

# HOW TO **ENABLE USERS** AND **HAND OVER AI** **AGENT OPERATIONS** TO THE BUSINESS

AI Agent Deployment, Integration & Performance Management

## WHERE THIS IS USED

- AI Studio agent deployments
- Venture Studio programs with AI-enabled ventures
- Corporate Incubators
- Foundry-as-a-Service engagements
- CVC portfolio companies with AI product components

## AUDIENCE

- AI Studio Agent Leads
- BU Operations Leads receiving the agent
- BU Managers and end users
- Venture Studio Program Managers
- EIRs / GMs
- Venture Board Team Sponsors

## PHASE

Phase Three: Build and Launch → User Enablement & Operational Handover (Weeks 10-14, concurrent with G3)

Believe in better.®

ICDeaux

# EXECUTIVE SUMMARY

Most AI deployments that fail in corporate organizations do not fail because the technology stopped working. They fail because the people who were supposed to use the technology – the business unit staff, the operations managers, the daily users – were never properly equipped to adopt it, own it, or recover when something went wrong.

The AI Studio team built the agent. The Tech Lead deployed it. The agent is stable in production. None of that means the business is ready to operate it. Readiness to operate is different from readiness to receive: it requires users who understand what the agent does and does not do, an Operations Owner who can resolve day-to-day issues without calling the AI Studio team, a Human-in-the-Loop protocol that defines when human judgement overrides AI output, and a formal handover that transfers accountability – not just access – from the venture team to the business.

This guide runs in parallel with Guide G3 from the first day of stable production. Change management planning begins in G1 (before go-live). User onboarding begins in the first week of production. The transition period – during which the AI Studio team and the Operations Owner operate the agent together – lasts 2–4 weeks. Handover acceptance is a documented sign-off against binary criteria. It is not a calendar date.



# THE CORE PROBLEM


The Venture Studio's job is to build and validate. The business unit's job is to operate and scale. The gap between these two mandates — the transition from build-mode to operate-mode — is where most AI deployments in corporate organizations fail.

## The failure patterns in the human side of AI deployment are consistent:

- The agent is deployed to users who were told it was coming but never trained on how to use it, what to do when it escalates, or how to signal when it produces a wrong output. Usage falls to near zero within 30 days.
- The Operations Owner is named on a handover document but was never involved in the transition period. Their first real encounter with a production issue is their first day operating the agent alone. They call the AI Studio team. The AI Studio team has moved on to the next sprint.
- The Human-in-the-Loop protocol exists in the technical specification but was never explained to users. When the agent escalates a call, the user does not know what to do. The escalation sits unresolved.
- Middle management resists adoption. The BU manager's team is measured on throughput metrics. The agent changes how work gets done. The manager sees it as a threat and quietly discourages use — not loudly, just passively. No one escalates because no one is watching.
- The handover is treated as a date on a project plan rather than a readiness milestone. The AI Studio team hands over on the planned date. The Operations Owner is not ready. Neither party says so because the project plan says the handover is done.

## The underlying cause:

Technology handover and human enablement are treated as the same event. They are not. Technology handover is the transfer of access and configuration. Human enablement is the transfer of understanding, confidence, and accountability. Both are required for a sustainable deployment. This guide addresses the human side.



# PREREQUISITES

## Must Be Complete Before Starting:

- Guide G1 completed – Go-Live Report filed, agent confirmed stable in production, first 48-hour monitoring window passed with no rollback triggers
- Guide E2 confirmed – Venture Board Team Sponsor is active, weekly health check cadence is running, antibody resistance vectors have been mapped
- Guide E3 confirmed – BU Operations Lead has been identified and is engaged; the BU relationship is formalized through the pilot commitment framework

## Readiness Checks Before Starting:

- BU Operations Owner named – a specific individual, not a team. This person will own the agent in production after handover. For BUs receiving multiple agents or a multi-agent workflow, the Operations Owner covers the full portfolio – confirm their capacity before proceeding.
- BU end users identified – the people who will interact with the agent daily. For larger BUs: a representative cohort of 5–10 users for the initial training sessions.
- Executive mandate confirmed with the BU head – the BU head has communicated to their team that the agent is being adopted, not piloted. Ambiguity about permanence is the primary driver of passive resistance.
- Baseline AI literacy confirmed – users have received or will receive foundational AI literacy as part of this program: what AI agents can and cannot do, how hallucinations occur, and what data must never be submitted to an AI tool. This does not require a dedicated course – it can be covered in 15 minutes within the onboarding session. But it must be covered explicitly.



# EXPECTED OUTPUT/ SUCCESS CRITERIA

You have completed this guide when the following are true:

- ✓ Change Management Plan completed – resistance vectors mapped, BU head mandate confirmed, communication plan sent
- ✓ All identified BU end users have completed the workflow-specific onboarding session
- ✓ BU Operations Owner has passed the readiness gate – able to perform all core operations tasks independently
- ✓ Human-in-the-Loop protocol distributed and confirmed understood by all users and the Operations Owner
- ✓ Support contact card distributed to all users before the transition period begins
- ✓ Operational Handover Document completed – all required sections, version-controlled
- ✓ Handover Acceptance Certificate signed by the BU Operations Owner, the BU Head, and the AI Studio Agent Lead
- ✓ Post-handover support boundary active – AI Studio team on advisory only, Operations Owner resolving daily issues independently



# STEP-BY-STEP INSTRUCTIONS

## STEP 1 BUILD THE CHANGE MANAGEMENT PLAN

Change management for an AI agent deployment is not a communications exercise. It is a resistance management exercise. The corporate organization generates predictable friction patterns when new automation touches existing workflows. Identifying those patterns before they manifest — and assigning a named owner to resolve each one — is what prevents quiet non-adoption from undermining a technically successful deployment.

1.1

### Map the resistance vectors specific to this BU and this agent

Draw from the five antibody patterns established in Guide E3 and apply them to the specific BU receiving this agent:

RESISTANCE PATTERN	HOW IT MANIFESTS	WHAT THE USER FEELS	RESOLUTION OWNER	RESOLUTION MECHANISM
Workflow protection	Users continue the manual process alongside the agent "just to check." The agent becomes parallel overhead, not a replacement.	Fear of making a mistake by trusting something they do not fully understand. Safety in the familiar.	BU Manager with AI Studio Agent Lead support	Define a clear cutover date per workflow step. After cutover, the manual process is not an option. Communicate this before training.
Manager incentive misalignment	The BU manager's team KPIs measure manual output volume. The agent changes the output type. The manager quietly discourages use.	My team's performance metrics are built around the old way of working. The agent makes my numbers look different, not better.	Team Sponsor (E2) + BU Head	Team Sponsor escalates within 48 hours. Executive mandate restated with explicit new success metrics for the manager's team.
Competence anxiety	Users avoid the agent or use it minimally to avoid exposing a skills gap.	If this tool replaces what I do, what is my value? What if I use it wrong in front of my manager?	AI Studio Agent Lead + BU Operations Owner	Training frames the agent as capability amplification. Lead with the specific outcome it enables that was previously impossible in the available time.
Trust deficit	Users override outputs manually without logging the correction. Prior bad experiences with AI tools.	I have seen these tools fail before. I do not want to pass something wrong to a customer or a manager.	AI Studio Agent Lead	Transparent demonstration: show correct outputs, then show escalations, then show a wrong output and how to spot it. Make the failure mode visible.
Technical friction	An awkward login, format requirement, or interface mismatch creates a step that costs more effort than it saves.	This takes longer than doing it myself. It is not worth the hassle for what I get out of it.	Tech Lead + BU Operations Owner	Identify friction from usage logs in the first week. A drop-off rate or 3x latency spike is a friction signal — fix the workflow step, not the user.

1.2

**Define adoption metrics before the first training session** — Adoption is not self-reporting. Define leading indicators that can be observed from usage data, set target ranges, and name the person who reviews them weekly:

ADOPTION METRIC	WHAT IT MEASURES	TARGET RANGE	SIGNAL OF SUCCESS	SIGNAL OF FAILURE
Weekly active users per role	Whether users are returning to the agent — not just activating once	> 70% of trained users active by week 3	Stable or growing WAU after week 2	WAU declining week-on-week after an initial spike — indicates training without habit formation
% of workflow runs via agent	Whether the agent is replacing the manual process or running alongside it	> 80% of eligible workflow runs via agent by week 4	Manual process usage declining	High agent usage but manual process still running — workflow protection pattern active
Feedback mechanism usage rate	Whether users are logging corrections or flagging quality issues	> 20% of sessions generate at least one feedback signal	Feedback volume stable or growing	Zero or near-zero feedback — users are either overriding silently or not checking outputs
Escalation close rate	Whether escalations are being resolved within SLA — indicating the HITL protocol is working	> 90% of escalations resolved within SLA	Close rate stable, SLA breaches < 10%	Low close rate or growing SLA breach rate — HITL protocol is not operational in practice

1.3

**Confirm the executive mandate is specific and visible** — A generic "we are adopting AI" message does not reduce resistance. The mandate must state which agent is being deployed, in which workflow, for which team, with what outcome expected, and by when. This message comes from the BU head — not from the innovation team or the AI Studio. The Team Sponsor is responsible for ensuring this message is delivered before any user training begins.



## 1.4 Build the internal communication plan for this deployment

AUDIENCE	MESSAGE	CHANNEL	TIMING	OWNER
BU Head	Why this agent is being adopted, what outcome it enables for the BU, and what their role is in making adoption succeed. Confirmation that this is permanent, not a pilot.	1:1 briefing from EIR/GM or AI Studio Agent Lead	Before any user communication	EIR / GM
BU Managers	What changes for their team, what the new success metrics are, and how to handle team questions. Explicit: "This is not optional for your team."	BU head-led team meeting	Week before training begins	BU Head + Team Sponsor
End Users	What the agent does, what it does not do, when to trust it, and when to escalate. Not a technology announcement – a workflow change announcement.	Team meeting + written one-pager	Day before training session	BU Operations Owner
BU Operations Owner	Full operations briefing: what they will own after handover, what the support structure is, and what the transition period looks like.	Dedicated session with AI Studio Agent Lead	Before production go-live	AI Studio Agent Lead
All users – ongoing	Monthly update from the BU Operations Owner: adoption metrics, improvements made, issues resolved, and what changes next. Success stories from real users. Keeps the agent visible and improving in the team's eyes.	Team channel or brief monthly email	Monthly post-handover, indefinitely	BU Operations Owner

### AI PROMPT – Change Management Plan

I am building the Change Management Plan for an AI agent deployment to a business unit. Venture: [describe]. Agent: [name and function]. BU receiving the agent: [describe – size, role type, current workflow the agent replaces]. Known resistance risks: [list – identify which of the five patterns apply]. Executive Sponsor: [role]. Team Sponsor: [role]. For each identified resistance pattern: (1) describe how it is likely to manifest for this specific BU and agent, (2) name the resolution owner, (3) write the specific resolution action with a timeline, (4) define the signal that confirms the resistance has been resolved. Output as a one-page Change Management Plan ready for the Team Sponsor review.

## STEP 2 DESIGN AND DELIVER USER ONBOARDING

User onboarding for an AI agent is not a technology training session. It is a workflow change session. The user's job is not to understand how the agent works technically – it is to understand how their workflow changes, when to trust the output, and what to do when the agent cannot help. Any onboarding session that spends more than 10% of its time on technical architecture has the wrong focus.

## 2.1 Design the onboarding session around three questions, not three features

- **Question 1** – What is different about my work now? Show the before and after for the user's specific workflow. The agent's role is to do [specific task] that previously required [specific manual effort]. Use real examples from the F2 pilot or from the first week of production.
- **Question 2** – When should I trust the output, and when should I question it? This is the most important part of the onboarding. Show high-confidence correct outputs. Show edge cases where the agent escalates. Show what a wrong output looks like and how to identify it. Then add two things that are frequently skipped: explain that AI agents can be confidently wrong – producing a plausible-sounding output that is factually incorrect – and demonstrate this with a real example. And cover the data safety boundary explicitly: what users must never paste into the agent (personal data, confidential financials, regulated information) based on the data policies confirmed in G1.
- **Question 3** – What do I do when something goes wrong? Walk through the escalation flow step by step. What does the user see when the agent escalates? What button do they press? Where does it go? Who follows up with them and when? This should take no more than 5 minutes – because the answer should be simple.

- 2.2 Apply the role-based onboarding structure** – Different user roles interact with the agent differently. A single all-hands training session produces mediocre outcomes for everyone. Group users by role and workflow interaction type:

USER ROLE	WHAT THEY NEED TO LEARN	SESSION FORMAT	DURATION
Daily operators (primary users of the agent output)	The full three-question onboarding. Hands-on practice with the agent on their specific task type. The feedback mechanism. The escalation flow. AI literacy basics: hallucinations, confident errors, and data safety rules.	Live session with hands-on component. Maximum 8 users per session.	60–90 minutes
Supervisors and reviewers (users who review or act on agent outputs)	How to identify a high-confidence output versus one requiring review. The escalation log and how to use it. Their role in the Human-in-the-Loop protocol.	30-minute briefing with Q&A	30 minutes
BU Operations Owner	Full onboarding plus: how to read the monitoring dashboard, how to handle a first-level incident, when to escalate to Level 2, and how to interpret the adoption metrics dashboard.	Dedicated session with AI Studio Agent Lead	2–3 hours across 2 sessions
BU Manager	What the agent changes for their team, the new success metrics, how to interpret the weekly performance summary, and how to handle user questions – including how to answer, "will this replace my job?"	30-minute briefing with the BU Head present if possible	30 minutes

- 2.3 Produce the user-facing onboarding materials before the first session** – Three required materials: a one-page workflow change summary (what is different about your work, for this specific role), a visual escalation guide (what to do when the agent cannot help – one page, no more than 5 steps), and the support contact card (who to call, in what order, for what type of issue). These materials are handed to users before the session, not at the end of it.

2.4

**Run a minimum of two onboarding sessions before declaring user readiness** —

The second session — held one week after the first — addresses questions that only emerge after users have interacted with the agent in their real workflow. The first session teaches the workflow change. The second session resolves the friction that the first session could not anticipate.

2.5

**Build prompt literacy and feedback habits into the hands-on component** — Two specific skills that are not obvious but significantly affect output quality: (1) how to give the agent better context when the first output is not quite right — not re-running the same input hoping for a different result, but providing a more specific instruction or example, and (2) how to use the feedback mechanism and why it matters — every correction logged improves the agent for everyone. Make the feedback link between one user's correction and the team's experience of the agent visible in the session. For larger or distributed teams: produce a short, recorded version of the escalation demo and a quick-reference card for users who could not attend the live session.

**AI PROMPT — User Onboarding Session Design**

I am designing the user onboarding session for an AI agent deployment. Agent: [name and function]. Primary user role being onboarded: [describe — role, daily workflow, how the agent changes their work]. The three most likely user concerns for this role are: [list]. The most common edge case or escalation scenario for this agent is: [describe]. Design: (1) the 60-90 minute session agenda with time allocations, (2) the three hands-on practice tasks that represent real workflow scenarios for this role, (3) the one-page workflow change summary for this role, (4) the visual escalation guide (5 steps or fewer), (5) the two questions to ask at session close that will identify users who are not yet confident. Keep all materials at the workflow level — avoid technical detail.

**STEP 3****IMPLEMENT THE HUMAN-IN-THE-LOOP PROTOCOL**

The Human-in-the-Loop protocol is the operational policy that defines when a human must review, override, or take over from the AI agent. It is not a technical configuration — it was built into the agent in G1. It is a human policy: who is responsible for acting on an escalation, what they are authorized to decide, and how their decision is logged.

The protocol has two components: the system-side escalation (the agent flags and hands off) and the human-side response (the user or reviewer acts). Guide G1 configured the system side. This step configures the human side.

### 3.1 Define the four human-side escalation roles

ROLE	WHO HOLDS IT	WHAT THEY ARE AUTHORIZED TO DO	RESPONSE TIME EXPECTATION
First Responder	The daily operator or user who is working at the point of escalation	Accept a manual fallback response for this specific call. Log the escalation outcome. Flag recurring patterns to the Operations Owner.	Immediate – within the same session or workflow step
Reviewer	A designated reviewer (supervisor or senior user) for escalations that require domain judgement	Review the escalated input and produce the correct output. Approve or reject the agent's fallback suggestion. Log the outcome with the reason.	Customer-facing: within 15 minutes. Internal: within 1 hour.
Operations Owner	The named BU individual responsible for the agent after handover	Handle escalations that exceed the Reviewer's authority. Manage pattern-level issues (multiple escalations of the same type). Trigger Level 2 support when the escalation indicates a system-level problem.	Within 4 hours for pattern-level issues
Level 2 Support (AI Studio)	AI Studio Agent Lead (during transition) / retained support contact (post-handover)	Diagnose and resolve system-level escalation causes. Revise the system prompt if a pattern indicates a configuration issue. Deploy the fix through the change control process.	Within 1 business day. Not a real-time role post-handover.

### 3.2 Write the Human-in-the-Loop protocol document – one-page, plain language –

The protocol document answers three questions for each user role: (1) When will I see an escalation? (2) What do I do with it? (3) How do I log what I decided? It must be readable and actionable by a user who reads it once before their first shift with the agent. Include one additional clarification: not all escalations block the user in real time. Some route to an async review channel (a team queue, a Slack message, or an email dashboard) where a reviewer acts within the SLA without the original user waiting. Make the difference explicit – users who expect an immediate response from a non-blocking escalation become confused when nothing happens immediately.

### 3.3 Run a tabletop escalation exercise before the transition period begins –

Simulate three escalation scenarios with the BU Operations Owner and a representative group of users. Walk through each scenario step by step: the agent escalates, the first responder receives it, the reviewer acts, the outcome is logged. The purpose is not to test the system – it is to build the human confidence that comes from having rehearsed the response before it happens in production. Include at least one scenario where the escalation routes to an async queue so users experience both patterns.

### 3.4 For high-stakes or regulated agent functions: apply the override logging requirement –

Every human override of an AI-generated output must be logged with the agent's output, the human's override decision, the reason code, and the decision-makers identifier. This log serves three purposes: the audit trail for regulated environments, the input for the Operations Owner's monthly feedback review, and the improvement data for the G3 monitoring cycle. Every override is a signal – it either indicates a system prompt gap that G3 will address, or a user training gap that the next onboarding session will address. Neither use is possible if the override is not logged.

**AI PROMPT – Human-in-the-Loop Protocol Document**

I am writing the Human-in-the-Loop Protocol for an AI agent deployment. Agent: [name and function]. The escalation trigger conditions are: [list from G1]. The four human roles in the escalation chain are: [First Responder: describe role; Reviewer: describe role; Operations Owner: describe; Level 2 Support: describe]. Write the one-page protocol document for this specific agent: (1) a plain-language description of when each role will see an escalation, (2) exactly what each role does – step by step, in no more than 5 steps per role, (3) how each role logs their decision, (4) the three tabletop scenarios the team will rehearse before the transition period begins. No technical language. Format for a non-technical user.

**STEP 4****PRODUCE THE OPERATIONAL HANDOVER DOCUMENT**

The Operational Handover Document is the institutional memory of the agent deployment. It contains everything the Operations Owner and the BU need to run the agent independently after the AI Studio team steps back. It is produced during the transition period – not written retrospectively after handover. Every section is completed as the knowledge exists, not in a final week of document assembly.

4.1

**Assign the document to a named owner** – the AI Studio Agent Lead – with a completion deadline at the start of the transition period – The handover document is not a post-project output. It is a live document that the Operations Owner should be reading and validating throughout the transition period. If the Operations Owner cannot understand a section, that section needs to be rewritten before handover.

**Operational Handover Document – Required Sections:**

SECTION	CONTENT	WHO WRITES IT	WHO VALIDATES IT
<b>0. Relationship to Other Agents and Tools</b>	For BUs running multiple agents or integrated AI workflows: which upstream or downstream agents' does this agent depend on or interact with? What breaks if one of those agents is unavailable? This section is optional for standalone single-agent deployments but required for any multi-agent or integrated context.	AI Studio Agent Lead	BU Operations Owner + Tech Lead
<b>1. Agent Overview</b>	Name, function, the workflow it automates, the users it serves, the BU that owns it. One paragraph. No technical language.	AI Studio Agent Lead	BU Operations Owner – must be able to read this to a new team member without explanation
<b>2. What the Agent Does and Does Not Do</b>	The specific tasks the agent handles. The specific tasks it does not handle and why (by design, not by failure). The escalation conditions and what they mean in operational terms.	AI Studio Agent Lead	BU Operations Owner

SECTION	CONTENT	WHO WRITES IT	WHO VALIDATES IT
3. Daily Operations Checklist	The tasks the Operations Owner performs to confirm the agent is running normally: monitoring dashboard check, escalation log review, daily accuracy summary confirmation. Include the key business impact metrics the agent is expected to move – average handling time, error rate, cost per transaction, or the equivalent for this BU – and where to find them. Format: step by step, with expected values and warning signal definitions.	AI Studio Agent Lead	BU Operations Owner – must complete this checklist independently in under 15 minutes
4. Escalation and Incident Response	The escalation chain: who does what, in what order, in how much time. The incident response steps for the three most likely failure modes. The rollback trigger conditions and who authorizes rollback post-handover. The agreed SLO/SLA for agent uptime and acceptable escalation rate.	AI Studio Agent Lead	BU Operations Owner + Tech Lead
5. Feedback and Continuous Improvement	How user feedback is captured. How the Operations Owner accesses the feedback log. How override logs feed into the G3 monthly review. The process for passing findings to Level 2 support for system prompt updates.	AI Studio Agent Lead	BU Operations Owner
6. Known Limitations and Workarounds	The specific input types the agent handles less reliably. The manual workaround for each limitation. The trigger that would prompt escalation of a limitation to Level 2 for resolution.	AI Studio Agent Lead	BU Operations Owner – must agree each workaround is feasible for their team
7. Contact and Support Reference	Level 1: BU Operations Owner. Level 2: AI Studio retained contact. Level 3: Tech Lead. Activation condition and response time for each level.	AI Studio Agent Lead	BU Operations Owner
8. Version and Change Log	Agent version at handover. Date of last system prompt update. Known issues under investigation. How changes are requested post-handover – the change-control process: who proposes, who reviews risk, who approves scope changes, and who deploys. Changes to the agent after handover require the same change-control discipline as the original deployment – not an informal prompt edit.	AI Studio Agent Lead	AI Studio Agent Lead + Tech Lead

4.2

**Version-control the document and store it in the content library from E2** – The handover document is a living reference. Post-handover, the Operations Owner owns updates. The document version is recorded alongside the agent configuration version from G1 so that any change to the agent is reflected in the document.



## STEP 5 RUN THE TRANSITION PERIOD

The Operational Handover Document is the institutional memory of the agent deployment. It contains everything the Operations Owner and the BU need to run the agent independently after the AI Studio team steps back. It is produced during the transition period – not written retrospectively after handover. Every section is completed as the knowledge exists, not in a final week of document assembly.

**Operations Owner Competency Model – what this person must be able to do independently after handover:**

- **Monitor:** access and interpret the agent performance dashboard without assistance
- **Triage:** distinguish between a user training issue, a known limitation, and a system-level problem requiring Level 2 escalation
- **Respond:** follow the incident response protocol for the three most likely failure modes
- **Communicate:** report agent performance to the BU team monthly and flag material issues to the AI Studio Level 2 contact with a correctly formatted request
- **Govern:** process a change request through the correct channel – not an informal prompt edit

5.1

### Set the transition period duration based on agent complexity and Operations Owner experience

- **Simple internal agent (low complexity, low stakes):** Minimum 1 week. Operations Owner has prior experience with similar tools.
- **Customer-facing agent or high-volume internal agent:** Minimum 2 weeks. Operations Owner has no prior experience with AI agents.
- **Regulated-sector agent (financial services, healthcare, HR):** Minimum 3–4 weeks. Additional sign-off from Legal and Compliance required before handover acceptance.
- **Coverage:** For agents that run continuously or serve customer-facing workflows: name a deputy Operations Owner before the transition period begins. The primary Operations Owner is unavailable at some point. The deputy must be able to perform the monitoring and triage competencies from the model above.

5.2

### Define the transition period operating rhythm

- **Daily:** AI Studio Agent Lead and Operations Owner jointly review the previous day's monitoring dashboard, escalation log, and any flagged feedback. Duration: 15 minutes. Location: the monitoring dashboard together, not a separate report.
- **Every 2–3 days:** The Operations Owner performs the daily operations checklist independently, then walks the AI Studio Agent Lead through their findings. This is not a check-up – it is a supervised practice run.
- **Weekly:** A structured review of one real incident or escalation from the week: what happened, what the Operations Owner did or would have done, and whether the response matched the protocol.

5.3

### Apply the readiness gate – the Operations Owner must demonstrate each of the following independently before handover

TASK	WHAT THE OPERATIONS OWNER MUST DO INDEPENDENTLY	PASS CRITERION
Daily monitoring check	Access the monitoring dashboard, interpret the key metrics (accuracy, escalation rate, latency, cost), identify whether the agent is performing within normal range, and correctly flag any metric that would require investigation.	Completes without prompting from AI Studio Agent Lead. Correctly interprets all dashboard metrics. Correctly identifies a simulated out-of-range metric.
Escalation response	Receive a test escalation, follow the Human-in-the-Loop protocol, log the response, and identify whether the escalation pattern indicates a system-level issue requiring Level 2 support.	Follows the protocol steps without reference to the document. Correctly identifies which escalation types require Level 2 escalation.
Incident response	Given a simulated incident (error rate spike, latency alert, or high escalation rate), identify the incident type, follow the response procedure from the handover document, and determine whether to escalate to Level 2 or manage at Level 1.	Correctly identifies the incident type. Takes the correct next step within the response time SLA. Documents the incident correctly.
Feedback review	Access the feedback log, identify the two or three most frequently occurring user corrections from the past week, and describe what they indicate about agent performance without prompting.	Correctly identifies the top feedback themes. Can articulate whether they represent a user training issue, a system prompt issue, or a known limitation.
Change request	Identify a change to the agent that would improve performance (e.g., a new edge case to handle), document it in the format required for Level 2 review and correctly route it to the appropriate support contact.	Change request is complete and correctly formatted. Routed to the correct level of support.

5.4

**Do not advance handover if any readiness gate task is not passed** — If the Operations Owner cannot complete a task independently by the end of the transition period: extend the period for that specific task. Do not conduct a full handover with an unresolved gap. A signed handover document with an unresolved operational gap is a liability, not a milestone.

## STEP 6

### ESTABLISH THE POST-HANDOVER SUPPORT MODEL

After handover, the AI Studio team is no longer the first point of contact for operational issues. They are advisory only — a Level 2 resource engaged for system-level problems that the Operations Owner cannot resolve. This boundary must be stated clearly, in writing, before handover acceptance is signed. An unclear support boundary produces an Operations Owner who defaults to calling the AI Studio team for every issue, which is the same as not handing over at all.

## 6.1 Define the three-level support model in the handover document

SUPPORT LEVEL	WHO HANDLES IT	WHAT IT COVERS	RESPONSE TIME	ACTIVATION CONDITION
Level 1 – Operational	BU Operations Owner	Day-to-day monitoring. User questions. Known limitation workarounds. Escalation log review. Routine feedback capture and routing to Level 2.	Same business day	Default – all first-contact issues
Level 2 – Technical	AI Studio Agent Lead (retained post-handover under a defined support arrangement)	System prompt updates. New edge case patterns identified from the escalation log. Integration issues. Regression in accuracy or latency. Requests to extend the agent's scope.	1 business day	Initiated by Operations Owner via a formal change or incident request – not a phone call
Level 3 – Infrastructure	Tech Lead	API failures, infrastructure outages, data pipeline issues, security incidents. Rollback if required.	As defined in the G1 escalation protocol – not changed by handover	Initiated by Level 2 or by an automated alert triggered in the G3 monitoring system

**6.2 Produce and distribute the support contact card before the transition period begins** – One page: the three support levels, named contacts at each level, the activation condition for each level, and the response time expectation. This card is distributed to all end users, all BU supervisors, and the Operations Owner. Users who do not have the contact card will contact the wrong person when something goes wrong.

**6.3 Define the support engagement format for Level 2** – Level 2 support is not a helpdesk. It is a structured engagement between the Operations Owner and the AI Studio team for issues that require system-level attention. The Operations Owner logs a request using a standard format: the agent name, the observed symptom, the frequency (number of occurrences), the relevant log entries, and the business impact. Without this format, Level 2 support receives fragmented information that is slow to diagnose.

**6.4 Agree the Level 2 support retention period and scope at handover** – The AI Studio team's post-handover Level 2 availability is a defined arrangement – not an open-ended obligation. Agree in writing before handover acceptance: the duration of the Level 2 support arrangement, the scope (what is included and what is not), the handover from AI Studio to a permanent internal or retained technical resource when the arrangement ends.

## STEP 7

## EXECUTE THE FORMAL HANDOVER AND ACCEPTANCE SIGN-OFF

Handover acceptance is a formal event — not a conversation, an email, or an implicit assumption that the transition period is over. It is a documented sign-off against binary criteria. The Operations Owner is not signing to say they are comfortable. They are signing to confirm that they have demonstrated specific competencies and accept specific accountabilities.

- 7.1 **Confirm all readiness gate tasks have been passed** — Every task in Step 5.3 must be marked as passed with evidence before the handover session is scheduled. If any task is outstanding: the handover session is not scheduled.
- 7.2 **Confirm the Operational Handover Document is complete** — Every section in Step 4 must be complete and validated by the Operations Owner. The Operations Owner must confirm that every section is accurate and that they could use it to resolve an issue without contacting the AI Studio team.
- 7.3 **Hold the handover session with three attendees: the BU Operations Owner, the BU Head, and the AI Studio Agent Lead** — The BU Head must attend. Their presence signals organizational endorsement of the accountability transfer. Their absence signals that accountability is not real.
- 7.4 **Complete the Handover Acceptance Certificate**

**Handover Acceptance Certificate — Required Content:**

- Agent name, version, and deployment date
  - BU Operations Owner name, role, and acceptance of the operational accountability described in Section 7 of the Handover Document
  - BU Head confirmation of executive sponsorship and internal escalation path
  - AI Studio Agent Lead confirmation that the readiness gate has been passed and the Handover Document is complete and current
  - Agreed SLO/SLA: agent expected uptime, maximum acceptable escalation rate within SLA, and the agreed Level 2 response time — both parties confirm these expectations before signing
  - Post-handover support arrangement: Level 2 contact name, retention period, and scope
  - First post-handover check-in date: a scheduled 30-day review between the Operations Owner and the AI Studio Agent Lead
  - 90-day value review date: a scheduled review between the BU Head, the Operations Owner, and the AI Studio Agent Lead to assess business impact, usage trajectory, and whether the agent's scope should be expanded, adjusted, or retired
  - Date, signatures
- 7.5 **File the Handover Acceptance Certificate in the content library (E2) alongside the Operational Handover Document, the Go-Live Report from G1, and the agent configuration version tag**

### AI PROMPT – Handover Acceptance Certificate Draft

I am preparing the Handover Acceptance Certificate for an AI agent deployment. Agent: [name and version]. BU Operations Owner: [name and role]. BU Head: [name]. AI Studio Agent Lead: [name]. Handover Document: [confirm complete – yes/no and last updated date]. Readiness gate results: [list each task and Pass/Fail]. Post-handover support arrangement: [Level 2 contact, retention period, scope]. Known limitations outstanding: [list]. Draft the complete Handover Acceptance Certificate ready for signature. Include a plain-language summary of what the Operations Owner is accepting – what they are now responsible for and what the AI Studio team is no longer responsible for. Both parties must be able to read this and have no ambiguity about where accountability sits after the date of signature.



# 6

## TROUBLESHOOTING

SYMPTOM	LIKELY CAUSE	FIX
Users are bypassing the agent and reverting to the manual process	Training covered how the agent works but not why it was introduced. Users do not see the benefit to their specific workflow.	Rerun the workflow-specific training using before/after comparisons. Make the time or effort saved visible: "This step previously took 20 minutes. The agent completes it in 40 seconds." Secure a manager endorsement that is workflow-specific – not a general innovation message.
The BU Operations Owner accepts the handover but cannot resolve production issues independently	Handover happened on schedule but the Operations Owner was not equipped to perform the core operations tasks before sign-off	The readiness gate must be passed before handover – not on a calendar date. Return to Step 5 (transition period). The Operations Owner must perform each core task in Step 5.3 without assistance before handover acceptance proceeds.
Middle management blocks user adoption after the agent is live	BU managers are incentivised on throughput metrics that the agent changes. The agent is seen as a threat to their team's value, not as a tool that helps them.	This is the antibody resistance pattern from E3. Escalate to the Team Sponsor immediately. The executive mandate must be restated specifically for this BU. Work with the BU manager to define what success looks like for their team – then tie the agent adoption metric to that outcome.
The AI Studio team is drawn back into day-to-day operations after handover	Handover acceptance was signed but the Operations Owner continues to escalate operational issues to the AI Studio team instead of following the support protocol.	Apply the support boundary from Step 6 strictly. The AI Studio team does not respond to operational issues that fall within the Operations Owner's remit. Restate the support boundary in writing. If the Operations Owner genuinely cannot handle an issue category, return to Step 5 for additional transition period work on that specific task.
Users report that the agent produces outputs that feel wrong but do not escalate them	No friction-free feedback mechanism exists. Users are solving the problem themselves by ignoring the output or manually correcting it, without logging the correction.	Activate the feedback path designed in F1. The feedback mechanism must require minimal effort – a single click or button, not a form. Users who manually correct outputs without logging are generating valuable training data that is being lost. Make the feedback path visible in the agent interface.
The support escalation path is unclear during an incident – multiple teams contacted simultaneously	The Level 1 / Level 2 / Level 3 support boundary was documented but not communicated to users or the Operations Owner before handover.	Distribute the support contact card (from Step 6) to every user and to the BU Operations Owner before the transition period begins – not at handover. Hold a 30-minute support briefing before the first real user interaction.
Usage is flat or declining after handover; feedback volume is near zero	No ongoing reinforcement. The agent is not visible in the team's daily experience after the initial training period. Benefits are not being reported. The agent may not be embedded in the primary tools the team uses.	Activate the monthly ongoing communication from Step 1. The BU Operations Owner reports the adoption metrics – WAU, workflow coverage rate, feedback volume – to the team each month. Highlight a specific outcome the agent enabled. Check whether the agent entry point is embedded in the tools the team already uses, or whether it requires a separate step that creates friction.
Handover acceptance is delayed because the Operations Owner is waiting for the agent to be "perfect"	The acceptance criteria were not defined as binary and observable. The Operations Owner is evaluating whether they feel comfortable, not whether the agent meets the agreed standard.	Return to the acceptance criteria in Step 7. Each criterion must have a binary observable test. "The Operations Owner feels confident" is not a criterion. "The Operations Owner performed the incident response procedure independently without guidance from the AI Studio team" is a criterion.

# VALIDATION STEPS

Confirm each of the following before declaring the handover complete:

Change Management Plan completed – resistance vectors mapped, executive mandate confirmed, communication plan delivered to all audiences



User onboarding completed – all identified user roles have attended a role-specific session; a second session held one week after the first



User-facing materials distributed: workflow change summary, escalation guide, support contact card



Human-in-the-Loop Protocol document distributed and confirmed understood by all users and the Operations Owner



Tabletop escalation exercise completed with the Operations Owner and representative users



Operational Handover Document complete – all 8 sections, validated by the Operations Owner, version-controlled



All readiness gate tasks passed independently by the Operations Owner – no task marked passed with AI Studio assistance



Post-handover support boundary communicated in writing to all users, supervisors, and the Operations Owner



Handover Acceptance Certificate signed by BU Operations Owner, BU Head, and AI Studio Agent Lead



Certificate filed in content library (E2) with all related documents



30-day post-handover check-in scheduled



# NEXT STEPS

Upon completing this guide with a signed Handover Acceptance Certificate:

- **GUIDE G3** – How to Monitor AI Agent Performance, Detect Drift, and Report to the Venture Board (the Operations Owner is now a participant in the G3 monthly review, not a recipient of a summary)

The 30-day post-handover check-in is not a formality. It is the first structured evidence of whether the handover worked. The Operations Owner presents their first month of independent monitoring data, any incidents handled, and any patterns observed in the escalation log. If the Operations Owner cannot present this independently, the handover was premature – address the gap then, not at the 90-day review.

The 90-day value review is different in focus. It asks: is the agent delivering the business impact it was designed to deliver? Are adoption metrics healthy? Are there opportunities to expand scope, or signals that the agent should be retired or replaced? The BU Head attends this review – not just the Operations Owner. If the answer at 90 days is that the agent is running but not visibly moving any business metric, that is a signal for the G3 monitoring guide, not a sign that the handover failed.

The Operate-to-Transfer principle applies here: the goal is not for the AI Studio team to be needed indefinitely. The goal is for the BU to own and improve the agent with decreasing AI Studio involvement over time. The G3 monitoring guide closes the loop – it gives the Operations Owner the tools to detect and respond to performance changes without requiring AI Studio involvement for every decision.



# MASTER CHECKLIST

## A. CHANGE MANAGEMENT

- Resistance vectors mapped for this specific BU and agent – all five antibody patterns assessed, user motivation and worry captured for each applicable pattern
- Resolution owner named for each identified resistance vector
- Adoption metrics defined before first training session: WAU, % workflow runs via agent, feedback usage rate, escalation close rate – with target ranges and named reviewer
- Executive mandate confirmed as specific and visible – states which agent, which workflow, for which team, with what outcome
- Mandate delivered by BU Head to BU Managers before any user communication
- Internal communication plan executed: BU Head briefed, BU Managers briefed, end users briefed, Operations Owner briefed
- Ongoing communication cadence confirmed: monthly update from Operations Owner to team, post-handover

Team Sponsor (from E2) confirmed as the escalation owner for manager incentive misalignment and trust deficit resistance patterns

## B. USER ONBOARDING

- Onboarding sessions designed around three questions: what changed, when to trust / question (including hallucination and confident-error risk), what to do when it fails
- Data safety rules confirmed and included in Question 2: what users must not submit to the agent
- Role-based sessions designed: daily operators (60–90 min), supervisors (30 min), Operations Owner (2–3 hours), BU Managers (30 min)
- Three user-facing materials produced: workflow change summary, visual escalation guide (5 steps max), support contact card
- Materials distributed to users before the session – not at the end
- Hands-on component includes prompt literacy task (improving output through better context) and feedback mechanism practice
- Scalable formats prepared for distributed teams: recorded escalation demo, quick-reference card
- First onboarding session completed for all user roles
- Second onboarding session completed one week after the first

### C. HUMAN-IN-THE-LOOP PROTOCOL

- Four human-side escalation roles defined with named individuals: First Responder, Reviewer, Operations Owner, Level 2 Support
- Response time SLAs confirmed for each role
- Real-time (blocking) vs asynchronous escalation paths documented – users briefed on which pattern applies
- Human-in-the-Loop Protocol document written in plain language – one page, no technical terminology
- Document distributed to all users and the Operations Owner before the transition period begins
- Tabletop escalation exercise completed with three scenarios – at least one async escalation scenario included
- Override logging requirement confirmed; override log explicitly linked to G3 improvement loop and monthly Operations Owner review

### D. OPERATIONAL HANDOVER DOCUMENT

- Section 0 completed (if multi-agent): relationship to other agents and tools, upstream/downstream dependencies
- All core sections completed: Overview, What agent does/does not do, Daily operations checklist (incl. business metrics), Escalation and incident response (incl. SLO/SLA), Feedback and continuous improvement, Known limitations, Contact and support reference, Version and change log (incl. change-control process)
- Each section validated by the Operations Owner – they can use it independently
- Document version-controlled and stored in content library (E2)
- Agent configuration version at handover recorded in the change log

### E. TRANSITION PERIOD

- Operations Owner competency model confirmed: monitor, triage, respond, communicate, govern
  - Deputy Operations Owner named for customer-facing or continuously-running agents
  - Transition period duration confirmed based on agent complexity: minimum 1 week (simple internal), 2 weeks (customer-facing or high-volume), 3–4 weeks (regulated sector)
  - Daily transition operating rhythm active: joint monitoring review (15 min), supervised practice runs every 2–3 days, weekly incident review
  - All five readiness gate tasks completed independently by the Operations Owner: daily monitoring check, escalation response, incident response, feedback review, change request
  - No readiness gate task marked passed while AI Studio Agent Lead was present and assisting
- Transition period extended for any task not yet independently passed

## F. POST-HANDOVER SUPPORT MODEL

- Three-level support model documented in the Handover Document: Level 1 (Operations Owner), Level 2 (AI Studio retained), Level 3 (Tech Lead)
- Support contact card distributed to all users and supervisors before transition period begins
- Level 2 support engagement format confirmed: named contact, activation condition, request format, response time
- Level 2 support retention period and scope agreed in writing before handover acceptance
- Support boundary communicated in writing: AI Studio is advisory only after handover date

## G. HANDOVER ACCEPTANCE

- All readiness gate tasks confirmed passed – evidence on file for each
- Operational Handover Document confirmed complete and current by Operations Owner
- Handover session held with BU Operations Owner, BU Head, and AI Studio Agent Lead all present
- Handover Acceptance Certificate completed with all required content including agreed SLO/SLA
- Certificate signed by BU Operations Owner, BU Head, and AI Studio Agent Lead
- Certificate filed in content library (E2) with Go-Live Report, Handover Document, and agent configuration version tag
- 30-day post-handover check-in scheduled in the E2 operating cadence calendar
- 90-day value review scheduled with BU Head, Operations Owner, and AI Studio Agent Lead

