Agentic Enterprise Stack: 90-Day Control Checklist

Use this before allowing AI agents to move from assistant mode into supervised or autonomous execution. The goal is not more agents. The goal is verified outcomes per controlled agent.

Layer	Minimum control before production	Owner
1. Identity + Work Graph	Agent has a unique non-human identity, named business owner, technical owner, workflow trigger, action boundary, and exit criteria.	Business + IT
2. Enterprise Context	Approved sources are mapped, source priority is defined, retrieval respects permissions, and stale/conflicting data has an escalation rule	Data + Process Owners
3. Tool Contracts + MCP	Every tool call has schema, risk tier, allowed action class, error handling, rate limit, and rollback rule.	IT + Security
4. Orchestration Runtime	Workflow can pause, resume, retry, time out, escalate, and preserve task state without relying on a chat transcript.	Platform Owner
5. Memory + State	Task memory, user preference memory, and organizational learning are separated with retention and privacy rules.	Data + Legal
6. Observability + Evals	Offline eval set exists; production dashboard tracks tool calls, cost, latency, override rate, policy failures, and business outcomes.	AI lead + Ops
7. Governance + Security	Approval gates, audit export, incident playbook, kill switch, and retirement criteria are documented and tested.	Risk + CISO

90-Day Rollout Gate

Days	Decision gate
1-15	Workflow mapped. If the team cannot draw the work graph, do not build the agent.
16-30	Control baseline complete: identity, permissions, data sources, tool contracts, eval set, escalation rules.
31-50	Shadow mode on real cases. Compare agent outputs to human decisions and measure failures.
51-70	Supervised execution only on low-risk actions with logging, approval, and rollback.
71-90	Expand autonomy only where evidence supports it. Retire or pause weak agents.

Board Question

Can the company produce a list of every production agent within one hour, including owner, purpose, permissions, connected systems, action tiers, latest eval score, monthly cost, incident history, and business outcomes?