




INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

# Claude Science Private HPC Deployment for Pharma

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 · 2026-10-18 · [claude science](#) · [private hpc deployment](#) · [claude science ssh](#) · [aws systems manager](#) · [pharma research](#) · [least privilege](#) · [hpc security](#) · [ai agent governance](#)

Original article: <https://intuitionlabs.ai/articles/claude-science-private-hpc-pharma-deployment>



Claude Science Private HPC Deployment for Pharma

[intuitionlabs.ai](https://intuitionlabs.ai)

---

## In this article

1. [Executive Summary](#)
  2. [Introduction and Background](#)
  3. [Key Changes](#)
  4. [Implementation Considerations and Process Changes](#)
  5. [Least-Privilege Target Architecture](#)
  6. [Credentials and Approval Points](#)
  7. [Evidence, Qualification, and Response](#)
  8. [Data Analysis and Evidence](#)
  9. [Implications and Future Directions](#)
  10. [Conclusion](#)
- 

## Executive Summary

**Claude Science private HPC deployment** is an access-control and evidence-design problem. Anthropic introduced the workbench in beta on **June 30, 2026**, with access to local computers, remote machines through Secure Shell (SSH), and high-performance computing (HPC) login nodes. ([www.anthropic.com](http://www.anthropic.com)) The decisive boundary is the remote operating-system account: current documentation says remote jobs execute outside the local sandbox, with that account's host permissions. Keeping compute on a private cluster therefore requires independently enforced restrictions on identity, files, scheduler actions, and outbound traffic. ([claude.com](http://claude.com)) ([www.cyber.gov.au](http://www.cyber.gov.au))

For eligible workloads on Amazon Web Services (AWS), a useful candidate is **SSH over AWS Systems Manager (SSM)** to a designated private login instance. Session Manager supports access without opening inbound ports or maintaining a bastion. ([docs.aws.amazon.com](https://docs.aws.amazon.com)) This arrangement preserves separate transport and SSH authorization decisions; it does not supply command-content recording. AWS explicitly excludes SSH and port-forwarding session content from Session Manager logging, making scheduler, host, storage, and application records necessary complementary evidence. ([docs.aws.amazon.com](https://docs.aws.amazon.com)) ([slurm.schedmd.com](http://slurm.schedmd.com))

**Local datasets do not imply local-only processing.** Anthropic's current data documentation identifies prompts, read-file contents, code and command output, connector results, and other conversation context as model-request data. ([claude.com](http://claude.com)) The deployment decision must therefore classify permitted context and results, as well as the source datasets. Joint cyber-agency guidance recommends agentic AI only for low-risk, non-sensitive tasks. A restricted pilot using approved, non-sensitive inputs is the conservative starting point. Keep protected health information (PHI) out of Claude Science and exclude workflows requiring Business Associate Agreement (BAA) coverage: Anthropic does not provide that coverage for Claude Science. Organizational approval and private HPC connectivity do not override these restrictions. ([Anthropic data documentation](#)) ([www.cyber.gov.au](http://www.cyber.gov.au))

The operational recommendation is to separate proposal, execution, and result promotion. Use a dedicated, attributable identity; enforce Slurm account and partition limits; constrain storage and network access; and preserve the code, environment, parameters, jobs, and review decisions needed to reconstruct each result.

Slurm limits require configured enforcement, while containers do not automatically reduce the user's file permissions. ([slurm.schedmd.com](https://slurm.schedmd.com)) ([apptainer.org](https://apptainer.org)) Qualification should follow intended use and applicable obligations. The Food and Drug Administration's (FDA) **February 2026** computer software assurance guidance directly concerns medical-device production and quality-management software, rather than all pharmaceutical research. It offers a useful risk-based pattern, not a blanket pharmaceutical research and development (R&D) mandate. ([www.fda.gov](https://www.fda.gov))

## Introduction and Background

Research information technology (IT) teams considering **Claude Science on private infrastructure** need to decide what the agent may observe, initiate, modify, and export. The important question is whether the connected account's authority matches the approved scientific task. Current availability includes Pro, Max, Team, and Enterprise plans, with organizational enablement requirements on ordinary Team and Enterprise plans. Availability alone does not establish suitability for a particular research environment. ([claude.com](https://claude.com)) ([www.nist.gov](https://www.nist.gov))

Platform assumptions also need a dated baseline. As checked for this report on **October 18, 2026**, current documentation lists macOS 13 or later, Windows 11 x64, and glibc-based Linux x64. A search for “Claude Science Linux” or “Claude Science for Windows” should therefore lead to current installation requirements, rather than treating the launch announcement as a complete support matrix. Linux has additional sandbox prerequisites that administrators should verify against the installed release. ([claude.com](https://claude.com))

This report examines connection patterns, least privilege, evidence preservation, and qualification. It does not benchmark scientific accuracy. The proposed controls are architectural recommendations unless explicitly identified as documented product behavior. The National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework is voluntary guidance for AI design, use, and evaluation; it is useful for organizing decisions without implying certification of a deployment. ([www.nist.gov](https://www.nist.gov))

IntuitionLabs' perspective here is that of an adjacent life-sciences consultancy: the decision belongs jointly to research computing, security, scientific owners, and quality teams. Applicable requirements depend on intended use. Exploratory analysis and a workflow supporting a regulated record can need different evidence, review, and change controls. The Medicines and Healthcare products Regulatory Agency (MHRA) frames data governance across good practice, or GxP, sectors, while the relevant sector and jurisdiction still determine the deployment's obligations. ([www.gov.uk](https://www.gov.uk)) ([database.ich.org](https://database.ich.org))

## Key Changes

### The remote account defines the execution boundary

Claude Science reads existing SSH configuration, authenticates using a key or SSH authentication agent, and installs nothing on the host itself. That integration can reuse established research-computing access, but it also makes the selected account and its existing entitlements central to the assessment. ([claude.com](https://claude.com)) An administrator should inventory that identity's groups, mounted filesystems, credentials, scheduler associations, and network routes before registering the host. ([www.cyber.gov.au](https://www.cyber.gov.au))

The local sandbox is not a remote-cluster security policy. Current documentation describes deny-by-default local network access through an allowlist proxy. ([claude.com](#)) On the remote host, jobs run outside that sandbox. Administrators should enforce remote restrictions through operating-system permissions, scheduler configuration, service wrappers, and network controls, rather than expecting local application settings to propagate automatically. ([claude.com](#))

The scheduler boundary needs verification too. With Slurm's submission command available, the documented integration uses Slurm; otherwise, work can run as detached processes directly on the host. Other schedulers, including Portable Batch System (PBS) and Load Sharing Facility (LSF), are not supported by the documented submission integration. ([claude.com](#)) A pilot should explicitly demonstrate that substantial computation reaches an approved compute allocation. The National Energy Research Scientific Computing Center's (NERSC) published policy, for example, prohibits compute-intensive and memory-intensive applications on shared login nodes. That is a useful operating pattern, not a claim about every laboratory's local policy. ([docs.nersc.gov](#))

## Data residency has several layers

A source dataset may stay on cluster storage while a selected table, file excerpt, image, error message, or derived result enters model context. Anthropic explicitly states that file content read to answer a prompt is sent to Anthropic. ([claude.com](#)) A data-flow review should therefore ask what the workflow reads and displays, including tool output, rather than examining only explicit file uploads. ([cheatsheetseries.owasp.org](#))

Commercial training terms and retention also require separate decisions. Team and Enterprise customer content is excluded from training unless expressly permitted, whereas consumer plans follow their own terms and settings. ([claude.com](#)) The data documentation distinguishes local records, model-call logs, and compliance transcripts; it also states that Claude Science does not provide inference in the customer's own cloud tenancy or through a third-party cloud platform. A private HPC connection should not be described as private model hosting. ([claude.com](#))

## Approval and connector assumptions need testing

Current documentation says Claude **might** propose a plan before multi-step work. It also describes permission cards and revocable standing permissions. A deployment should test which specific operations trigger approval after permissions have been saved; a conversational plan is not sufficient evidence of a mandatory authorization gate. ([claude.com](#)) ([cheatsheetseries.owasp.org](#))

BioTeam's current repository still describes its SSM approach as a workaround and native SSM as a direction of travel. Official documentation opened for this report established SSH integration but did not independently establish native SSM availability. Treat the published workaround as a version-specific implementation example and verify the installed connector before relying on it. ([github.com](#))

# Implementation Considerations and Process Changes

## Choose the access pattern from the trust boundary

**Native SSH** can fit an existing private network path, including a virtual private network (VPN), where the client can reach a controlled login host. **ProxyJump** introduces a jump host that forwards traffic to the final destination; administrators must configure the jump host and destination appropriately. OpenSSH notes that destination settings do not generally configure jump hosts automatically. ([man.openbsd.org](https://man.openbsd.org))

For AWS-hosted clusters, **SSH over SSM** moves the transport into an identity-authorized management channel. AWS documents a ProxyCommand using its command-line interface and the AWS-StartSSHSesssion document. This remains SSH inside the tunnel, with a separate SSH authentication decision.

([docs.aws.amazon.com](https://docs.aws.amazon.com)) ([github.com](https://github.com))

*Table 1* compares deployment patterns. Cost descriptions identify operational components to budget; they are not measured price comparisons.

| Pattern                                | Access surface and operability                                                                                                                                    | Evidence and cost drivers                                                                                                                                                                                 | Principal condition to test                                                                                                                          |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Direct SSH over a private route or VPN | A reachable SSH service remains necessary. SSH provides encrypted transport and user authentication.<br>( <a href="https://www.openssh.org">www.openssh.org</a> ) | Budget private-network operation, identity administration, and host logging.                                                                                                                              | Denied users and unintended destinations remain unreachable; host identity is verified.<br>( <a href="https://man.openbsd.org">man.openbsd.org</a> ) |
| SSH through a bastion                  | ProxyJump adds a forwarding host and another managed access boundary.<br>( <a href="https://man.openbsd.org">man.openbsd.org</a> )                                | Budget bastion operation, patching, availability, and correlation of hop records.                                                                                                                         | Jump-host rights do not unintentionally provide broader downstream access.                                                                           |
| SSH over SSM                           | Session Manager supports access without inbound ports or a maintained bastion.<br>( <a href="https://docs.aws.amazon.com">docs.aws.amazon.com</a> )               | Budget management-agent operation, endpoints, identity controls, and independent command evidence. SSH content is not logged by SSM.<br>( <a href="https://docs.aws.amazon.com">docs.aws.amazon.com</a> ) | IAM and SSH permissions are both enforced; loss of either prevents access.                                                                           |
| Disconnected, manual promotion         | An approved operator executes reviewed material without giving the agent a live cluster credential. This is a proposed alternative.                               | Budget operator effort, transfer review, and reconciliation of executed files.                                                                                                                            | Exported context and returned results remain within the approved information policy.                                                                 |

The connection pattern should follow workload eligibility and evidence requirements. SSM can remove an inbound-access requirement, but access to the management service, SSH authentication, remote authorization, and scientific review remain separate design responsibilities. The disconnected alternative deserves consideration when the permitted context cannot include sensitive material, consistent with the conservative scope of joint cyber-agency guidance. ([www.cyber.gov.au](https://www.cyber.gov.au)) ([www.nist.gov](https://www.nist.gov))

## Establish and test the complete path

The following sequence is a proposed implementation procedure, based on the documented AWS and SSH mechanisms:

- **Identify the destination:** Register the intended private login instance and confirm its lifecycle, ownership, host identity, and replacement procedure.
- **Separate caller and node roles:** Distinguish permission to start a session from the instance profile that allows the managed node to participate in Systems Manager. ([docs.aws.amazon.com](https://docs.aws.amazon.com))
- **Provide service reachability:** Use appropriate private endpoints where required. AWS documents private-IP access to Systems Manager APIs; without the relevant endpoints, its guidance calls for outbound Hypertext Transfer Protocol Secure (HTTPS) on port 443. ([docs.aws.amazon.com](https://docs.aws.amazon.com))
- **Verify transport prerequisites:** Check the installed management agent, client plugin, command-line interface, and SSH client before adding the application.
- **Verify SSH trust:** Pre-provision trusted host keys and require strict checking. OpenSSH can refuse connections when a host key changes. ([man.openbsd.org](https://man.openbsd.org))
- **Test from the application:** Confirm that the desktop application's environment, profile selection, and SSH configuration produce the same intended destination as the approved terminal test.

AWS documents minimum versions of **SSM Agent 2.3.672.0** and **Session Manager plugin 1.1.23.0** for this SSH pattern. These are compatibility floors from the documentation, not recommendations to deploy old versions. ([docs.aws.amazon.com](https://docs.aws.amazon.com)) Also inspect configuration precedence: OpenSSH says the first specified ProxyCommand or ProxyJump can prevent the other from taking effect. An inherited configuration should be tested as a complete configuration, not assembled by assumption. ([man.openbsd.org](https://man.openbsd.org))

## Least-Privilege Target Architecture

### A constrained execution lane

For an eligible AWS pilot, the proposed target is a private login instance reached through scoped SSM authorization, a separate attributable operating-system identity, a dedicated Slurm account and partition, restricted project storage, and controlled egress. AWS Identity and Access Management (IAM) policies can specify both the managed nodes reachable by an identity and the Session Manager operations allowed. ([docs.aws.amazon.com](https://docs.aws.amazon.com)) The design should additionally limit the usable session documents and prohibit unnecessary forwarding capabilities.

BioTeam's example deserves close reading: its IAM policy permits **one instance and three AWS-managed documents**, including forwarding documents. Copying the example unchanged would therefore authorize more than an SSH-only interpretation suggests. Review each document against the intended use and retain only capabilities the workflow needs. ([github.com](https://github.com)) This is a configuration decision, not a reason to assume that an instance-scoped policy alone constrains every reachable service.

A practical least-privilege specification should state each identity's host, scheduler action, path, and network scope:

- **Researcher identity:** May authorize an approved task and review results; does not share a personal SSH key with other users.
- **Agent execution identity:** May reach the designated login host and execute the approved workflow; receives no administrative elevation.

- **Scheduler association:** Binds the execution user to the intended cluster, account, and partition. Slurm associations explicitly include these dimensions. ([slurm.schedmd.com](https://slurm.schedmd.com))
- **Input zone:** Presents approved datasets read-only and excludes unrelated projects, home-directory secrets, and general shared storage.
- **Working zone:** Permits project-specific scratch writes with quota, ownership, cleanup, and job-to-directory mapping.
- **Result zone:** Accepts candidate outputs; a separate review identity promotes accepted results into the authoritative repository.
- **Network scope:** Permits named dependencies and approved result destinations; rejects arbitrary outbound transfer and access to unnecessary internal services. ([cheatsheetseries.owasp.org](https://cheatsheetseries.owasp.org))

These are proposed controls implementing the general principle of granting only task-required permissions. OWASP likewise recommends restricting tool-server filesystem access to required directories and disabling network access unless needed. Its Model Context Protocol (MCP) guidance informs the principle; it is not evidence that every Claude Science connector uses MCP. ([docs.aws.amazon.com](https://docs.aws.amazon.com)) ([cheatsheetseries.owasp.org](https://cheatsheetseries.owasp.org))

## Scheduler, container, and filesystem enforcement

Set scheduler-side concurrency, outstanding-submission, wall-time, and resource limits. **MaxJobs**, **MaxSubmitJobs**, and **MaxWallDurationPerJob** address distinct dimensions, so limiting running jobs alone does not bound the submission queue. ([slurm.schedmd.com](https://slurm.schedmd.com)) Verify AccountingStorageEnforce settings: Slurm documents that configured association limits can otherwise remain unenforced. ([slurm.schedmd.com](https://slurm.schedmd.com))

A container helps capture a software environment, but should not be assumed to reduce the account's authority. Apptainer documents that the user remains the same inside and outside the container, and that container applications can directly access files the user owns. Use deliberate bind mounts and independently enforced file permissions. ([apptainer.org](https://apptainer.org))

Shared scratch also requires an explicit boundary. BioTeam's setup calls for a writable path visible to every compute node. Place that shared path within the approved project's writable area, rather than granting a broad shared-filesystem root. ([github.com](https://github.com)) Test read-only inputs and cross-project denials from inside the actual job environment. ([apptainer.org](https://apptainer.org))

## Credentials and Approval Points

### Separate identities and credential lifecycles

Cloud authorization, SSH login, scheduler identity, and scientific approval should have distinct responsibilities. AWS recommends temporary credentials and IAM roles for workloads where supported. ([docs.aws.amazon.com](https://docs.aws.amazon.com)) BioTeam's published guide, however, describes a Claude Science AWS credential field that accepts a static key pair without a session-token field. That is implementation-maintainer evidence about its documented setup, rather than a guarantee about every current installation. ([github.com](https://github.com))

If the installed workflow cannot support the organization's credential-lifetime policy, use a different access method or retain manual execution. An exception should identify its owner, restricted resources, rotation process, and termination test. Do not describe short-lived authentication as implemented merely because it appears in a target-state diagram. ([www.cyber.gov.au](http://www.cyber.gov.au))

For regulated records, distinguish a dedicated workload account from an anonymous shared login. Preserve the initiating person's identity, approval, session identifier, and job linkage. MHRA advises against shared logins or generic access for systems generating, changing, or storing GxP data. A service account requires an attribution design and applicable quality assessment; naming it “agent” does not satisfy that requirement. ([assets.publishing.service.gov.uk](http://assets.publishing.service.gov.uk)) Joint cyber-agency guidance similarly recommends distinct agent identities and permissions restricted to the approved task. ([www.cyber.gov.au](http://www.cyber.gov.au))

## Approve concrete actions

The proposed approval policy should require enough information to understand the operation. OWASP recommends displaying full tool-call parameters when requesting confirmation for sensitive actions. A friendly summary should be accompanied by the actual target, files, command, and material side effects. ([cheatsheetseries.owasp.org](http://cheatsheetseries.owasp.org))

- **New host:** Approve the destination, account, trust fingerprint, network route, and project scope before registering access.
- **Job submission:** Approve the script, input manifest, partition, account, resource request, expected runtime, and output location.
- **Job cancellation:** Verify ownership and the exact job identifiers; distinguish cancellation from deleting output.
- **Data export:** Approve destination, content classification, selected files, permitted metadata, and any required redaction.
- **Package installation:** Approve the source, version, environment location, and reproducibility impact before changing the execution environment.
- **External network call:** Approve the service and transmitted material, including command output and diagnostic context.
- **Result promotion:** Require an accountable reviewer to accept candidate outputs into the authoritative scientific record. ([assets.publishing.service.gov.uk](http://assets.publishing.service.gov.uk))

These are recommended approval points to implement and test; the report does not claim that the application natively enforces this entire policy. Authorization should remain effective even if text in a document or tool response instructs the model to bypass it. OWASP treats tool responses, including those from approved servers, as untrusted data and recommends independent authorization in trusted code.

([cheatsheetseries.owasp.org](http://cheatsheetseries.owasp.org))

# Evidence, Qualification, and Response

## Preserve the chain from request to result

Application history is useful evidence, but completeness must be demonstrated for the intended workflow. Anthropic documents beta Compliance application programming interface (API) coverage for read-only session transcripts. Coverage, retention, export, and linkage to executed work still need testing against the organization's record requirements. ([claude.com](https://claude.com)) ([assets.publishing.service.gov.uk](https://assets.publishing.service.gov.uk))

AWS CloudTrail captures Session Manager API activity, while SSH-over-SSM content logging is unavailable. Those statements describe different evidence layers, and should appear separately in the design.

([docs.aws.amazon.com](https://docs.aws.amazon.com)) Slurm accounting can report running or terminated job resource usage when the relevant collection and storage components are configured. It must be deliberately enabled and retained. ([slurm.schedmd.com](https://slurm.schedmd.com))

Table 2 is a proposed data-flow register. Each retention value is a policy decision to document; none should be inferred from “the data stays local.”

| Artifact and route                                                              | Sensitivity and retention decision                                                                                                               | Protection to verify                                                           | Evidence owner                 |
|---------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|--------------------------------|
| Approved input: governed storage to remote job                                  | Dataset classification and authorized project retention                                                                                          | Read-only mount, restricted account, storage encryption, object-access records | Data owner and storage team    |
| Prompt, read content, command output: application to model service              | Review all context allowed outside the lab; Anthropic identifies these as model-request content. ( <a href="https://claude.com">claude.com</a> ) | Verify contracted handling, permitted fields, transport, and redaction         | Security and application owner |
| Credential: identity system to client or application store                      | Record credential type, expiry, rotation, and revocation                                                                                         | Verify secret storage, local device controls, and absence from prompts or logs | Identity owner                 |
| Candidate result: remote job to staging or application                          | Classify derived results and metadata before transfer                                                                                            | Approved destination, integrity checks, and access restrictions                | Scientific owner               |
| Transcript, job, environment and review records: systems to evidence repository | Define record classes, retention, deletion holds, and reviewer access                                                                            | Time correlation, protected exports, hashes, and tested retrieval              | Platform and quality owners    |

The register should be completed using actual configuration and contracts. Claude Science does not offer zero data retention (ZDR); exclude workflows requiring ZDR, even with organizational approval or private HPC connectivity. ([Anthropic data documentation](#)) For example, the data documentation describes **30-day retention for designated Covered Model requests**, but that is not a universal retention promise for local files, all model logs, or compliance transcripts. Record classes and applicable exceptions must be checked individually. ([claude.com](https://claude.com)) MHRA describes audit trails as metadata supporting reconstruction of record changes, which is a broader requirement than retaining an application's conversation screen. ([assets.publishing.service.gov.uk](https://assets.publishing.service.gov.uk))

For each accepted analysis, the proposed evidence bundle should include:

- **Request and approval:** Intended use, initiating person, permitted inputs, decision record, and session linkage.
- **Executed material:** Code revision, submitted script, parameters, input manifest, dependency sources, and environment or container digest.
- **Execution evidence:** Host identity, scheduler job and array identifiers, timestamps, exit status, resource usage, and relevant storage events.
- **Scientific review:** Acceptance criteria, comparison results, deviations, reviewer identity, and result-promotion decision.
- **Preserved outputs:** Candidate and accepted artifacts, integrity hashes, repository location, and retention classification.

Reproducible Builds recommends defining or recording relevant environment information. Applying that principle here supports reconstruction; its bitwise software-build definition does not establish that stochastic or floating-point scientific workflows will produce identical outputs. Define scientifically meaningful tolerances where appropriate. ([reproducible-builds.org](https://reproducible-builds.org))

## Qualify intended use and challenge failure paths

The qualification package should begin with workflow boundaries, risk classification, supplier assessment, installation checks, operational challenges, and scientific performance acceptance criteria. The Organisation for Economic Co-operation and Development's (OECD) good laboratory practice (GLP) guidance uses a lifecycle approach to computerized-system validation and operation. ([www.oecd.org](https://www.oecd.org)) FDA's computer software assurance guidance offers a risk-based assurance pattern within its specific device-production and quality-system scope. ([www.fda.gov](https://www.fda.gov))

*Table 3* proposes negative tests and the evidence expected from a restricted pilot. Expected outcomes are design requirements to demonstrate, rather than reported product test results.

| Challenge                             | Expected result                                                                                     | Evidence to retain                                           |
|---------------------------------------|-----------------------------------------------------------------------------------------------------|--------------------------------------------------------------|
| Credential expiry or revocation       | New access fails; handling of established sessions and queued work follows the documented procedure | Identity events, session termination evidence, affected jobs |
| Denied host or changed host key       | Connection is rejected without silently accepting a replacement                                     | SSH diagnostics and denied destination                       |
| Denied path or cross-project input    | Read or write fails without exposing the restricted content                                         | Filesystem decision, job output, access record               |
| Malicious instructions in tool output | No unauthorized export, credential use, or command occurs                                           | Full parameters, denied action, transcript                   |
| Job timeout or resource excess        | Scheduler-enforced limit applies and incomplete output remains unpromoted                           | Scheduler state, exit status, partial-output handling        |
| Network loss during submission        | Submission status is reconciled before retry                                                        | Submission identifier and scheduler reconciliation           |

| Challenge                          | Expected result                                                                | Evidence to retain                                    |
|------------------------------------|--------------------------------------------------------------------------------|-------------------------------------------------------|
| Duplicate submission               | An idempotency record prevents accidental repeated work or flags it for review | Request identifier, job mapping, reviewer decision    |
| Connector or configuration upgrade | Critical access, denial, approval, and evidence tests pass before broader use  | Version baseline, regression results, change approval |

Interpret results against the workflow's consequence and intended use. International Council for Harmonisation (ICH) Q9(R1) says risk-management effort, formality, and documentation should be commensurate with risk. ([database.ich.org](https://database.ich.org)) European Commission Annex 11's published text distinguishes application validation from infrastructure qualification and bases supplier-audit need on risk assessment. Check the applicable adopted text and local obligations when building the quality plan. ([health.ec.europa.eu](https://health.ec.europa.eu)) MHRA also notes that OECD guidance takes precedence for United Kingdom (UK) GLP data governance, with MHRA guidance supplementary. ([www.gov.uk](https://www.gov.uk))

## Respond without losing attribution

The proposed response procedure is to suspend new actions, revoke affected credentials, terminate applicable sessions, restrict the execution lane, and preserve records before cleanup. Identify touched datasets, running and queued jobs, exported results, and promoted records. Quality and scientific owners should assess whether affected outputs need review or invalidation. ([www.oecd.org](https://www.oecd.org))

Credential revocation alone should not be assumed to stop existing computation. Slurm documents that reducing limits does not cancel already running jobs. Test and document the separate operations needed to stop sessions and workloads, then reconcile every affected job with its outputs. ([slurm.schedmd.com](https://slurm.schedmd.com)) Record the review and its outcome: MHRA expects documented data review, including relevant metadata and audit trails according to risk. ([assets.publishing.service.gov.uk](https://assets.publishing.service.gov.uk))

## Data Analysis and Evidence

### Quantify configured authority

This report does not present a measured deployment-risk probability, scientific benchmark, adoption estimate, or cost-saving figure. The quantitative evidence concerns documented configuration dimensions and defaults. These are useful for prioritizing controls, but cannot establish the safety of a scientific result.

Slurm associations provide **four dimensions**: cluster, account, user, and optional partition. That structure supports explicit access reviews, provided the relevant enforcement settings are enabled. ([slurm.schedmd.com](https://slurm.schedmd.com)) Separate limits on running jobs, outstanding submissions, and wall time should be recorded in the deployment baseline and challenged independently. ([slurm.schedmd.com](https://slurm.schedmd.com))

Current Claude Science documentation specifies a default **100-job concurrency limit**, **30-minute job timeout**, and approximately **100 megabyte (MB) output-transfer threshold**, with larger files left on the host. These are application defaults to verify in the installed release, rather than a substitute for scheduler capacity limits or an information-export policy. A small file may still contain sensitive information, while a large file left behind may still be summarized into context. ([claude.com](https://claude.com))

AWS documents a default Session Manager idle timeout of **20 minutes**, configurable between **1 and 60 minutes**. ([docs.aws.amazon.com](https://docs.aws.amazon.com)) Maximum session duration is a different control: it is unset by default and can be configured between **1 and 1,440 minutes**. ([docs.aws.amazon.com](https://docs.aws.amazon.com)) Test their actual behavior for the selected SSH transport and client versions, including established sessions and independent queued work. An idle timer and credential expiry should not be treated as interchangeable.

## A bounded prioritization heuristic

**Hypothetical Example:** define an exposure index as reachable hosts multiplied by permitted action classes multiplied by credential lifetime in hours. A broad configuration with **6 hosts × 5 action classes × 8 hours = 240 host-action-hours** can be compared with **1 host × 3 action classes × 1 hour = 3 host-action-hours**. The arithmetic gives an **80-fold difference** in this deliberately simplified index.

These are analyst-selected inputs and arithmetic, not observed deployment measurements. The index is neither a validated risk model nor a probability: it ignores data value, action severity, alternate credentials, network topology, established sessions, and scientific consequence. Its purpose is to make reviewers ask why an identity needs another host, action, or hour. AWS's least-privilege guidance supports examining actions, resources, and conditions separately; it does not validate this formula. ([docs.aws.amazon.com](https://docs.aws.amazon.com)) ([www.nist.gov](https://www.nist.gov))

For an actual pilot, measure denied actions, approval bypass attempts, orphaned jobs, missing evidence fields, and successful reconstruction against predefined acceptance criteria. Keep the denominators and test scope visible. A favorable test result supports the tested configuration and intended use; it does not justify broadening access without reassessment, consistent with OECD's risk-centered lifecycle approach. ([www.oecd.org](https://www.oecd.org))

## Implications and Future Directions

The deployment decision should have a written eligibility threshold before infrastructure work begins. Joint cyber-agency guidance recommends low-risk, non-sensitive uses; it should not be represented as an endorsement of unrestricted agent access to sensitive research systems. ([www.cyber.gov.au](https://www.cyber.gov.au)) A practical first phase is therefore an approved non-sensitive workflow, constrained execution identity, manual promotion, and demonstrated evidence retrieval.

Broader use should depend on observed control performance and the consequences of an incorrect or unauthorized action. NIST's voluntary framework can organize ownership, measurement, and review across the AI lifecycle. ([www.nist.gov](https://www.nist.gov)) Within regulated activities, the quality system should determine the applicable assurance package rather than treating a generic AI checklist as sufficient. OECD's computerized-system guidance likewise frames validation and operation as lifecycle activities. ([www.oecd.org](https://www.oecd.org))

Connector changes warrant particular attention. BioTeam's repository describes an evolving SSM approach, so a future native implementation could change credential handling, transport setup, or evidence collection. ([github.com](https://github.com)) Re-run the critical denial and reconstruction tests whenever those boundaries change. A simpler connection experience should be evaluated for its actual identity and logging behavior.

The most useful future product capabilities would make authorization and evidence easier to verify: clearly bounded temporary credentials, explicit remote-action policies, exportable execution manifests, and reliable links between sessions and scheduler jobs. These are desired capabilities, not claims that the present product provides them. Independently enforced authorization remains valuable even as application interfaces evolve; OWASP's guidance makes that separation explicit for tool-mediated actions. ([cheatsheetseries.owasp.org](https://cheatsheetseries.owasp.org))

A deployment review should end with a concrete record:

- **Eligibility:** Identify the approved datasets, scientific tasks, and prohibited information flows.
- **Authority:** Record accountable identities, destinations, scheduler limits, filesystem permissions, and credential lifetimes.
- **Evidence:** Assign owners for transcripts, execution records, storage events, environment capture, and review decisions.
- **Acceptance:** Define the successful and denied operations that demonstrate the intended boundary.
- **Change:** Specify review intervals and the configuration, connector, or use-case changes that trigger reassessment.

That record gives research, security, and quality teams something measurable to approve and periodically reassess. ([assets.publishing.service.gov.uk](https://assets.publishing.service.gov.uk))

## Conclusion

A useful Claude Science private HPC design begins with an approved information boundary and a narrowly defined scientific task. The cluster connection is one part of that design. Account permissions, scheduler enforcement, file access, outbound destinations, approvals, and result review determine what authority the agent actually receives.

For an eligible AWS workload, SSH over SSM is a candidate architecture when the organization wants private access without an inbound management surface. Its suitability depends on independently proving the transport, authentication, authorization, and evidence paths. A private route does not by itself make model context private to the laboratory, and an application transcript does not by itself reconstruct an entire analysis.

The strongest pilot is small enough to explain and test. It uses approved inputs, an attributable identity, a constrained execution lane, documented failure behavior, and a reviewer-controlled path into the scientific record. It demonstrates denied actions as well as successful jobs, and preserves the material needed to reproduce or assess each accepted output.

Where those boundaries cannot be enforced, disconnected execution and manual result promotion remain legitimate design choices. Expansion should follow evidence that the configured controls work for the intended use. The deployment should be judged by what it permits, prevents, and records, with the same discipline applied again when credentials, connectors, datasets, or scientific purpose change.

## Source links

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

1. <https://www.anthropic.com/news/claude-science-ai-workbench>
2. <https://claude.com/docs/claude-science/remote-compute-clusters>
3. <https://www.cyber.gov.au/publication/careful-adoption-of-agentic-ai-services>
4. <https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>
5. <https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-getting-started-enable-ssh-connections.html>
6. <https://slurm.schedmd.com/accounting.html>
7. <https://claude.com/docs/claude-science/how-claude-science-works-with-your-data>
8. <https://apptainer.org/docs/user/latest/introduction.html>
9. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/computer-software-assurance-production-and-quality-management-system-software>
10. <https://claude.com/docs/claude-science/overview>
11. <https://www.nist.gov/itl/ai-risk-management-framework>
12. <https://www.gov.uk/government/publications/guidance-on-gxp-data-integrity>
13. [https://database.ich.org/sites/default/files/ICH\\_Q9%28R1%29\\_Guideline\\_Step4\\_2025\\_0115\\_0.pdf](https://database.ich.org/sites/default/files/ICH_Q9%28R1%29_Guideline_Step4_2025_0115_0.pdf)
14. <https://claude.com/docs/claude-science/core-concepts>
15. <https://docs.nersc.gov/policies/resource-usage>
16. [https://cheatsheetseries.owasp.org/cheatsheets/MCP\\_Security\\_Cheat\\_Sheet.html](https://cheatsheetseries.owasp.org/cheatsheets/MCP_Security_Cheat_Sheet.html)
17. <https://github.com/bioteam/claude-science-hpc-integrations>
18. [https://man.openbsd.org/ssh\\_config](https://man.openbsd.org/ssh_config)
19. <https://github.com/bioteam/claude-science-hpc-integrations/blob/main/iam-user-for-ssm-sessions/README.md>
20. <https://www.openssh.org/manual.html>
21. <https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-getting-started-instance-profile.html>
22. <https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-getting-started-privatelink.html>
23. <https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-getting-started-restrict-access.html>
24. [https://slurm.schedmd.com/resource\\_limits.html](https://slurm.schedmd.com/resource_limits.html)
25. <https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html>
26. <https://github.com/bioteam/claude-science-hpc-integrations/blob/main/Connecting-Claude-Science-to-AWS.md>
27. [https://assets.publishing.service.gov.uk/media/5aa2b9ede5274a3e391e37f3/MHRA\\_GxP\\_data\\_integrity\\_guide\\_March\\_edited\\_Final.pdf](https://assets.publishing.service.gov.uk/media/5aa2b9ede5274a3e391e37f3/MHRA_GxP_data_integrity_guide_March_edited_Final.pdf)
28. <https://reproducible-builds.org/docs/recording>
29. [https://www.oecd.org/en/publications/application-of-glp-principles-to-computerised-systems\\_e3c9cda5-en.html](https://www.oecd.org/en/publications/application-of-glp-principles-to-computerised-systems_e3c9cda5-en.html)
30. [https://health.ec.europa.eu/document/download/8d305550-dd22-4dad-8463-2ddb4a1345f1\\_en?filename=annex11\\_01-2011\\_en.pdf&prefLang=sl](https://health.ec.europa.eu/document/download/8d305550-dd22-4dad-8463-2ddb4a1345f1_en?filename=annex11_01-2011_en.pdf&prefLang=sl)
31. <https://docs.aws.amazon.com/systems-manager/latest/userguide/session-preferences-timeout.html>
32. <https://docs.aws.amazon.com/systems-manager/latest/userguide/session-preferences-max-timeout.html>

## 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>

---

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.