

OPERATING TEMPLATE

AI agent SLA template for revenue workflows

Define the service before you set targets. Replace every bracketed field with evidence from the workflow you actually run.

Use this template to define the service before setting targets. Replace every bracketed field with evidence from the actual workflow. This document is an operating template, not legal advice or proof of regulatory compliance.

1. Service definition

- Workflow name:
- Business purpose:
- Service owner:
- Workflow owner:
- Data owner:
- Approval owner:
- Incident/recovery owner:
- Effective date:
- Review cadence:

In scope

- Eligible trigger:
- Required inputs:
- Permitted agent steps:
- Required output and evidence:
- Designated system of record:
- Teams/segments covered:

Out of scope and prohibited

- Prohibited actions:
- Restricted records or fields:
- Decisions reserved for humans:
- Conditions that make work ineligible:

2. Workflow states

State	Entry condition	Exit condition	Required timestamp/evidence	Accountable owner
Eligible	Trigger and prerequisites pass validation	Accepted or bounded rejection	Eligibility result and start time	
Acknowledged	Work is accepted	Prepared, failed or escalated	Acceptance state and reason	
Prepared	Permitted agent steps are complete	Human decision or direct completion	Artifact, sources and policy version	
Reviewed	Approval is required	Approved, edited, rejected or escalated	Reviewer, decision and timestamp	
Recorded	Approved result is ready	System-of-record write verified	Before/after state and write result	
Verified	Completion checks pass	Service item closed	Final status, evidence and exceptions	

3. Service-level indicators and objectives

Set thresholds from a measured baseline. Report agent processing time and human approval time separately even when the end-to-end commitment includes both.

Indicator	Definition	Numerator	Denominator	Target	Window	Data source	Owner
Acknowledgement time	Eligible to accepted or bounded rejection		Eligible items				
Preparation time	Accepted to agent-prepared output		Accepted items				
Trustworthy completion	All required steps, evidence and records complete	Complete items	Eligible items due				
Evidence completeness	Required evidence present and reviewable	Items with complete evidence	Prepared items				
Approval latency	Ready for review to human decision		Items requiring approval				
Escalation rate	Work routed to the documented exception owner	Escalated items	Eligible or accepted items				
Recovery time	Detected failure to verified safe state		Recoverable incidents				

4. Exclusions and clock rules

Reason code	When it applies	Clock behavior	Evidence required	Approver
Missing prerequisite		Stop / pause / count		
Source unavailable		Stop / pause / count		
Waiting for approval		Stop / pause / count		
Policy exception		Stop / pause / count		
Planned maintenance		Stop / pause / count		

No exclusion may be applied without a reason code and supporting evidence.

5. Permission and approval policy

Action	Read scope	Write/action scope	Approval required	Approver	Required evidence	Prohibited condition

6. Escalation and recovery

- Exception categories:
- First-response owner:
- Escalation time:
- Business notification path:
- Safe fallback behavior:
- Rollback method:
- Verification after recovery:
- Incident evidence retained:
- Conditions requiring workflow suspension:
- Conditions required before restart:

Recovery rehearsal

Failure scenario	Expected safe behavior	Recovery owner	Test date	Result/evidence	Follow-up
Required source is stale					
A partial write succeeds					
Human approval is overdue					
A prohibited action is attempted					

7. Change control

The following changes require documented review before production:

- Model or model-routing change.
- Prompt, policy or approval-threshold change.
- New data source, object, field or tool.
- Expanded role, segment, geography or action right.
- Modified definition of eligibility or trustworthy completion.
- Changed evidence, logging, escalation or recovery behavior.

Change	Risk/impact	Evaluation evidence	Approvers	Release date	Rollback plan

8. Operating dashboard

Review at least:

- Eligible, accepted, completed, excluded, failed and recovered counts.
- Attainment by workflow, priority, team and reason code.
- Evidence completeness, human acceptance, edit and rejection rates.
- Approval queue age and decision latency.
- Open exceptions, oldest exception and recovery-time distribution.
- Policy, context, model, tool and threshold changes.
- The workflow's existing business KPI, without unsupported attribution.

9. Review decision

- Service objective met? Yes / No
- Critical control failure? Yes / No
- Scope stays the same / narrows / widens / pauses
- Owner accepting the decision:
- Evidence reviewed:
- Corrective actions and due dates:

RevTech provides fully managed AI agents for GTM. RevTech runs the agents; customers govern objectives, permissions, approvals and business decisions.

RevTech provides fully managed AI agents for GTM: RevTech runs the agents, your team governs them. This worksheet accompanies revtech.ai/resources/ai-agent-sla-compliance and is an operating template, not legal, regulatory or security advice.