


PROBLEM

AI agents are becoming operational workers, but real work is not a chat transcript.

When an agent touches production systems, teams need to know what happened, what evidence was used, what risk was found, who approved it, and how work can be resumed safely.

AgentOps Case Control Room

Turn AI-agent work into a UiPath-governed case with evidence, risk, approval, robot work items, and human handoff.




SOLUTION

A UiPath-governed case for human-AI operations.

- Human opens an exception case.
- AI agent creates a read-only investigation plan.
- Robot gathers ticket, pull request, and owner evidence.
- Policy gateway blocks risky production action.
- Action Center-style tasks route approval to humans.

Core thesis: agents can move fast, but UiPath makes the work governable.



ARCHITECTURE

Maestro Case

Case stages: intake, agent investigation, robot/API evidence, risk review, approval, handoff.

Robot Work Items

Evidence collection is modeled as queue-ready robot work items.

Action Center

Human approval tasks preserve evidence, risk, and decision context.

Implementation artifacts include a BPMN-style process blueprint, case data model, Action Center form schema, robot queue model, state machine, transaction logs, and simulated case run.

uipath-package/maestro-process.bpmn

runtime/case-state-machine.json

runtime/orchestrator-transaction-log.jsonl

Repeatable local verification is included.

- 26 shared AgentOps events
- 3 cases
- 5 approval gates
- 1 blocked production action
- 2 robot work items
- 3 Action Center tasks
- 6 Orchestrator transaction events
- 9 Action Center decision events
- Final state: ClosedRejected

```
cd uipath-agenthack
```

```
bash scripts/run_uipath_local_checks.sh
```

```
verify_ok
```

```
uipath_verify_ok
```

WHY NOW

AgentOps Case Control Room shows how human, AI, robot, and API work can be reviewed as one case. That is the bridge between a

Human-AI Operations Control Plane

One event stream for humans, AI agents, robots, APIs, approvals, risks, costs, and evidence. This is the shared prototype behind the UiPath, Splunk, FIND EVIL, Google Cloud, and Microsoft hackathon lanes.

Events
26

Cases
3

Approvals
5

Blocked
1

High/Critical
3

Est. Cost
\$4.86

UiPath Case Workbench

1 Intake
Human opens the case
evt-0001

2 Agent Investigation
AI agent plans and tests read-only
evt-0002, evt-0005

3 Robot Evidence
UiPath robot and APIs gather records
evt-0003, evt-0004, evt-0009

4 Risk Review
Control plane escalates and blocks risk
evt-0006, evt-0007

5 Human Approval
Humans ask for evidence and reject release
evt-0008, evt-0010

6 Handoff
Report is generated from evidence IDs
evt-0011

Human Approval Queue

evt-0007 codex-ops-agent
Attempted production deployment was blocked by policy before execution.
BLOCKED BY POLICY

evt-0008 operations-lead
Human requested more evidence instead of approving the production deployment.
NEEDS MORE EVIDENCE

evt-0010 service-owner
Service owner rejected the release until the retry regression is fixed.
REJECTED

evt-0018 security-analyst
Human approved containment of the reported message and requested no endpoint isolation.
APPROVED

evt-0025 support-manager
Human approved the workflow with a rule that expensive model calls are escalation-only.
APPROVED

Contest Adapters

UiPath
case_packet.json

Splunk
hec_events.jsonl

FIND EVIL
dfir_report.md

Google
gemini_mcp_workflow.json

Microsoft
copilot_architecture.md

CASE-AI-OPS-001
AI Agent Operations Exception Review
critical

11 events3 approvals

1 blocked\$0.006

1. evt-0005 Regression test failed around duplicate retry suppression.
2. evt-0006 Production payment behavior change plus failing regression test.
3. evt-0007 Production deployment is destructive and requires human approval.
4. evt-0008 Human requested more evidence instead of approving the production deployment.

CASE-DFIR-002
Evidence-Locked Phishing Investigation
high

8 events1 approvals

0 blocked\$0.012

1. evt-0015 Initial model answer overclaimed malware execution before process evidence was present.
2. evt-0016 Unsupported certainty detected in analyst draft.
3. evt-0017 Hash lookup returned low-confidence phishing kit association; raw indicator was redacted in report output.

CASE-CLOUD-003
Human-Controlled Cloud Agent Workflow
medium

7 events1 approvals

0 blocked\$4.842

1. evt-0024 Projected monthly model spend exceeded the manager's prototype budget.

Event Timeline							Case	All cases	
EVENT	PHASE	ACTOR	TYPE	STATUS	RISK	SUMMARY			
evt-0001	intake	human operations-lead	task_start	info	none	Human opened an AI-agent operations case after a release bot proposed a production change.			
evt-0002	planning	ai_agent codex-ops-agent	plan_update	success	low	Agent created a read-only investigation plan with explicit approval gates before any production action.			
evt-0003	evidence_collection	robot uipath-evidence-robot	robot_task	success	low	Robot gathered the change ticket, linked pull request, deployment window, and service owner.			
evt-0004	evidence_collection	api github-api	api_call	success	low	API confirmed the pull request touches payment retry behavior and has one approving review.			
evt-	investigation	ai_agent	tool_call	failed	medium	Agent found a failing retry regression test that was not mentioned in the change			

prototype on a laptop and agentic automation that enterprises can trust.

Next: import the generated process, queue, and approval models into UiPath Automation Cloud / Maestro and verify a live run.

0005		codex-ops-agent		failure	medium	ticket.
evt-0006	risk_review	ai_agent codex-ops-agent	risk_signal	warning	high	Case risk escalated because the proposed release affects payments and has failing verification.
evt-0007	execution	ai_agent codex-ops-agent	tool_call	blocked	critical	Attempted production deployment was blocked by policy before execution.
evt-0008	approval	human operations-lead	approval_gate	success	medium	Human requested more evidence instead of approving the production deployment.
evt-0009	evidence_collection	robot ulpath-evidence-robot	robot_task	success	low	Robot requested service-owner signoff and attached the failed test evidence.
evt-0010	approval	human service-owner	approval_gate	success	low	Service owner rejected the release until the retry regression is fixed.
evt-0011	handoff	system agentops-recorder	handoff_report	success	none	Generated a case handoff report with cited evidence IDs and final rejection decision.
evt-0012	intake	human security-analyst	task_start	info	none	Security analyst opened a suspicious email investigation with attached endpoint telemetry.
evt-0013	planning	ai_agent dfir-analysis-agent	plan_update	success	low	Agent proposed an evidence-first DFIR plan: preserve artifacts, list hypotheses, then verify each claim.
evt-0014	evidence_collection	robot sift-collector	evidence_captured	success	low	Collected email headers, attachment hash, browser download timeline, and endpoint process list.
evt-0015	investigation	ai_agent dfir-analysis-agent	ai_call	warning	medium	Agent generated an initial hypothesis but flagged one claim as unsupported by current evidence.
evt-0016	risk_review	ai_agent dfir-analysis-agent	risk_signal	warning	high	System marked the draft finding as provisional because execution evidence was missing.
evt-0017	evidence_collection	api threat-intel-api	api_call	success	medium	Hash lookup returned low-confidence phishing kit association; raw indicator was redacted in report output.
evt-0018	approval	human security-analyst	approval_gate	success	low	Human approved containment of the reported message and requested no endpoint isolation.
evt-0019	handoff	system agentops-recorder	handoff_report	success	none	Generated DFIR report with evidence IDs, unsupported claims, and human containment decision.
evt-0020	intake	human support-manager	task_start	info	none	Support manager requested an agent that answers customer questions and escalates uncertain cases.
evt-0021	planning	ai_agent gemini-ops-navigator	plan_update	success	low	Agent created a workflow plan using retrieval, MCP tools, cost guardrails, and human escalation.
evt-0022	evidence_collection	api knowledge-base-mcp	api_call	success	low	MCP retrieval returned the current refund policy and source URL for citation.
evt-0023	execution	ai_agent gemini-ops-navigator	ai_call	success	low	Agent drafted a customer answer grounded in the retrieved refund policy.
evt-0024	risk_review	ai_agent gemini-ops-navigator	cost_signal	warning	medium	Cost guardrail warned that the high-quality model should be reserved for escalations.
evt-0025	approval	human support-manager	approval_gate	success	low	Human approved the workflow with a rule that expensive model calls are escalation-only.
evt-0026	handoff	system agentops-recorder	task_end	success	none	Closed cloud-agent workflow case with source citations, cost guardrail, and human approval captured.