial Agent-based designs. Candidates accepted algorithms more readily when they understood how their responses affected outcomes.

Designing the Hybrid Middle



Institutional Governance & Stewardship

SUSTAINABLE AI

1. Architecture Review Boards

Reverse the burden of proof (The Microsoft Model).
Presume a workflow implementation unless a proposer can provide economic justification for agentic autonomy.

2. Employee Voice in Design

Utilize participatory design. Involve employee advisory groups to identify edge cases that purely technical teams miss.

3. Continuous Learning

Conduct annual token-cost and capability reviews. Architectures must evolve as model token costs drop and reasoning reliability improves.

Precision Over Hype

80 - 90%

The percentage of current HR AI use cases best served by transparent, cost-effective workflows or hybrid architectures.

The leaders of the next wave of HR AI will not be those who deployed the most autonomous agents. They will be those who developed the judgment to match architectural sophistication to problem complexity, preserving human dignity in technology-mediated people processes.