Growth Performance

AI + LEADERSHIP TOOL

Agentic AI Blueprint Planner

A structured planning template for leaders designing human-autonomous collaboration. Use this to map your AI agents to business outcomes, define decision-making scope, establish governance boundaries, and prevent agent sprawl before it starts.

What this planner is for

Agentic AI systems, those that can plan, execute, and adapt across multi-step tasks without constant human instruction, are being deployed faster than most organisations have designed governance for. This blueprint helps you define the scope, oversight, and handoff design for each agent in your area before deployment, not after the problems emerge. Based on McKinsey (2026), Deloitte (2026), and EY (2025) research on the human factors of agentic AI.

LEADER / OWNER NAME

FUNCTION / TEAM

DATE

SECTION 1: AGENT INVENTORY

List every AI agent currently in use or under consideration in your area. Include both purpose-built agents and general-purpose tools being used in agentic ways (automated pipelines, scheduled AI tasks, multi-step workflows). Be exhaustive, agent sprawl typically begins with undocumented deployments.

#	AGENT NAME / TOOL	PRIMARY FUNCTION	BUSINESS OUTCOME IT SERVES	STATUS (LIVE / PILOT / PLANNED)
1				
2				
3				
4				
5				
6				

AGENTS I SUSPECT EXIST BUT CANNOT YET CONFIRM

SECTION 2: DECISION SCOPE MATRIX

For each agent, define precisely what decisions it can make autonomously, which require human approval before action, and which must always escalate to a named human. Clarity here prevents the gradual expansion of autonomous authority that typically goes unnoticed until something goes wrong.

#	AGENT NAME	DECIDES ALONE (NO APPROVAL NEEDED)	REQUIRES HUMAN APPROVAL BEFORE ACTING	ALWAYS ESCALATES TO HUMAN
1				
2				
3				
4				
5				

Key question: For every decision marked "decides alone"are you certain that the consequences of an error are within acceptable bounds without human review? If not, move it to the approval column.

SECTION 3: GOVERNANCE BOUNDARIES

Governance is not a constraint on AI performance, it is the condition under which you can trust AI performance. Define the oversight intervals, escalation triggers, and audit mechanisms for each agent before it operates at scale.

Escalation triggers (define at least three per active agent)

An escalation trigger is a condition under which the agent must stop acting autonomously and route the situation to a named human. Examples: confidence below threshold, financial impact above limit, action affecting a customer record, output flagged by another system.

#	AGENT NAME	ESCALATION TRIGGER CONDITION	ESCALATES TO (NAMED ROLE / PERSON)
1			
2			
3			
4			
5			
6			

OVERSIGHT REVIEW INTERVAL (HOW OFTEN AGENT OUTPUTS ARE REVIEWED)

AUDIT TRAIL LOCATION (WHERE AGENT DECISIONS ARE LOGGED)

SECTION 4: HUMAN-AGENT HANDOFF DESIGN

Every point where an agent hands work back to a human, or where a human must review, approve, or correct agent output, is a handoff point. Poor handoff design is the most common source of error in human-autonomous collaboration. It creates invisible gaps where nobody is checking and nobody knows who is accountable.

#	AGENT NAME	HANDOFF POINT (WHAT TRIGGERS THE RETURN TO HUMAN)	HUMAN RESPONSIBLE AT HANDOFF	HANDOFF PROTOCOL (HOW THE TRANSFER IS SIGNALLED)
1				

#	AGENT NAME	HANDOFF POINT (WHAT TRIGGERS THE RETURN TO HUMAN)	HUMAN RESPONSIBLE AT HANDOFF	HANDOFF PROTOCOL (HOW THE TRANSFER IS SIGNALLED)
2				
3				
4				
5				

GAPS IN HANDOFF DESIGN I NEED TO ADDRESS BEFORE SCALING

SECTION 5: AGENT SPRAWL PREVENTION

Agent sprawl occurs when the number and scope of active agents grows faster than governance can track. It is one of the most significant risks in agentic AI deployment, creating overlapping functions, undefined accountability, and escalating operational complexity. Use this section to assess and constrain sprawl before it becomes structural.

SPRAWL RISK AUDIT, CHECK ALL THAT APPLY

- [ ] Two or more agents perform overlapping functions with no documented rationale for both existing.
- [ ] At least one agent in my inventory was deployed by a team member without my awareness.
- [ ] I cannot name the human accountable for the outputs of at least one active agent.
- [ ] At least one agent does not have a defined sunset or review date.
- [ ] Agent outputs from at least one deployment are not being reviewed at a regular interval.
- [ ] The total number of active agents in my area has grown in the last 90 days without a deliberate decision to expand.

AGENTS TO CONSOLIDATE OR RETIRE (NAME THEM)

SUNSET OR REVIEW DATE FOR EACH ACTIVE AGENT

PROPOSED POLICY FOR NEW AGENT DEPLOYMENT APPROVAL IN MY AREA

SECTION 6: BLUEPRINT REVIEW RHYTHM

An agentic AI blueprint is not a document, it is a living governance commitment. Define how and when you will review it, who attends, and what will trigger an unscheduled review.

SCHEDULED REVIEW FREQUENCY

STANDING REVIEW ATTENDEES (ROLES)

TRIGGER CONDITIONS FOR AN UNSCHEDULED REVIEW (E.G. NEW AGENT DEPLOYMENT, ESCALATION EVENT, SIGNIFICANT ERROR)

ONE THING I WILL CHANGE OR CLARIFY IN THIS BLUEPRINT BEFORE ITS FIRST USE

Based on: McKinsey & Company (2026) Six Shifts to Build the Agentic Organization of the Future; Deloitte (2026) Global Human Capital Trends; EY (2025) EY Survey Reveals Majority of Workers Are Enthusiastic About Agentic AI; Buell, R. and Kagan, J. (2026) What Leadership Looks Like in an Agentic AI World. Growth Performance © 2026.

Want help putting this into practice?

Growth Performance builds the human advantage in an AI-driven world, with people, teams and leaders.

growthperformance.co.uk/contact
hello@growthperformance.co.uk