



A.G.E.N.T.S.

Enterprise AI Agent Implementation Package


From business opportunity to governed production deployment



USAII x Jim Markunas

Built around Jim Markunas' A.G.E.N.T.S. Enterprise
Agent Operating Model

A Authority	G Guardrails	E Evidence	N Network & Integrations	T Transfer & Escalation	S Success & Accountability
-----------------------	------------------------	----------------------	---------------------------------------	--------------------------------------	---



Find the right opportunity. Define the job. Govern the agent. Prove it’s ready.

USAII x Jim Markunas

A.G.E.N.T.S.

Six questions every enterprise AI agent needs answered before it can be trusted with real work.

A chatbot can provide information. An enterprise agent can **make decisions, use tools, change systems, trigger workflows, and create business consequences.**

That changes the product problem. The team must decide what authority the agent receives, what constrains it, what evidence it leaves behind, what systems it may access, when a person takes over, and who owns the result. That is the purpose of the A.G.E.N.T.S. Enterprise Agent Operating Model.

A AUTHORITY What may it do? Define what the agent may observe, recommend, prepare, decide, and execute. Authority should be assigned action by action - not granted as a blanket permission.	G GUARDRAILS What must it never do? Define prohibited actions, approval thresholds, financial or customer-impact limits, policy constraints, and the conditions that require the agent to stop.
E EVIDENCE How will we know what it did? Preserve enough evidence to reconstruct what the agent knew, what rules applied, what it decided, what action it took, and what happened afterward.	N NETWORK & INTEGRATIONS What systems may it touch? Define the systems, data, APIs, identities, permissions, and specific capabilities the agent needs - without giving it broader access than necessary.
T TRANSFER & ESCALATION When does a person take over? Define when the agent must stop, who receives the work, and what context, evidence, and attempted actions travel with the handoff.	S SUCCESS & ACCOUNTABILITY Did it create value - and who owns the result? Measure the business outcome the agent was created to improve while also monitoring quality and risk. Assign clear ownership for the result.

THE A.G.E.N.T.S. TEST

Before an enterprise agent reaches production, the team should be able to answer all six questions clearly:

1.

What may it do?

2.

What must it never do?

3.

How will we know what it did?

4.

What systems may it touch?

5.

When does a person take over?

6.

Who owns the result?

If any of those answers are unclear, the operating model is incomplete.

NEXT · Use the implementation process on Page 3 to take one real agent opportunity from business problem to production decision.

A.G.E.N.T.S. ENTERPRISE AGENT OPERATING MODEL · 03

From Opportunity to Production

Use one real business problem. Work in order.

This package takes a potential AI-agent use case from “**we think this could help**” to a clear implementation decision. Start with the work, not the model. Use real processes, systems, constraints, and owners. If something is unknown, write **UNKNOWN** rather than guessing.

FIND FIND THE RIGHT OPPORTUNITY Understand the environment and decide whether the problem is worth solving, and whether an agent is the right tool.	PAGE 4 Systems & Tools Inventory Map the systems, sources of truth, permissions, constraints, and manual work. PAGE 5 A.G.E.N.T.S. Value Test Test whether the opportunity can make money, save money, or make work easier. PAGE 6 Current State → Agent-Enabled Process Show how the work happens today and what would actually change. PAGE 7 Should This Actually Be an Agent? Choose agent, automation, copilot, or stop.
DEFINE DEFINE THE JOB Turn the opportunity into a concrete operating assignment before deciding the technology.	PAGE 8 A.G.E.N.T.S. Agent Charter Define the objective, trigger, information, decisions, actions, systems, boundaries, and owner.
GOVERN DESIGN THE BOUNDARIES Run the proposed agent through the six parts of A.G.E.N.T.S.: Authority, Guardrails, Evidence, Network & Integrations, Transfer & Escalation, and Success & Accountability.	PAGE 9–14 A · G · E · N · T · S Answer the six operating questions that turn AI capability into an operating model.
PROVE PROVE READINESS Test whether those boundaries are actually defined, implemented, testable, measurable, and owned.	PAGE 15–16 Production Readiness + Final Decision Finish with a documented GO · GO WITH CONDITIONS · NO GO decision.
WORKED EXAMPLE Customer Order Exception Agent	A fictional enterprise agent that investigates delayed or failed customer orders across customer, order, inventory, and payment systems. It may resolve defined low-risk exceptions automatically and prepare or escalate higher-risk cases for human review. The same example appears throughout the package so each decision can be seen in context.

ONE RULE FOR THE ENTIRE PACKAGE If you can’t explain why the agent should act, what it’s allowed to do, and who owns the consequence, it’s not ready for production.

Systems & Tools Inventory

Map the environment before you design the agent.

Inventory the entire enterprise: your systems, tools, data sources, and manual work. The goal is to understand where information lives, which systems are authoritative, what access the agent would need, and where it could create a customer, financial, or operational consequence.

01 INFORMATION Where the relevant facts live.	02 AUTHORITY Which system is authoritative.	03 ACCESS What the agent must be able to read or do.	04 CONSEQUENCE Where an action changes the real world.
--	--	---	---

ANSWER THESE BEFORE YOU COMPLETE THE INVENTORY

01

If two systems disagree, which one wins?

Name the source of truth by decision type. One system does not have to be authoritative for everything.

02

What information currently exists only in a person’s head, inbox, spreadsheet, or manual workaround?

Capture the knowledge the process depends on but the agent cannot reliably access today.

03

Where would the agent need permission to create a customer, financial, or operational consequence?

Examples: refunding money, changing an order, sending a customer message, updating a record, or triggering downstream work.

WORKED EXAMPLE

Customer Order Exception Agent

Salesforce: customer context

SAP: order source of truth

Stripe: payment source of truth

When complete: Enterprise Systems Map → feeds the Agent Charter and Network & Integrations.

Systems Inventory Worksheet

Complete one row for each system or tool involved in the target process.

[illegible]

ENTERPRISE SYSTEMS MAP · Use the completed inventory when defining the Agent Charter and the agent's Network & Integrations.

AI Opportunity Inventory

Start with business ideas, not AI features.

Use this page to capture processes, ideas, initiatives, or new tools that business teams believe AI could improve. **Do not solve the technical design yet.** Describe the desired outcome in plain language, then mark every business-value category the idea could materially improve.

HOW TO USE THIS PAGE

- 01

Add one idea per row.
Use a process, initiative, problem, or proposed tool.
- 02

Say what it should do.
Write the business outcome in plain language.
- 03

Mark the value case.
Place an X in every category the idea could materially improve.

WHAT THE THREE VALUE COLUMNS MEAN

INCREASE ROI Creates, protects, or accelerates revenue or return.	DECREASE COST Reduces spend, labor, rework, or operating expense.	OPERATIONAL EFFICIENCY Reduces time, handoffs, friction, or effort.
---	---	---

One strong value case can be enough to investigate further. Don't force an X into every column.

Initiative / Idea / New Tool	What Should It Do?	Increase ROI	Decrease Cost	Operational Efficiency
		X		

OUTPUT AI Opportunity Inventory · Carry the strongest candidates forward for deeper evaluation.

Opportunity Shortlist

Choose which ideas deserve a closer look.

Take the strongest ideas from the AI Opportunity Inventory and compare them at a high level. **This is not a business case.** The goal is to decide which opportunities are worth investigating further.

HOW TO USE THIS PAGE

BUSINESS VALUE How meaningful is the expected impact on ROI, cost, or operational efficiency?	FEASIBILITY How realistic does delivery look based on the systems, access, data, and ownership we know today?	PRIORITY How important is this opportunity relative to the other ideas on the list?
---	---	---

Score each idea **High/Medium/Low** for the three signals, then decide whether it should move forward.

Initiative, Idea, or New Tool	Business Value H / M / L	Feasibility H / M / L	Priority H / M / L	Move Forward? ☐ Yes ☐ No
				☐ Yes ☐ No
				☐ Yes ☐ No
				☐ Yes ☐ No
				☐ Yes ☐ No
				☐ Yes ☐ No
				☐ Yes ☐ No
				☐ Yes ☐ No
				☐ Yes ☐ No
				☐ Yes ☐ No

SELECTED OPPORTUNITY FOR THE REST OF THIS PACKAGE

OUTPUT Opportunity Shortlist · Carry the selected opportunity forward to map how the work happens today.

Current-State Process Map

What happens today?

Before deciding what an agent should change, document how the selected work currently happens. Capture the major steps, who performs them, which systems they use, and where the process creates friction.

HOW TO USE THIS PAGE

- 1

Start with the event that triggers the process.
- 2

Add the major steps in the order they happen.
- 3

Identify who performs each step and the primary system or tool used.
- 4

Note where the work is slow, manual, repetitive, error-prone, or dependent on handoffs.

Do not redesign the process on this page. Capture the current state as it exists.

Step	What Happens?	Who Does It?	System/Tool	Friction/Problem
1				
2				
3				
4				
5				
6				
7				
8				

WHERE IS THE BIGGEST SOURCE OF DELAY, COST, MANUAL EFFORT, OR FRUSTRATION?

OUTPUT Current-State Process Map · Next: decide what the agent should change, remove, accelerate, or leave with a human.

Agent-Enabled Process Map

What should change?

Now redesign the same process around the selected opportunity. Decide where an agent should take over work, where a human should remain responsible, which systems are involved, and what changes compared with today.

HOW TO USE THIS PAGE

- 1 Use the same process steps from the Current-State Process Map where possible.

3 Keep human ownership where judgment, approval, or accountability is still required.
- 2 Describe the work the agent should perform at each step.

4 State what actually changes: remove a handoff, shorten a delay, automate an action, or leave the step unchanged.

Do not assign permissions or guardrails yet. This page defines the proposed operating flow.

Step	What Should the Agent Do?	What Should a Human Do?	System/Tool	What Changes?
1				
2				
3				
4				
5				
6				
7				
8				

WHAT IS THE MOST IMPORTANT CHANGE THIS AGENT WOULD MAKE TO THE PROCESS?

OUTPUT Agent-Enabled Process Map · Next: decide whether this opportunity should actually be implemented as an agent.

Should This Actually Be an Agent?

Does this work require agentic behavior?

Not every AI opportunity needs an agent. Use an agent when the work has to interpret changing context, choose between possible actions, use multiple systems, handle exceptions, or take action. If the process is predictable and follows fixed rules, a simpler workflow, automation, or copilot may be better.

THIS PROBABLY NEEDS AN AGENT IF IT MUST:

- ☐ Interpret changing context
- ☐ Gather information from multiple systems
- ☐ Choose between different possible actions
- ☐ Adapt when conditions change
- ☐ Handle exceptions that cannot all be scripted
- ☐ Take action—not just provide information
- ☐ Stop or escalate at the limits of its authority

THIS MAY NOT NEED AN AGENT IF:

- ☐ The workflow is fixed and predictable
- ☐ The same rules apply every time
- ☐ Conventional automation could complete the task
- ☐ The main problem is search or information access
- ☐ A human only needs recommendations, not autonomous action

WHAT SHOULD WE BUILD?

Implementation Option	Best Fit	Select
Enterprise AI Agent	Context changes, multiple systems/actions, judgment and exceptions required	☐
Human-in-the-Loop Agent	Agent can do most of the work, but consequential actions require approval	☐
Copilot / Assistant	Human remains responsible for decisions and execution	☐
Rules-Based Automation	Process is predictable and can be expressed with fixed logic	☐
Do Not Pursue Yet	Value, readiness, or problem definition is not strong enough	☐

WHY IS THIS THE RIGHT IMPLEMENTATION PATH?

OUTPUT Implementation Approach · Next: define exactly what the selected agent is being asked to do in the A.G.E.N.T.S. Agent Charter.

A.G.E.N.T.S. Agent Charter

Define the job before defining the technology.

Turn the selected opportunity into a clear operating assignment. Describe what the agent is expected to accomplish, what starts the work, what information and systems it needs, what decisions and actions are in scope, and who owns the result. Keep the answers in plain business language.

COMPLETE THE CHARTER

Define	Answer
Agent Name	
Business Objective	
Primary Value Case	
Trigger	
Who / What It Serves	
Information It Needs	
Decisions It Makes	
Actions It Takes	
Systems It Touches	
Expected Outcome	
What It Must Never Do	
When It Escalates	
Accountable Owner	

REVIEWED BY

☐ Business Owner

☐ Product Owner

☐ Technology / AI Owner

☐ Operations Owner

OUTPUT Agent Charter · Next: use the charter to define the agent’s Authority.

A: Authority

What may the agent do?

Authority is a ladder, not a toggle. An agent can have different levels of authority for different actions. List the important actions from the Agent Charter and mark the highest level the agent should receive for each one.

AUTHORITY LADDER

OBSERVE Read information and monitor state	RECOMMEND Propose an action for a human	PREPARE Assemble or stage the action	DECIDE Select the action or outcome	EXECUTE Carry out the action in the system
--	---	--	---	--

AUTHORITY MATRIX

Agent Action	Observe	Recommend	Prepare	Decide	Execute	If Wrong, What Happens?
	☐	☐	☐	☐	☐	
	☐	☐	☐	☐	☐	
	☐	☐	☐	☐	☐	
	☐	☐	☐	☐	☐	
	☐	☐	☐	☐	☐	
	☐	☐	☐	☐	☐	
HIGHEST-CONSEQUENCE ACTION What is the highest-consequence action the agent may take?						
HUMAN INVOLVEMENT Where must a human remain involved?						

WORKED EXAMPLE Refund ≤ \$50 → Execute · Refund > \$50 → Prepare / Human Approval

OUTPUT Authority Matrix · Next: define the limits that constrain that authority in G: Guardrails.

G: Guardrails

What must constrain the agent?

Authority defines what the agent may do. Guardrails define the boundaries around that authority: what it may do on its own, what requires approval, and what it must never do.

OPERATING BOUNDARIES

MAY DO	APPROVAL REQUIRED	MUST NEVER DO

HARD LIMITS

Financial / Transaction Limit	Customer-Impact Limit	Privacy / Data Boundary	Other Hard Stop

WHAT CONDITION SHOULD CAUSE THE AGENT TO STOP IMMEDIATELY RATHER THAN CONTINUE OR GUESS?

WORKED EXAMPLE Refunds over \$50 require approval; never refund a payment already reversed or act on conflicting payment status.

OUTPUT Guardrail Set · Next: define what evidence must be preserved in E: Evidence.

E: Evidence

How will we know what it did?

A consequential agent should leave enough evidence to reconstruct what it knew, what constraints applied, what it decided, what action it took, and what happened afterward.

EVIDENCE CHAIN

SOURCE | STATE | CONSTRAINT | DECISION | ACTION | OVERRIDE | OUTCOME

WHAT MUST BE PRESERVED?

Evidence	What Must Be Recorded?	Where Will It Live?
Source		
State		
Constraint		
Decision		
Action		
Human Override		
Outcome		

SIX-MONTH TEST

Six months from now, could someone reconstruct why the agent took a consequential action? ☐ Yes ☐ Partially ☐ No
If not, what evidence is missing?

OUTPUT Evidence & Audit Requirements · Next: define the systems and capabilities the agent may use in N: Network & Integrations.

N: Network & Integrations

What systems may it touch?

Give the agent the specific capabilities it needs to do its job - not blanket access. Define what it may read, what it may write, how it authenticates, and what should happen when a dependency is unavailable.

AGENT CAPABILITY MAP

System / Tool	What Does the Agent Need to Do?	Read	Write	Identity / Access	If Unavailable
		☐	☐		
		☐	☐		
		☐	☐		
		☐	☐		
		☐	☐		
		☐	☐		

FAILURE & CONFLICT RULES

If two systems disagree, which source wins - and when should the agent stop?	
If an action only partially succeeds, how will the agent detect, reconcile, or reverse it?	
What should the agent do when a required source of truth is unavailable?	

PRINCIPLE Give the agent capabilities, not blanket access.

OUTPUT Agent Capability Map · Next: define when the agent stops, where work goes, and what travels with it in T: Transfer & Escalation.

T: Transfer & Escalation

What happens when the agent cannot continue?

Escalation is a routing problem. Define what causes the agent to stop, what it must do next, who receives the work, how quickly they need it, and whether the agent may resume afterward.

ESCALATION ROUTING MATRIX

Escalation Scenario	What Must the Agent Do?	Escalates To	Urgency / SLA	Can Agent Resume?
Outside authority				
Approval required				
Conflicting information				
Required system unavailable				
Policy / rule unclear				
Unexpected or high-risk condition				

WHAT MUST TRAVEL WITH THE CASE?

- ☐ Objective / requested outcome
- ☐ Systems checked
- ☐ Actions already attempted
- ☐ Specific decision or approval needed
- ☐ Current state and relevant context
- ☐ Evidence considered
- ☐ Reason the agent stopped
- ☐ Any deadline or customer-impact timing

HUMAN HANDOFF OWNER

Primary destination: _____ Backup destination: _____

WORKED EXAMPLE Refund > \$50 → prepare refund, route to Customer Operations Lead, same queue cycle, resume after approval.

PRINCIPLE A human should inherit the work - not restart it.

OUTPUT Transfer & Escalation Plan · Next: measure whether the agent created the business value selected earlier in S: Success & Accountability.

S: Success & Accountability

Did the agent create the value we expected?

Measure the same business-value buckets used earlier in the package. Complete the rows that apply to this initiative and compare the result before and after implementation.

SUCCESS MEASUREMENT

Success Metric	Before	After	Improvement
Increase ROI / Revenue			
Decrease Cost			
Operational Efficiency / Streamlining			

ACCOUNTABILITY

Business Owner: _____

Agent / Product Owner: _____

PRINCIPLE Execution can be delegated. Accountability cannot.

OUTPUT Success & Accountability Record · This closes the A.G.E.N.T.S. operating model.

Production Readiness Gate

Is the agent ready for real consequences?

Use the completed A.G.E.N.T.S. pages as the final production check. If an area is not ready, record the issue or condition before approving launch.

A.G.E.N.T.S. PRODUCTION CHECK

Area	Ready?	Issue / Condition
A — Authority Authority is defined action by action.	☐ Yes ☐ No	
G — Guardrails Limits, approvals, and prohibited actions are defined.	☐ Yes ☐ No	
E — Evidence Consequential actions can be reconstructed.	☐ Yes ☐ No	
N — Network & Integrations Required access and failure behavior are tested.	☐ Yes ☐ No	
T — Transfer & Escalation Escalation routing and handoff work end-to-end.	☐ Yes ☐ No	
S — Success & Accountability Business-value measures and owners are defined.	☐ Yes ☐ No	

FINAL DECISION

☐ GO

☐ GO WITH CONDITIONS

☐ NO GO

CONDITIONS / UNRESOLVED ISSUES

Decision Owner: _____ Target Launch / Review Date: _____

FINAL RULE If any A.G.E.N.T.S. control is unclear, the agent is not ready.

OUTPUT Production Decision · GO · GO WITH CONDITIONS · NO GO

Run the Agent Through A.G.E.N.T.S.

Six questions to use before an enterprise AI agent is trusted with real work.

A	Authority	What may it do?
G	Guardrails	What must constrain it?
E	Evidence	How will we know what it did?
N	Network & Integrations	What systems may it touch?
T	Transfer & Escalation	When does a person take over?
S	Success & Accountability	Did it create value — and who owns the result?

IF ANY ANSWER IS UNCLEAR, THE AGENT IS NOT READY.

USAI x Jim Markunas

Jim Markunas' A.G.E.N.T.S. Enterprise Agent Operating Model