IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Claude Managed Agents GxP Human Approval Tool Permissions

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-09-27 · claude managed agents · gxp · human approval · tool permissions · mcp · audit trail · validation · anthropic · life sciences

Original article: https://intuitionlabs.ai/articles/claude-managed-agents-gxp-approval-control-matrix

In this article

Executive Summary

As of September 27, 2026, **Claude Managed Agents** is a public-beta hosted agent service with a newly expanded permission model. Anthropic introduced the `auto` policy on September 10: the server can allow, pause, or deny each enabled agent or Model Context Protocol (MCP) tool call after considering its input and the session context. The same release added live intervention through `ant beta:sessions connect` in CLI version 1.32.0. Those capabilities are useful control inputs, but neither the server's risk judgment nor a terminal approval is itself a validated GxP authorization or electronic signature. (platform.claude.com) (claude.com) (www.ecfr.gov) (www.iso.org)

The decision rule is direct. Use `always_ask` where a named, authorized person must approve **every** call before execution. Use `auto` only where preapproved operating rules permit some calls to execute without review. An `auto` result of `ask`, with reason code `indeterminate`, can be approved or denied through `user.tool_confirmation`; an `auto` result of `deny`, with reason code `high_risk`, cannot be overridden. Missing or disabled tools are denied before policy evaluation. Custom tools sit outside this permission mechanism and require application controls. The agent toolset defaults to `always_allow`, while MCP toolsets default to `always_ask`, so the organization must inspect effective configuration, not infer it from a pilot transcript. (platform.claude.com)

The recommended evidence unit is one attempted tool call: immutable event ID, session and agent version, tool and input digest, policy configuration, top-level permission outcome, full evaluation object and reason code, reviewer identity and decision where applicable, final tool result, record identifiers, and timestamps. The API exposes tool-use and idle events, but an enterprise must join them to its own identity, change, retention, and business-record controls. A waiting approval has no documented timeout, and webhook delivery can be duplicated or out of order. A scripted consumer should reconcile against persisted session events rather than treat a webhook arrival as the complete audit record. (www.ecfr.gov)

For a **GxP pilot**, keep literature retrieval and draft creation in a bounded read or draft environment; require explicit approval before record mutation; and reserve regulated submission or batch-release actions for authorized business systems and their established signatures. Test every matrix path, including denial, partial

confirmations, unknown schema values, reconnect, and custom-tool handling. Report observed rates with their denominators, without converting an `auto` allow share into a quality or compliance score. Part 11 demands validated intended performance, access restriction, and time-stamped audit trails where it applies; EU Annex 11 similarly calls for validated applications and recorded access changes. The fit decision therefore belongs to the regulated organization's intended-use assessment and controls, not to product labeling. (www.ecfr.gov) (health.ec.europa.eu) (health.ec.europa.eu) (airc.nist.gov)

Introduction and Background

GxP is shorthand for regulated good practices across functions such as manufacturing, clinical research, and laboratories. This report addresses platform owners, quality assurance (QA), information security, computer-system validation, and regulatory operations evaluating Claude Managed Agents as a tool-execution layer. The question is narrow: which tool calls can run, which pause for a person, which are denied, and what evidence must a regulated organization retain? The MHRA describes its GxP data-integrity guidance as covering pharmaceutical-lifecycle and good laboratory practice activities; its scope illustrates why a single generic permission profile cannot cover every workflow. (www.gov.uk) (database.ich.org) (www.who.int)

Anthropic announced Managed Agents in public beta on April 8, 2026. A session has an agent and environment, retains conversation history, and can preserve progress across disconnections. The product's event stream exposes tool and session activity, while the organization remains responsible for deciding whether those events satisfy its own record and approval requirements. The distinction matters because 21 CFR Part 11 applies to certain electronic records under US Food and Drug Administration (FDA) requirements, and its closed-system provisions address validation, access, and audit trails. It does not certify a vendor product by name. (claude.com) (www.ecfr.gov) (www.ecfr.gov)

The September release creates a practical design choice that older, broad descriptions of human oversight miss. `auto` is a per-call evaluation, and two invocations of the same tool can produce different outcomes. Therefore, a reviewer cannot approve a tool class once and assume that all later calls will pause. If every execution needs review, Anthropic explicitly points to `always_ask`. The useful unit of analysis is the **individual attempted call**, with its input, evaluated policy, decision, result, and relation to a regulated record. (platform.claude.com) (airc.nist.gov)

IntuitionLabs describes its work as governed life-sciences workflows and an information layer with identity, permissions, retrieval, and evaluation. That is an adjacent advisory perspective here: the consultancy is not a Managed Agents permission-policy alternative. The analysis below treats the vendor mechanics as documented behavior and the GxP control design as an implementation proposal that a regulated owner must verify for its own context of use. (intuitionlabs.ai) (intuitionlabs.ai) (www.fda.gov)

Key Changes

September 10 permission evaluation

On **September 10, 2026**, Anthropic added `auto` for agent and MCP tool calls and added an `evaluation` object beside `evaluated_permission` on the corresponding events. The top-level outcome is `allow`, `ask`, or `deny`. For `auto`, the nested object records the same outcome and, for `ask` or `deny`, a reason code. This is evidence about how the platform handled a call, not a deterministic classification of the call's GxP criticality. The server considers the named tool, the input, and session content, so a change in prompt or call arguments can change the result. (platform.claude.com)

An `auto` allow executes before human review. An `auto` ask pauses because the server made no determination; its reason code is `indeterminate`. An `auto` deny stops a call evaluated as high risk, returns an error result to the agent, and cannot be reversed through the confirmation event; its reason code is `high_risk`. Clients must accept future `evaluation.type` or `reason_code` values without silently treating them as approval. A safe implementation stores the raw value and routes unknown combinations to a controlled exception queue. That last routing rule is this report's proposed control, not an Anthropic guarantee. (platform.claude.com) (database.ich.org)

Live session connection

The `ant` CLI **1.32.0** added `ant beta:sessions connect`; the guide requires that version or later and an interactive terminal for the terminal view. An operator can follow a session, message it, and allow or deny a call waiting on approval. This is valuable for supervised development, replay observation, and a small pilot. It is not, by itself, an enterprise approval register: a production control normally needs reviewer authentication, duty assignment, reason capture, durable evidence export, and reconciliation with the business record. NIST's AI Risk Management Framework calls for human-oversight processes to be defined and documented. (platform.claude.com) (airc.nist.gov) (www.ecfr.gov)

The event API provides a different integration surface. A client can consume events, branch on `ask`, capture the approver's decision, and send `user.tool_confirmation`. CLI scripts can write each event to standard output as it arrives. A webhook can wake a workflow, but Anthropic says webhook delivery is not ordered or a durable log. A robust approval service should fetch or stream persisted events, deduplicate notification IDs, and reconcile state after disconnects.

Boundary between server and application

Permission policies apply to the **server-executed agent and MCP toolsets**. Custom tools execute in the customer's application; their `agent.custom_tool_use` events lack both permission fields. For custom tools, the application decides whether to execute before returning `user.custom_tool_result`. That boundary is especially relevant when a custom connector can write to a validated document or quality system. The application must implement its own authentication, authorization, approval, idempotency, and audit trail. (platform.claude.com) (www.ecfr.gov) (health.ec.europa.eu)

Policy configuration also has lifecycle consequences. A toolset has a default and may use per-tool `configs`; no toolset uses `auto` by default. A disabled or absent tool is a separate denial case. Agent updates apply to newly created sessions, while session-specific updates require an idle session and replace the supplied tool or MCP array. The validation baseline should therefore identify the effective agent version and session configuration for every test, then retest material changes. (airc.nist.gov)

Permission Outcomes and Control Matrix

Table 1 maps documented platform outcomes to the proposed GxP control response. “Override” means whether `user.tool_confirmation` can change the outcome of that individual call; it does not mean whether an administrator can later change policy configuration.

Policy or path	Event outcome and evaluation	Execution and human action	Confirmation override	Evidence to retain
`always_allow`	`allow`; `evaluation.type=always_allow`	Runs automatically. Restrict to preapproved low-impact capability.	No confirmation target.	Effective policy, input digest, tool result, actor context, time. (platform.claude.com)
`always_ask`	`ask`; `evaluation.type=always_ask`	Pauses; authorized reviewer allows or denies each call.	Yes, for that `ask` event.	Blocking ID, reviewer, decision, rationale, confirmation event, result. (platform.claude.com)
`auto` allow	`allow`; nested `auto/allow`	Runs before review; monitor and sample according to risk plan.	No confirmation target.	Full evaluation, input/context reference, result, later review status. (platform.claude.com)
`auto` ask	`ask`; `auto/ask`, `indeterminate`	Pauses; reviewer decides.	Yes, for that `ask` event.	Reason code, blocking ID, reviewer, rationale, result. (platform.claude.com)
`auto` deny	`deny`; `auto/deny`, `high_risk`	Does not run; session may continue after error result.	No; platform denies override.	Denial, reason, attempted input, error result, exception disposition. (platform.claude.com)
Tool absent or disabled	`deny`; no `evaluation`	Does not run; investigate configuration or agent behavior.	No.	Enabled-tool snapshot, attempted name, denial, change ticket. (platform.claude.com)
Custom tool	Neither permission field on `agent.custom_tool_use`	Application decides before execution and returns result.	Outside this confirmation policy.	Application authorization, approval, execution, record audit trail. (platform.claude.com)

The matrix separates a **platform result** from a **quality disposition**. A server denial is an observation to record, not proof that every unsafe action was stopped; an allow is an execution decision, not proof that a record is fit for use. For mandatory checkpoints, `always_ask` gives the clearer invariant. For a custom tool, the enterprise application owns that invariant. Part 11 authority checks and Annex 11 access controls concern who may perform operations, not merely whether a model was interrupted. (health.ec.europa.eu)

An important edge case is historical data. Anthropic says `evaluation` may be absent for older events recorded before that field existed; missing evaluation can also signal an unavailable tool on a current denial. A consumer should use event time, top-level permission, agent version, and schema version to disambiguate. It should not infer that every absent `evaluation` means a disabled tool. Where source events are incomplete, classify the record as an exception and preserve the raw payload. (www.oecd.org)

Approval States and Evidence Design

An approval-needed `agent.tool_use` or `agent.mcp_tool_use` event is followed by `session.status_idle` with `stop_reason.type=requires_action`; `stop_reason.event_ids` lists the blockers. The session waits indefinitely. The client sends one `user.tool_confirmation` per blocking event, with the referenced `tool_use_id` and `allow` or `deny`; it can batch multiple confirmations in one events request. Only after all blockers resolve does the session return to `running`. If fewer than all are resolved, the API re-emits idle with the remainder. (platform.claude.com) (platform.claude.com)

This sequence implies an application state machine: **observed call → pending review → decision recorded → confirmation sent → result reconciled**. The pending state needs an owner, an escalation time, and a disposition path because waiting has no platform timeout in the documented flow. The decision should bind to the exact event ID and proposed arguments, not to a generic tool label. Once a call has `allow` or `deny`, sending a confirmation yields HTTP **400**; an auto denial cannot be overridden. A test should verify that the approval UI never offers a confirmation button for those states. (airc.nist.gov)

An audit evidence store should preserve the following fields, with retention and access rules set by the regulated record owner: (www.ecfr.gov) (www.oecd.org)

- **Identity and baseline:** session ID, agent ID and version, environment, toolset and per-tool policy, configuration hash, initiating service principal, and business workflow ID.
- **Attempt:** immutable tool-use event ID, event type, tool name, original arguments or protected reference, input digest, record IDs, and processed timestamp.
- **Evaluation:** top-level outcome, complete `evaluation` object, raw `reason_code`, parser version, and classification of unknown values.
- **Human decision:** reviewer identity, role, authorization check, decision, contemporaneous reason, decision timestamp, linked `user.tool_confirmation`, and any `deny_message`. (www.ecfr.gov)
- **Execution and reconciliation:** tool result, success or error, affected record version, downstream system audit ID, final session state, and duplicate-event resolution. (www.ecfr.gov)

This is a proposed data model, not a claim that every field is emitted by Anthropic. In particular, reviewer identity and GxP record linkage must come from the organization's systems. For electronic signatures where Part 11 applies, the regulation requires signature meaning, signer identification, and controls such as distinct identification components; an `allow` click should not be represented as a compliant signature unless the integrated process supplies the necessary controls. EU Annex 11 likewise expects electronic signatures to remain linked to their records. (www.ecfr.gov) (health.ec.europa.eu)

Event transport also affects evidence quality. Persisted session events include processing timestamps, while webhook notifications can arrive more than once and out of order. The webhook event identifies a resource but is not the complete source object. A collector can use the webhook to trigger retrieval, deduplicate by event ID, and compare its local sequence to the session event history. The reconciliation job should record gaps and replay outcomes. This is an implementation recommendation consistent with the platform's delivery semantics and with guidance that audit trails be reviewable. (health.ec.europa.eu) (picscheme.org)

Implementation Considerations and Process Changes

Match policy to the intended operation

Risk tiering should start with the effect of a **specific action** on an authoritative record, not the perceived safety of a general tool name. A browser read of public literature differs from a connector update to a study master file, even if both use an MCP server. ICH Q9(R1) calls for quality-risk-management effort and documentation proportionate to risk; FDA's draft AI credibility framework similarly ties assessment to a stated context of use. Neither source says that the Managed Agents `auto` evaluator performs a GxP risk assessment. (database.ich.org) (www.fda.gov)

Table 2 proposes a starting policy profile. It is a design hypothesis to test and approve locally, not a vendor default or a regulatory classification. (database.ich.org) (airc.nist.gov)

Example tier and operation	Proposed permission	Required external control	Pilot evidence
Read-only public literature retrieval	`auto` or `always_allow` after egress and source restrictions	Read scope, allowed domains, provenance and citation review	Inputs, fetched-source IDs, evaluation, output review. (platform.claude.com)
Draft generation in a segregated workspace	`auto` for bounded read/write to drafts; `always_ask` if the write has wider effect	Draft labeling, versioning, reviewer assignment	Draft version, source references, human review disposition.
Change to a controlled quality, clinical, or regulatory record	`always_ask` for the specific agent/MCP operation	Named role, least privilege, transaction validation, downstream audit trail	Exact proposed mutation, approver and reason, before/after record IDs. (www.ecfr.gov) (picscheme.org)
External submission, release, or other reserved business action	Keep the agent unable to execute; use the authorized system's workflow	Formal review and any required electronic signature in the system of record	Decision package, signature record, submission or release receipt. (health.ec.europa.eu) (www.ecfr.gov)

The read-only tier still needs boundaries: it should prevent a retrieval tool from reaching confidential sources beyond the approved scope, and the output should identify provenance. The draft tier needs a clear distinction between a reversible workspace artifact and a controlled record. A mutation tier cannot rely on `auto` as a universal human checkpoint, because an `auto` allow executes immediately. For reserved actions, the organization can disable the tool entirely or leave the final action in its validated business system. The matrix must be implemented as actual tool exposure, identities, and downstream rights, then confirmed by tests. (www.ecfr.gov) (www.oecd.org)

MCP and custom-tool coverage

MCP is the protocol boundary for server-executed tools in this permission system. An MCP toolset defaults to `always_ask`, but a per-tool configuration can alter its effective policy. A custom tool is different: permission events do not decide its execution. A custom connector that writes into an electronic document, laboratory, or quality application needs an application-level gate before the write, plus evidence from the destination system afterward. The organization should inventory each route to the same business operation; an `always_ask` MCP action is ineffective if an unrestricted custom tool can perform the same write. (platform.claude.com) (www.ecfr.gov) (picscheme.org)

The custom-tool gap checklist is practical:

- **Inventory:** list every custom action, destination system, record type, and service credential. (www.oecd.org)
- **Authorize:** verify a named person or service role against the requested operation and record state. (picscheme.org)
- **Approve:** bind approval to the exact input and record version before execution. (airc.nist.gov) (picscheme.org)
- **Execute once:** use transaction identifiers and idempotency to handle retries and reconnects. (www.oecd.org)
- **Reconcile:** connect application results to destination audit entries and retain the reason for the action. (picscheme.org) (www.ecfr.gov)

These are controls to design and validate, not built-in custom-tool features. The organization should also prevent a shared service account from erasing individual attribution at the system of record. MHRA and OECD guidance emphasize role-appropriate access and discourage generic logins where regulated data are generated or amended. (assets.publishing.service.gov.uk) (www.oecd.org)

Roles, configuration, and change management

The business owner defines which operations a workflow may perform; QA approves the regulated intended use and evidence plan; security controls credentials and destinations; the platform team implements and tests the policies. A reviewer of a controlled record should be separate from the agent identity that proposed the change when the procedure calls for independent verification. PIC/S guidance supports separating data generation from verification. ICH Q10 places ultimate responsibility for the pharmaceutical quality system with senior management; delegation of an approval task does not transfer that accountability to the agent platform. (picscheme.org) (database.ich.org) (airc.nist.gov)

Treat policy and tool inventory as versioned configuration. Capture the `default_config`, individual `configs`, enabled-tool set, agent version, and session override; test whether an idle-session update changed the full intended array. A change that adds a write-capable tool, changes `always_ask` to `auto`, or modifies a connector credential can change the approved control boundary. Revalidate those affected paths before exposing them to a regulated workflow. ISPE describes GAMP as a life-cycle approach rather than a one-time checklist, and NIST calls for repeatable testing and documented tools. (ispe.org) (airc.nist.gov)

Validation Tests and Pilot Gate

A useful qualification package has three layers: configuration inspection, event-path tests, and end-to-end business-record tests. Configuration inspection confirms that each agent and MCP tool has the intended default and override, that prohibited tools are absent, and that custom tools have their own gate. Event-path tests assert the exact transition and reason fields. Business-record tests verify identity, record versioning, downstream audit trails, and any signature workflow. This follows Part 11's emphasis on validated intended performance and Annex 11's call for documented risk assessment and validated applications. (www.ecfr.gov) (health.ec.europa.eu) (health.ec.europa.eu)

Minimum scripted cases should include:

- **Automatic allow:** allowed call executes once, produces the expected result, and has a retained evaluation. (picscheme.org)
- **Mandatory ask:** `always_ask` pauses before execution, exposes a blocker ID, and remains idle until a valid decision.
- **Indeterminate auto ask:** an `auto/ask` reason is retained; both human allow and human deny branches produce the expected result.
- **High-risk auto deny:** no execution occurs, an error result is recorded, and a confirmation attempt is rejected. (platform.claude.com)
- **Absent tool:** a deny without `evaluation` is recognized as a different branch from high-risk auto denial.
- **Partial and multiple blockers:** remaining IDs stay pending until resolved; duplicate or late messages are reconciled.
- **Unknown schema value:** parser preserves the raw payload and routes it for review without assuming `allow`. (airc.nist.gov)
- **Disconnect and webhook replay:** event order and duplicate notification do not create duplicate execution or lose a decision.
- **Custom-tool action:** the application authorization gate acts before execution and returns an attributable result. (www.ecfr.gov)
- **Policy change:** a new or updated session uses the intended tool array, while the old evidence remains tied to its actual configuration.

The pilot gate should be a documented decision by the workflow owner, QA, and security, based on observed cases and unresolved exceptions. NIST calls for documented test sets, metrics, and evaluation tools; ICH Q9(R1) calls for formality proportionate to risk. There is no public, primary-source pass percentage that converts these platform outcomes into GxP acceptance. A team should set acceptance criteria in advance for its own use case, then report failures and remediation without retroactively changing denominators. (airc.nist.gov) (database.ich.org) (www.fda.gov)

Data Analysis and Evidence

The most useful quantitative evidence for this decision is a **run-level denominator**, rather than search demand or a generic AI benchmark. Record every attempted agent or MCP call in the observed pilot window, including calls to absent or disabled tools, and separately record custom-tool attempts. Define `N` as all

attempted agent/MCP calls with a captured permission outcome. Let `A` be automatic allows under `always_allow` or `auto`; `H+` be human-approved asks; `H-` be human-denied asks; `D` be server-denied auto calls; and `M` be missing-tool denials. Let `P` be asks still pending at snapshot time. Reconcile `A + H+ + H- + D + M + P = N`, with each attempted call counted once. The equation is this report's measurement design; it should be tested against raw event IDs. (airc.nist.gov)

Table 3 gives the pilot dashboard specification. No counts are claimed here because no organization-specific runs were supplied. The example calculation shows how to report measured shares once a pilot exists. (airc.nist.gov) (database.ich.org)

Measure	Calculation and denominator	Interpretation
Automatic execution share	`A / N`	Exposure to execution before human approval; segment by tool and risk tier.
Human approval share	`H+ / N`	Calls actually allowed after an `ask`, not all calls reviewed later.
Human denial share	`H- / N`	Reviewed calls stopped by an authorized person; retain reasons.
Server denial share	`D / N`	Auto high-risk denials; examine attempted actions and recurrence.
Missing-tool share	`M / N`	Configuration or agent-behavior signal; do not merge with `D`.
Pending-ask share	`P / N`	Calls still awaiting confirmation at the snapshot; report the count and age separately.
Pending-age distribution	elapsed time from `requires_action` to confirmation, among resolved asks	Detect indefinite waits and escalation workload.
Evidence completeness	fully joined event packages / all attempted calls	Confirm traceability, record linkage, and reviewer attribution. (www.ecfr.gov) (www.oecd.org)

For illustration only (Hypothetical Example), if a pilot observed 200 calls, with 80 automatic allows, 70 human approvals, 20 human denials, 20 auto denials, and 10 missing-tool denials, the shares would be 40%, 35%, 10%, 10%, and 5%. Those are arithmetic examples, not measured product performance or suggested thresholds. The report should additionally show counts by tool, connector, workflow, agent version, and configuration version, because an aggregate can hide concentration in a single write-capable route. The denominator must exclude calls whose event was lost until reconciliation restores it; otherwise the rate is misleading. (airc.nist.gov) (www.oecd.org) (picscheme.org)

Claude Managed Agents pricing is secondary to control design but quantifiable. Anthropic bills Managed Agents for tokens and active session runtime; the published runtime line is **$0.08 per running session-hour** (platform.claude.com), while idle time waiting for approval is not counted as runtime. Web searches inside a session have a separate listed price of **$10 per 1,000 searches** (platform.claude.com). A pilot cost sheet should record token usage, active seconds, tool charges, and reviewer labor separately. It should also time

approval queues even when idle runtime is unbilled, because long waits are a process risk rather than a runtime-cost saving. Prices are vendor list rates as of September 2026 and should be checked before procurement.

Implications and Future Directions

The September permission change expands the range of workflows that can be **observed and gated**, but it does not collapse business authorization into a server risk code. For read and draft workflows, `auto` may reduce interruptions if the organization accepts automatic execution within explicit limits. For controlled record changes, `always_ask` and destination-system authorization offer a clearer human checkpoint. For custom tools, the application remains the enforcement point. A future product release could add new evaluation values, so parsers and validation tests should preserve unknown fields and fail into review rather than an implicit approval. (platform.claude.com)

The regulatory landscape itself is evolving. The European Commission consulted in 2025 on revisions to Annex 11 and a proposed Annex 22 for AI; proposal text refers to monitoring and human review, but a consultation is not the operative Annex 11. FDA's AI credibility guidance for drug and biological product regulatory decisions was also a draft as of this publication date. Organizations should version the requirements used for the pilot and revisit them when binding requirements or final guidance change. The current Annex 11 still calls for validation, access management, and audit trails; Part 11 remains the US electronic-records boundary where applicable. (health.ec.europa.eu) (www.fda.gov) (health.ec.europa.eu)

An evidence-first implementation can also improve governance beyond this vendor. A common event schema, record-linking method, reviewer authorization service, and exception queue can be reused across MCP and custom tools, provided each integration is validated for its own intended use. NIST treats human oversight as a defined process, OECD calls for traceability of AI processes and decisions, and ISPE frames computerized-system assurance as a life-cycle activity. The resulting control is the full regulated workflow, including the destination system and people, not a single permission flag. (airc.nist.gov) (www.oecd.org) (ispe.org) (picscheme.org)

Frequently Asked Questions (FAQs)

Can Claude Managed Agents provide a human-in-the-loop approval for every tool call?

Yes for an enabled agent or MCP tool configured as `always_ask`, subject to the customer's reviewer process. An `auto` tool does not guarantee that every call pauses: it may execute, pause, or be denied. Custom tools need an application approval gate. A permission confirmation is not automatically a Part 11 electronic signature. (platform.claude.com) (www.ecfr.gov)

Can an operator override an `auto` high-risk denial?

No. The documented `auto/deny` path cannot be overridden, and the API rejects `user.tool_confirmation` for a call whose outcome is not `ask`. A later configuration change is a separate controlled change, not an override of the denied event.

Does session connect replace a validated approval workflow?

The CLI can show the live transcript and let an operator allow or deny a waiting call, with `ant` version 1.32.0 or later. A regulated deployment still needs named identity, role checks, reason retention, record linkage, and reconciliation appropriate to its intended use. (www.ecfr.gov) (picscheme.org)

How should a team handle missing evaluation data?

A current missing-tool denial can have `deny` with no `evaluation`, while older events can also lack the object. Preserve the raw event, inspect time and effective configuration, and classify it explicitly. Do not silently treat absence as an auto denial or a permitted call. (www.oecd.org)

Conclusion

Claude Managed Agents' September 2026 controls provide a useful permission and event layer for agent and MCP tools. The decisive distinction is between a mandatory `always_ask` checkpoint and an `auto` policy that may run a call without review. Only `ask` outcomes can receive `user.tool_confirmation`; high-risk auto denials and missing-tool denials cannot be overridden through that event. Custom tools require application controls. (platform.claude.com)

For GxP use, the appropriate pilot is narrow and measurable. Define the intended operation and authoritative record, configure the exact tool exposure, bind human decisions to attempted calls, retain the complete event and business-record evidence, and test every allow, ask, deny, and exception path. Include reviewer identity and downstream audit records rather than treating a terminal prompt or platform reason code as a completed regulated approval. Monitor observed shares and unresolved pending cases with explicit denominators, then approve expansion only through the organization's quality and security process. (www.ecfr.gov) (www.ecfr.gov) (airc.nist.gov) (database.ich.org) (www.who.int)

This leaves a clear procurement and validation question: can the organization demonstrate, for each intended use, that its full workflow reliably restricts execution, attributes decisions, preserves records, and reconstructs changes? The product's documented permission fields can support that evidence. They do not replace the controls imposed by the relevant record system, governing procedure, or regulation. (health.ec.europa.eu) (picscheme.org) (www.iso.org)

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://platform.claude.com/docs/en/release-notes/overview
2. https://claude.com/blog/claude-managed-agents
3. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11
4. https://www.iso.org/home/insights-news/resources/iso-42001-explained-what-it-is.html
5. https://platform.claude.com/docs/en/managed-agents/permission-policies
6. https://intuitionlabs.ai/articles/audit-trail-requirements-ai-gxp-compliance
7. https://health.ec.europa.eu/system/files/2016-11/annex11_01-2011_en_0.pdf

8. https://airc.nist.gov/airmf-resources/airmf/5-sec-core/
9. https://www.gov.uk/government/publications/guidance-on-gxp-data-integrity
10. https://database.ich.org/sites/default/files/ICH_Q9%28R1%29_Guideline_Step4_2023_0126_0.pdf
11. https://www.who.int/news/item/28-06-2021-who-issues-first-global-report-on-ai-in-health-and-six-guiding-principles-for-its-design-and-use
12. https://intuitionlabs.ai/articles/21-cfr-part-11-compliance-ai-systems
13. https://intuitionlabs.ai/
14. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological
15. https://platform.claude.com/docs/en/cli-sdks-libraries/cli/sessions-connect
16. https://www.oecd.org/en/topics/ai-principles.html
17. https://www.oecd.org/en/publications/2021/09/glp-data-integrity_c2f067ec.html
18. https://picscheme.org/layout/document.php?id=714
19. https://www.oecd.org/content/dam/oecd/en/publications/reports/2021/09/glp-data-integrity_c2f067ec/45779212-en.pdf
20. https://assets.publishing.service.gov.uk/media/5aa2b9ede5274a3e391e37f3/MHRA_GxP_data_integrity_guide_March_edited_Final.pdf
21. https://database.ich.org/sites/default/files/Q10_Guideline.pdf
22. https://ispe.org/topics/gamp
23. https://platform.claude.com/docs/en/about-claude/pricing
24. https://health.ec.europa.eu/consultations/stakeholders-consultation-eudralex-volume-4-good-manufacturing-practice-guidelines-chapter-4-annex_en

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai