RJL SOLUTIONS

PRACTITIONER GUIDE / AGENTIC AI

Agent Patterns: A Practitioner's Guide

Eighteen agent patterns on five dimensions: when to use each, the tools current in November 2026, what regulators and buyers will ask, and the controls and evidence to run it in production.

PATTERN · AUTONOMY · CONTROLS

Richard White
RJL Solutions · rjlsol.com

VERSION 1.1 · PUBLISHED EDITION
NOVEMBER 2026

01 / THE ANSWER FIRST

Most lists of agent types mix five different things.

Executives are shown tables of "agent types" that put a reflex rule, a memory feature, a browser and an autonomy level side by side as if they were alternatives. They are not. They answer five different questions: how an agent decides, how agents are organized, what capabilities are added, where the agent acts and how much it may do without a person. Mixing them leads to poor design choices, such as buying a "memory agent" or treating "autonomous" as a product feature.

What this guide does

- Separates the five dimensions, so a design can be described in one line: one choice from each.
- Describes 18 patterns in four groups, each with when to use it, when not to, the representative tools, the main risk and the minimum controls.
- Adds a five-level autonomy scale, from L0 assist to L4 bounded autonomy, and a suggested starting level for each pattern.
- Maps each pattern to examples in 100 AI Use Cases and Who Runs Them.
- Shows what a bank examiner, an insurance regulator or a growth company's buyers will ask, and sets out 22 controls with their owner in each line of defense and the evidence that proves them.

Three findings

- Most production value today sits in agentic workflows, where code fixes the path and a model performs named steps, with a person approving the consequential step.
- Risk is set by what the agent can touch and how much it may do alone, not by the sophistication of the model. Autonomy is a deployment decision, made per action.
- The tool market is moving fast. In 2026 alone, OpenAI discontinued its Atlas browser and is retiring Agent Builder, Microsoft made Agent Framework generally available, and A2A joined MCP under the Linux Foundation.

HOW TO USE IT Describe the workflow on the five dimensions (section 02), find the pattern (03 to 05), set the autonomy level (06), see what your regulator or buyers will ask (07), and assign the controls and evidence (08). The interactive version at rjlsol.com/library/agent-patterns produces a release pack for your agent.

---

02 / FIVE DIMENSIONS

Describe any agent with one choice from each dimension.

Every agent can be described on five independent dimensions. The first four are covered by the pattern groups in sections 03 to 05; the fifth is the autonomy scale in section 06.

DIMENSION	QUESTION IT ANSWERS	OPTIONS IN THIS GUIDE
A. Control pattern	Who decides the next step: code, or the model?	Rules-based automation, agentic workflow, router, tool-using agent, planner-executor, evaluator-optimizer
B. Topology	How many agents, and how do they coordinate?	Single agent, orchestrator-worker, peer hand-off team, cross-organization agents
C. Capability modules	What is added to the agent?	Agentic retrieval, memory, learning
D. Environment	Where does the agent act?	API (default), code, browser, computer use, voice, data, background events
E. Autonomy and oversight	What may it do without a person?	L0 assist to L4 bounded autonomy

A worked description

A claims intake service might be an agentic workflow behind a router (control pattern), a single agent (topology), with retrieval over the policy file (capability module), taking calls on voice (environment), at L2, with a person approving anything that changes the claim record (autonomy).

Written that way, each choice carries its own controls. The workflow needs per-step evaluation. The router needs misroute testing. Retrieval needs entitlement checks at the source. Voice needs warm transfer and disclosure. L2 needs a reviewer with the time and authority to say no.

IMPLICATION When a vendor describes its product as "an autonomous agent" or "a memory agent", ask for the other four dimensions. The answers tell you what you are buying and which controls you still have to build.

03 / CONTROL PATTERNS

Who decides the next step: code, or the model?

The control pattern is the most consequential choice. The further down this list, the more of the path the model decides, and the harder the agent is to test. Choose the highest row that does the job.

	PATTERN	HOW IT WORKS, AND WHEN	REPRESENTATIVE TOOLS	RISK AND MINIMUM CONTROLS
01	**Rules-based automation** BASELINE, NOT AN AGENT START AT L3	Fixed triggers and rules. No model reasoning, no memory. BEST FOR High-volume, stable, fully specified tasks. AVOID WHEN Inputs are unstructured or the rules keep multiplying.	Zapier, n8n, Power Automate, IFTTT. All now offer AI steps; adding one turns the flow into an agentic workflow. IN THE 100 Routing or acknowledging email. None of the 100 use cases is pure rules.	Silent breakage when an upstream format changes. CONTROLS Change control and run monitoring, as for any integration.
02	**Agentic workflow** MODEL STEPS ON A FIXED PATH START AT L2	Code fixes the sequence; a model performs named steps such as extract, classify or draft. BEST FOR Repeatable processes with unstructured inputs. Where most production value sits today. AVOID WHEN The path genuinely cannot be known in advance.	LangGraph; Microsoft Agent Framework workflows (1.0 GA, April 2026); Gemini Enterprise Agent Designer; n8n and Zapier AI steps. IN THE 100 #4 Commercial submission intake & triage; #3 GenAI claims correspondence drafting	Step-level errors compound along the chain. CONTROLS Evaluation set per step; schema-validated outputs; a person on the consequential step.
03	**Router (triage) agent** SENDS EACH REQUEST TO THE RIGHT PATH START AT L2	Classifies each request and hands it to the right workflow, agent or person. BEST FOR Front doors with many request types. AVOID WHEN There is only one path.	OpenAI Agents SDK handoffs; LangGraph conditional edges; Google ADK. IN THE 100 #84 Correspondence triage for vulnerable claimants	Misroutes stay invisible unless measured. CONTROLS Confusion-matrix testing; a low-confidence default route to a person.

When the model decides the path.

These three patterns let the model choose its next step. They suit open-ended work, and they need the strongest evaluation and the narrowest tool scopes.

	PATTERN	HOW IT WORKS, AND WHEN	REPRESENTATIVE TOOLS	RISK AND MINIMUM CONTROLS
04	**Tool-using agent** SINGLE-AGENT LOOP START AT L1	The model loops: reason, call a tool, read the result, repeat until done. BEST FOR Open-ended tasks across a few systems. AVOID WHEN Tools are unreliable, or a write cannot be undone.	Claude Agent SDK; OpenAI Agents SDK; Google ADK; AWS Strands Agents. Tools exposed through MCP. IN THE 100 #20 Agentic sales assistant for auto finance	Tool misuse; prompt injection arriving in tool results. CONTROLS Least-privilege tool scopes; approval bound to the exact action; idempotent writes; full trace.
05	**Planner-executor** LONG-HORIZON, INCLUDING DEEP RESEARCH START AT L1	Writes a plan, executes it step by step and revises as results arrive. Runs minutes to hours. Includes deep research agents. BEST FOR Multi-step goals with checkpoints: research, compliance packs, migrations. AVOID WHEN A one-step answer will do, or latency and cost matter more than depth.	ChatGPT Work plan mode (launched July 2026); Claude Managed Agents (public beta); OpenAI, Gemini and Claude deep research; LangGraph. IN THE 100 #98 Agentic tax compliance	Drift from the goal; runaway token cost. CONTROLS Plan approval before execution; step and budget limits; resumable checkpoints.
06	**Evaluator-optimizer** DRAFT, CRITIQUE, REVISE START AT L0	One model drafts; a second pass critiques against a rubric; the draft is revised until it passes. BEST FOR Quality-sensitive writing and code where a rubric exists. AVOID WHEN There is no objective rubric, or the critic shares the drafter's blind spots.	Reflexion and Self-Refine (research techniques); LangGraph loops; model-graded evaluators in evaluation tools. IN THE 100 #45 Clinical/regulatory document drafting	Self-agreement mistaken for correctness. CONTROLS Calibrate the critic against human-graded examples; cap iterations.

04 / TOPOLOGY

How many agents, and how they coordinate.

Add agents only when one agent cannot hold the context or the work splits cleanly. Each hand-off loses some context and adds cost, and accountability must stay with a named person however many agents are involved.

	PATTERN	HOW IT WORKS, AND WHEN	REPRESENTATIVE TOOLS	RISK AND MINIMUM CONTROLS
07	**Orchestrator-worker** HIERARCHICAL START AT L1	A lead agent splits the goal, delegates to specialist subagents with their own context and tools, then assembles the result. BEST FOR Broad tasks that divide cleanly into parallel parts. AVOID WHEN Subtasks are tightly coupled, or delegation adds nothing.	Claude Agent SDK subagents; Microsoft Agent Framework Magentic pattern (GA); CrewAI hierarchical process; Google ADK. IN THE 100 #93 M&A due diligence review	Context lost at hand-offs; cost multiplies with each worker. CONTROLS Typed hand-off contracts; no worker holds wider permissions than the lead; one trace across agents.
08	**Peer hand-off team** SPECIALISTS PASS WORK BY ROLE START AT L1	Specialist agents pass work to one another by role, with no single manager. BEST FOR Pipelines of distinct skills: researcher, drafter, reviewer. AVOID WHEN The task is small enough for one agent to hold.	OpenAI Agents SDK handoffs; CrewAI; Microsoft Agent Framework sequential and parallel patterns (AutoGen is now in maintenance mode). IN THE 100 #91 Contract drafting and review	Loops; unclear accountability for the outcome. CONTROLS A named owner for the outcome; termination conditions; a shared trace ID.
09	**Cross-organization agents** AGENT-TO-AGENT ACROSS FIRMS START AT L1	Agents from different vendors or firms discover each other and exchange tasks over a protocol. BEST FOR Partner, supply-chain and payment workflows. AVOID WHEN Counterparty identity and liability are not contracted.	A2A protocol (an Agentic AI Foundation project since August 2026); Agent Payments Protocol (AP2); agent registries and gateways. IN THE 100 Emerging. None of the 100 use cases is yet cross-firm.	Impersonation; data leaving the boundary. CONTROLS Agent identity and signed requests; contracted scopes; data-sharing rules enforced at the gateway.

What you add to an agent.

Retrieval, memory and learning attach to any pattern. They are not agent types on their own, and each brings a data obligation: entitlements for retrieval, retention for memory and release discipline for learning.

	PATTERN	HOW IT WORKS, AND WHEN	REPRESENTATIVE TOOLS	RISK AND MINIMUM CONTROLS
10	**Agentic retrieval** AGENTIC RAG START AT L0	The agent decides what to search, retrieves, checks relevance and searches again before answering with citations. BEST FOR Answers that depend on current or internal documents. AVOID WHEN The corpus is ungoverned, or built-in knowledge is enough.	LlamaIndex; LangChain; Snowflake Cortex Search; Databricks AI Search; Azure AI Search; Gemini Enterprise. IN THE 100 #16 Research Q&A assistant (capital markets); #96 Accounting standards research assistant	Retrieval exposes content the user is not entitled to see; stale sources. CONTROLS Enforce the caller's entitlements at retrieval, not in the prompt; show source dates and citations.
11	**Memory** CONTINUITY ACROSS SESSIONS START AT L0	Stores facts, preferences and past outcomes across sessions and recalls them when relevant. BEST FOR Assistants with ongoing relationships or long-running cases. AVOID WHEN Tasks are one-off, or retention rules forbid storage.	Letta; Mem0; Zep; LangMem; managed memory in AWS Bedrock AgentCore and Claude Managed Agents (research preview). IN THE 100 #15 Advisor meeting notes & follow-up (wealth)	Privacy breach; stale or poisoned memories. CONTROLS Retention and deletion policy; provenance on each memory; users can view and erase.
12	**Learning (optimization)** IMPROVES FROM GRADED FEEDBACK START AT L2	Performance improves from graded feedback, mostly offline: prompt optimization, fine-tuning, refreshed examples. Runtime self-modification is rare in production. BEST FOR High-volume tasks with a clear success signal. AVOID WHEN Behavior must stay stable and auditable between releases.	DSPy; OpenAI reinforcement fine-tuning; evaluation-driven prompt optimization. IN THE 100 #3 GenAI claims correspondence drafting	Drift; reward hacking; untested behavior reaching production. CONTROLS Every learned change ships as a new release, tested and approved like code.

05 / ENVIRONMENT

Where the agent acts.

Prefer an API whenever one exists. Each step from API to browser to computer use is slower, more fragile and harder to secure.

	PATTERN	HOW IT WORKS, AND WHEN	REPRESENTATIVE TOOLS	RISK AND MINIMUM CONTROLS
13	**Coding agent** REPOSITORIES, TESTS, PULL REQUESTS START AT L1	Reads a repository, writes and runs code and tests, and opens changes for review. BEST FOR Software delivery, legacy modernization, data pipelines. AVOID WHEN There are no tests, no sandbox and no reviewer.	Claude Code; OpenAI Codex (now inside the ChatGPT desktop app); GitHub Copilot coding agent; Cursor; Devin. IN THE 100 #17 Legacy code modernization (AI software agents)	Plausible but wrong code; secrets exposure. CONTROLS Sandboxed execution; mandatory review and CI gates; no production credentials.
14	**Browser agent** NAVIGATES SITES AND FORMS START AT L1	Navigates sites, reads pages, fills forms and clicks through flows. BEST FOR Web research, monitoring, portals with no API. AVOID WHEN An API exists, or the pages hold sensitive data under the user's sign-in.	Perplexity Comet and Comet Enterprise (March 2026); ChatGPT Work's built-in browser; Claude in Chrome; Browser Use (open source). ChatGPT Atlas was discontinued on 9 August 2026. IN THE 100 Competitor and regulatory monitoring; portal submissions where no API exists.	Indirect prompt injection from page content. CONTROLS Site allow-lists; signed-out mode by default; confirmation before any submit or purchase.
15	**Computer-use agent** DRIVES DESKTOP APPLICATIONS START AT L1	Sees the screen and drives mouse and keyboard in desktop applications. BEST FOR Legacy desktop software with no API; replacing brittle RPA scripts. AVOID WHEN An API or browser route exists, or the system is high-risk.	Claude computer use tool; computer use inside ChatGPT Work; Gemini computer use; RPA vendors such as UiPath adding agents. IN THE 100 Back-office re-keying between legacy screens.	Slow and error-prone, with a wide blast radius. CONTROLS Dedicated virtual machine and account; session recording; human takeover.

Voice, data and background agents.

These three environments bring their own controls: live conversation, governed numbers and actions taken while no one is watching.

	PATTERN	HOW IT WORKS, AND WHEN	REPRESENTATIVE TOOLS	RISK AND MINIMUM CONTROLS
16	**Voice agent** REAL-TIME CONVERSATION START AT L2	Speech in, speech out, with tool calls, in a live conversation. BEST FOR Contact centers, intake, scheduling, outreach. AVOID WHEN Conversations are complex, emotional or regulated advice.	OpenAI Realtime API; Azure Voice Live API (GA); Gemini Live API; ElevenLabs; contact-center platforms. IN THE 100 #1 Voice AI claim intake (FNOL); #40 Voice AI for patient access and outreach	Errors happen live; containment targets can hurt customers. CONTROLS Warm transfer to a person; recording and QA sampling; disclosure that the caller is speaking with AI.
17	**Data and analytics agent** GOVERNED QUESTIONS OVER A SEMANTIC LAYER START AT L0	Turns a business question into a governed query over a semantic layer, runs it in the database and returns the answer with its definition and date. BEST FOR Self-service metrics; finance and risk questions. AVOID WHEN Definitions are not agreed, or a causal question is asked of descriptive data.	Snowflake Cortex Analyst; Databricks Genie; dbt Semantic Layer; Looker with Gemini; Microsoft Fabric data agents. IN THE 100 #55 Demand forecasting at scale; #68 Demand forecasting and S&OP	Confident wrong numbers from ambiguous definitions. CONTROLS Verified query repository; row and column security enforced in the engine; every answer shows definition, source and date.
18	**Ambient (event-driven) agent** RUNS IN THE BACKGROUND START AT L2	Runs in the background on a trigger or schedule, with no prompt, then reports or acts. BEST FOR Monitoring, queue work, overnight preparation. AVOID WHEN Nobody owns the output.	ChatGPT Work scheduled tasks; Claude Managed Agents; AWS Bedrock AgentCore (GA); Microsoft Foundry Hosted Agents (GA). IN THE 100 #78 Interconnection queue processing	Unseen actions; spend that runs while no one watches. CONTROLS Agent identity in IAM; run logs; budget caps; kill switch.

Autonomy is a setting, not a type of agent.

Any pattern can run at any level. The level should follow the consequence of the action and the evidence that the agent performs it reliably, and it can differ by action within one agent: read freely, draft at L0, write only at L2.

LEVEL	THE AGENT	THE PERSON	USE WHEN	MINIMUM OVERSIGHT
L0 Assist	Suggests, drafts or answers.	Decides and acts.	Any new use case; anything touching people's rights or money.	Output review by the user; usage logging.
L1 Approve each	Proposes every action.	Approves each action before it runs.	Early production; high-consequence actions.	Approval bound to the exact proposal, with expiry.
L2 Approve the consequential	Acts on low-risk steps; stops before money, customer or record changes.	Approves the consequential step.	Most regulated workflows today.	Consequence tiers defined per action; a reviewer with the information, time and authority to say no.
L3 Monitor	Acts within bounds.	Reviews samples, exceptions and metrics.	Proven, reversible, high-volume work.	Sampling plan; alerting; fast rollback.
L4 Bounded autonomy	Runs end to end within contracted limits.	Owns the outcome and reviews performance.	Low-consequence work with strong evidence.	Hard limits on spend, scope and data; kill switch; periodic re-approval.

Regulators are moving the same way. In June 2026 the Financial Stability Board proposed sound practices that call for human approval of high-risk agent actions [2]. The suggested starting levels in sections 03 to 05 assume a regulated firm; raise a level only when the evaluation evidence and the incident record support it.

PRACTICE Write the autonomy level into the agent's manifest for each action, not once for the agent. Agent Onboarding sets out the manifest; AI Governance Evidence sets out what an examiner will ask to see.

What a bank examiner, an insurance regulator and a buyer will ask.

The same three questions decide how hard an agent will be examined: what happens if it is wrong, how much authority it holds, and how much a person still oversees. Regulators now say so explicitly.

Banks: the OCC examiner framework

On 4 November 2026 the OCC issued Bulletin 2026-54 and a supervisory framework that examiners may use to evaluate banks' use of what it now calls super intelligence [6]. It is an examiner guide and states that it does not establish supervisory expectations for banks. It suggests risk classification may rest on potential consequences, transfer of agency and the degree of human oversight rather than technical sophistication, distinguishes a person in the loop, who approves each action, from a person on the loop, who monitors with the authority and practical ability to intervene, and starts from a documented purpose statement. Its four illustrative tiers are below, with the author's reading of where a pattern lands by what it can affect and its autonomy level.

TIER	PROFILE AND OCC EXAMPLES
Tier 1: Critical	Direct action, low or no reversibility, high financial or customer impact, autonomous or partly autonomous. OCC examples: autonomous payment agent, real-time credit decisioning, fraud-hold release.
Tier 2: High	Direct action or high-weight decision support with a person on the loop, moderate to high impact. OCC examples: AI-assisted underwriting, agent-initiated customer communications, AML alert disposition.
Tier 3: Moderate	Decision support with a person in the loop, limited data sensitivity, reversible. OCC examples: internal summarization for credit memos, agent-assisted research.
Tier 4: Low	A person fully in control, no customer or financial data exposure. OCC examples: internal productivity tools, code assistance, general employee questions.

WHAT THE AGENT CAN AFFECT	L0	L1	L2	L3	L4
Internal information only	4	4	4	3	3
Drafts or recommendations a person acts on	3	3	3	2	2
Customer communications or decisions about customers	3	3	2	2	1
Money, credit, coverage or hard-to-reverse records	2	2	2	1	1

Illustrative. Your examiners and your own second line decide the tier.

Insurers, growth companies and the purpose statement.

Insurers answer to state regulators. The NAIC model bulletin, adopted by 24 states and the District of Columbia as of April 2026, expects a written AI Systems program covering governance, risk management and internal controls, and oversight of third-party AI, available on examination [7, 8]. Uses that make or support decisions about consumers need the strongest testing and documentation. Growth companies have no single regulator, but enterprise customers' security reviews, investors and acquirers ask the same questions; the essential controls in section 08 answer most of them.

ORGANIZATION	WHO ASKS	WHO DOES WHAT
Bank	OCC examiners	First line: The business owner and the team that runs the agent; Second line: Risk management, compliance and independent model validation; Third line: Internal audit
Insurer	State insurance regulators	First line: The business owner and the team that runs the agent; Second line: Risk, compliance and actuarial review under the written AI program; Third line: Internal audit
Growth company	Enterprise customers, investors, acquirers, auditors	Owner: The team that builds and runs the agent; Independent check: Someone who did not build it: a CTO, security lead, board member or fractional advisor; Outside assurance: Your SOC 2 or ISO/IEC 42001 auditor, enterprise customers' security reviews, and investor or acquirer diligence

Write the purpose statement first

Whatever the sector, start with a purpose statement, as the OCC framework suggests for banks. It should say what the agent is designed to do and not do, who or what may invoke it, the data and systems it may reach, the actions it may take, and when it needs human review or suspension. Treat activity outside that statement as a control failure even when the result looks reasonable; the framework describes the statement as a basis for identifying scope drift and unauthorized use.

PRACTICE The release pack at rjlsol.com/library/agent-patterns drafts the purpose statement from your design choices and leaves the specifics in brackets for the owner to complete.

Every control needs an owner in each line, and proof that it works.

A control that cannot be evidenced will not survive an examination, an audit or a customer's security review. For each control, the first line runs it, the second line sets the standard and challenges it, and internal audit tests it. Banks and insurers usually call this the three lines of defense; the Institute of Internal Auditors now calls it the Three Lines Model. In a growth company the same roles become the owner, an independent reviewer who did not build the agent, and outside assurance such as a SOC 2 or ISO/IEC 42001 auditor or investor diligence.

CONTROL	FIRST LINE RUNS IT	SECOND LINE CHALLENGES	THIRD LINE TESTS	EVIDENCE THAT PROVES IT
Purpose statement and permitted use *Essential for growth companies*	Business owner writes it and keeps it current	Approves the scope and checks use stays inside it	Tests that actual use matches the approved purpose	Approved, versioned purpose statement, including who may invoke the agent
Named owner and inventory entry *Essential for growth companies*	A named owner accountable for outcomes, with authority to stop the agent	Confirms the owner has the authority and capacity	Checks the inventory is complete and current	AI inventory entry with owner, tier, status and providers
Risk tier and consequence classification	Proposes the tier from consequence, autonomy and oversight	Challenges and sets the tier	Tests tiering consistency across the inventory	Tier record with rationale and approver
Agent identity and least-privilege tool scopes *Essential for growth companies*	Gives the agent its own identity and the narrowest tool scopes	Sets the access standard and approves any elevated scope	Samples actual access against the manifest	Agent manifest (identity, tools, scopes) and periodic access review sign-off
Data entitlements enforced at the source *Essential for growth companies*	Enforces the caller's permissions at the data source or engine, not in the prompt	Defines data classification and permitted uses	Re-performs allowed and denied requests	Entitlement test results with allowed and denied cases
Approved business meaning	Data owners approve the definitions the agent relies on	Checks the definitions used in regulated outputs	Traces an output back to approved definitions	Versioned glossary or ontology release with owner approvals
Evaluation gate before release *Essential for growth companies*	Builds and runs the evaluation set; release is blocked below threshold	Approves thresholds and reviews results for higher tiers	Confirms the gate ran and could not be bypassed	Signed evaluation report per release, with thresholds and results
Independent validation	Supplies documentation and test access	Validates independently of the builders	Reviews validation coverage and independence	Independent validation report and findings log

Controls for action, oversight and recovery.

CONTROL	FIRST LINE RUNS IT	SECOND LINE CHALLENGES	THIRD LINE TESTS	EVIDENCE THAT PROVES IT
Approval bound to the exact action	Reviewers approve the exact proposed action before it runs	Defines which actions need approval	Re-performs a sample of approvals	Approval log: reviewer, exact proposal, time, expiry and outcome
Meaningful human review	Staffs review so each reviewer has the information, time and authority to say no	Monitors override and approval rates for rubber-stamping	Tests whether review is meaningful, not a nominal click	Review volume per reviewer, override rates and reviewer training records
Monitoring, sampling and alerting	Monitors outcomes, samples results and responds to alerts	Sets sampling rates and alert thresholds	Checks alerts were raised and resolved	Monitoring dashboard, sampling log, and alert and incident tickets
Hard limits on spend, scope and rate	Configures spend, scope, rate and data limits	Approves the limits	Tests that a breach is blocked	Configured limits and a breach test record
Kill switch and fallback operation *Essential for growth companies*	Can stop the agent quickly and keep the process running without it	Sets the triggers for suspension	Reviews the drill and the time to stop	Kill-switch drill record with time to stop, and the fallback runbook
Request reconstruction *Essential for growth companies*	Records each request: inputs, sources, tools, outputs and approvals	Defines what must be kept and for how long	Rebuilds one allowed, one denied and one corrected request	Evidence index and request traces
Change control for prompts, models and tools *Essential for growth companies*	Versions prompts, models, tools and learned changes, and ships them as releases	Reviews material changes	Tests that production matches an approved release	Release record listing every component version
Model and tool provider oversight *Essential for growth companies*	Performs due diligence on providers and tracks their model changes	Approves providers and contract terms	Reviews provider oversight	Provider due diligence file and contract terms on data use, notice of model changes and audit

Controls for exposure, fairness and cost.

CONTROL	FIRST LINE RUNS IT	SECOND LINE CHALLENGES	THIRD LINE TESTS	EVIDENCE THAT PROVES IT
Prompt injection and unsafe content testing *Essential for growth companies*	Tests injection and unsafe content through every input path	Sets the test scope	Checks testing covered every path	Red-team and injection test results
Customer disclosure and complaint handling *Essential for growth companies*	Tells customers they are dealing with AI and routes complaints to a person	Reviews disclosures and complaint trends	Samples complaints end to end	Disclosure text, complaint log and resolutions
Fair outcome testing	Tests outcomes across customer groups before and after release	Sets fairness tests and reviews results (fair lending; unfair discrimination)	Re-performs a test	Outcome testing results and remediation
Typed hand-offs between agents	Defines typed hand-offs; no worker holds wider permissions than the lead	Approves cross-agent permissions	Traces one task across every agent	Hand-off schemas, tests and a cross-agent trace
Memory retention and deletion *Essential for growth companies*	Applies retention and deletion; users can view and erase	Sets the retention rules	Tests deletion end to end	Retention configuration and a deletion test
Cost per accepted task	Tracks cost per accepted task against the baseline	Finance reviews it against the business case	Checks benefit claims are supported	Monthly cost-per-accepted-task report

PRACTICE Which controls apply depends on the pattern, the autonomy level and what the agent can affect. The release pack at rjlsol.com/library/agent-patterns selects them for one agent and adds a draft purpose statement. Readiness questions Q01 to Q10 refer to the Production-Readiness Diagnostic; the evidence index follows AI Governance Evidence.

09 / CHOOSING

Choose the simplest pattern that works.

Most failed agent projects chose too much agency too early. Work down this list and stop at the first answer that fits [1].

1. Can fixed rules do it? Use rules-based automation and stop.
2. Is the path known in advance? Use an agentic workflow. Most value sits here, and it is the easiest to test.
3. Does the next step depend on what the agent finds? Use a tool-using agent, or a planner-executor if the task runs long.
4. Does the work split into independent parts? Add an orchestrator with workers. Add agents only when one agent cannot hold the context.
5. Add modules only when needed: retrieval for current documents, memory for continuity, learning for high-volume tuning.
6. Choose the environment in this order: API, then browser, then computer use. Each step down is slower and riskier.
7. Set autonomy from the consequence of the action, then raise it only on evidence.

How this maps to the RJL disciplines

DISCIPLINE	WHAT IT MEANS FOR AGENTS
Trusted Data	Retrieval, memory and data agents are only as good as the entitlements, definitions and provenance beneath them.
Tested Capability	Every pattern needs an evaluation set, and every learned or prompt change ships as a release.
Bounded Access	Tool scopes, agent identity and the autonomy level decide what an agent can touch, and so decide its risk.

NEXT STEP Take one workflow from 100 AI Use Cases and Who Runs Them, describe it on the five dimensions, and score it with the Production-Readiness Diagnostic before choosing a platform. Choosing an Enterprise Platform for AI Agents covers the platform decision.

Status of the main tools at publication.

Tool lists in this guide are representative, not endorsements or a complete market view. Status labels are as published by each vendor on 11 November 2026. Confirm the status of any tool before committing a production design.

TOOL	STATUS AT 11 NOVEMBER 2026	SOURCE
ChatGPT Atlas	Discontinued 9 August 2026. Browser-based agent work moved into ChatGPT and Codex.	OpenAI Help Center
ChatGPT Work	Agent mode launched 9 July 2026 on the same technology as Codex; succeeds the earlier ChatGPT agent.	OpenAI, 9 July 2026
OpenAI Agent Builder	Being wound down; unavailable on the platform from 30 November 2026. OpenAI points code workflows to the Agents SDK.	OpenAI AgentKit page
Perplexity Comet	Available; Comet Enterprise, with admin and security controls, launched 17 March 2026.	TestingCatalog, March 2026
Microsoft Agent Framework	1.0 GA in April 2026 as the successor to Semantic Kernel and AutoGen; Agent Harness and Foundry Hosted Agents since reached GA.	InfoQ, August 2026
Google ADK	Code-first builder inside the Gemini Enterprise Agent Platform, which replaced Vertex AI as Google's agent platform in April 2026.	PYMNTS, April 2026
AWS Bedrock AgentCore	Generally available; runs agents built with CrewAI, ADK, LangGraph, LlamaIndex, the OpenAI Agents SDK and Strands.	Constellation Research
Claude Agent SDK and Managed Agents	Agent SDK available as a library; Managed Agents, the hosted runtime, in beta since 8 April 2026.	Claude Platform docs
MCP and A2A	Both under the Linux Foundation's Agentic AI Foundation: MCP since 9 December 2025, A2A since 17 August 2026.	Agentic AI Foundation

A caution on voice: Commonwealth Bank of Australia reversed 45 contact-center job cuts tied to a voice bot in August 2025 after call volumes rose [3]. Containment targets are not the same as resolved customers.

Sources and notes.

Prepared by Richard White. Version 1.1, November 2026. Tool statuses come from the public sources in section 10 and were not confirmed with the vendors. Use-case examples are the author's classification of each use case; they are not claims about how the named organization built its system. Proposed guidance; not legal advice.

[1] Anthropic, Building effective agents, December 2024. anthropic.com

[2] Reuters via The Star, Global watchdog calls for tighter controls on agentic AI in finance, 10 June 2026. thestar.com.my

[3] Finextra, CommBank reverses plan to replace call centre staff with AI, August 2025. finextra.com

[4] Visual Studio Magazine, Microsoft ships production-ready Agent Framework 1.0, 6 April 2026. visualstudiomagazine.com

[5] Blck Alpaca, What the Linux Foundation and AAIF mean for MCP and agentic AI, October 2026. blckalpaca.at

[6] Office of the Comptroller of the Currency, Bulletin 2026-54, Super Intelligence: Supervisory Framework, 4 November 2026. occ.gov

[7] NAIC, Model Bulletin: Use of Artificial Intelligence Systems by Insurers, adopted 4 December 2023. content.naic.org

[8] Mayer Brown, NAIC Spring 2026 National Meeting highlights (adoption count), April 2026. mayerbrown.com

[9] RJL Solutions, 100 AI Use Cases and Who Runs Them, Version 1, 10 November 2026. Use-case numbers in this guide refer to it. rjlsol.com

ABOUT

Richard White

I was global firm chief data officer at McKinsey & Company, owning data, AI and platforms end to end, and before that chief data officer at New York Life, where I built the chief data office and reported data risk to risk committees at multiple levels. Earlier I led data at the Federal Reserve Bank of New York and enterprise data governance at Citigroup. I founded RJL Solutions in 2024. My book, *The Day After the Demo*, is free at rjlsol.com/book.

RJL Solutions connects strategy, architecture, engineering, governance and operation through Frame, Prove, Build, then Operate and transfer. Each phase is separately scoped and priced in a written proposal before you commit. The first sixty-minute working session is free.

USE CASES **100 AI Use Cases and Who Runs Them**
Ten industries, named organizations, public sources.

PLATFORMS **Choosing an Enterprise Platform for AI Agents**
Assistant, models, model access and agent platform.

RELEASE **Agent Onboarding**
Treating an agent like a data product, with a manifest and scorecard.

RJL SOLUTIONS

WEB rjlsol.com

EMAIL rwhite@rjlsol.com

Tool statuses checked against public sources on 11 November 2026 and not confirmed with vendors.
Tool lists are representative, not endorsements.