

Claude Science product guide

A practical deployment guide

[Chapter] Foreword

Life sciences organizations are at a turning point in their AI adoption journeys, and researchers on the ground are the ones responsible for putting this new technology to work.

Deloitte's [2026 Life Sciences Outlook](#), a survey of 280 biopharma and medtech leaders, found that 78% expect AI to play a central role in driving major change this year—yet only 14% report full implementation of AI tools into daily workflows, with another 40% still working toward it. Anthropic's own internal research, drawn from interviews with researchers across chemistry, physics, biology, and computational fields, found that 91% of scientists want more AI in their research, while 79% named trust and reliability as their number-one barrier to adoption.

A biologist's day spans literature, experiment design, data wrangling, analysis, figures, and writing, but the tools span PubMed, Jupyter, R, a cluster terminal, scientific data renderers, spreadsheets, and more. The labs and R&D orgs pulling ahead aren't the ones with the most compute or the biggest bioinformatics teams, they're the ones who've collapsed the distance between a scientist's question and a defensible result. Everyone else is still paying a toil tax on every experiment, and the gap only compounds.

[Claude Science](#) (in beta) is Anthropic's answer to this problem: an application for every digital step of life science, built to run next to the scientist's data and produce results that can be traced, reproduced, and defended. It sits inside a broader Claude product family—including Claude Chat, Claude Cowork, Claude Code, Claude for Microsoft 365, the Claude Platform, and Claude Managed Agents—that life sciences organizations like [Novo Nordisk](#), [the Garvan Institute](#), and [Benchling](#) use for the document, regulatory, and enterprise work that surrounds the science. This guide covers which surface to reach for when, then goes deep on deploying Claude Science inside a research organization, with customer examples and a rollout roadmap. Claude Science is intended for research use and is not designed for clinical or diagnostic decision-making. Items highlighted as being on Anthropic's roadmap are subject to change and do not represent a commitment.

Let's dive in.

[Chapter] When to use which Claude surface

Claude works for life sciences teams in several ways. While Claude Science is the surface built for scientists, other Claude products cover the document, software, and enterprise work that surrounds your team's scientific discovery. Most organizations deploy more than one.

Claude Science is an application for the digital steps of research, including literature review, experiment design, data analysis, figure generation, and writeups. It runs locally next to the lab's data and compute, ships with domain capabilities across genomics, single cell, proteomics, structural biology, cheminformatics, and more, and tracks the provenance of every artifact so results can be reproduced and defended. *Covered in depth in the next chapter.*

Claude Chat is the web and desktop chat interface, available at Claude.ai, for quick queries and drafting. A medical writer might use it to pressure-test a mechanism-of-action narrative, while a discovery scientist might use it to summarize a stack of papers ahead of a program review.

Claude Cowork is a desktop application where Claude works across local files and connected systems—Benchling, Veeva, Microsoft 365, PubMed—to complete multi-step document projects. Reach for it for study- and submission-level work that spans folders and apps: reviewing every site monitoring report in a study folder, reconciling a TMF section, drafting a CSR section from a folder of TLFs.

Claude Code is our command-line interface for software engineering teams, supporting scientific computing, biostatistics programming, and platform groups building and maintaining production pipelines and internal tools under version control. Use Claude Code when the output is software that ships to other teams; reach for Claude Science when the output is an analysis, a figure, or a result.

Claude for Microsoft 365 puts Claude inside Word, Outlook, Excel, and PowerPoint via add-ins, and lets Claude search Outlook, Teams, SharePoint, and OneDrive via the M365 connector. A regulatory writer can redline a Module 2 summary with tracked changes and turn the gap analysis into a slide without leaving the app.

Claude Platform is Anthropic's API for organizations embedding Claude into ELN, LIMS, CTMS, safety, and RWE systems. **Claude Managed Agents** runs any agent built on the Platform as a hosted service with long-running sessions, scoped permissions, and full execution tracing—for example, a Literature Surveillance agent monitoring PubMed and ClinicalTrials.gov for competitive readouts across a portfolio.

Claude product matrix: when to use what

Surface	Best for	Primary users	Where it runs	Example task
Claude Science	End-to-end research analysis with	Computational biologists, bioinformaticia	Local app (macOS, Linux); dispatches to	Run QC and clustering on this scRNA-seq

Surface	Best for	Primary users	Where it runs	Example task
	full provenance—literature, pipelines, figures, writeup	ns, research scientists, PIs	SSH, SLURM, or cloud compute	dataset, show me the UMAP, then subset to immune cells and redo the marker analysis.
Claude Chat	Conversational drafting, day-to-day questions and answers in a chat interface	All staff	Browser, desktop, mobile	Summarize this document and flag any contradictory findings.
Claude Cowork	Cross-app study and document work that touches files and multiple systems	Medical writing, clinical operations, regulatory, medical affairs	Claude desktop app	Review the SAE narratives in this folder against the protocol and flag any that meet expedited reporting criteria.
Claude Code	Agentic software engineering inside a repository	Scientific computing, biostatistics programming, platform engineering	Terminal, IDE	Migrate or refactor large codebases
Claude for Microsoft 365	In-place drafting, redlining, and review across the Microsoft suite	Medical writing, regulatory, clinical operations	Word, Outlook, Excel, PowerPoint (add-ins); Teams, SharePoint, OneDrive (connector)	Redline Module 2.5 against the new nonclinical data and produce a change summary for the submission team.

Surface	Best for	Primary users	Where it runs	Example task
Claude Platform (API)	Embedding Claude into ELN, LIMS, CTMS, safety, or RWE systems	Scientific computing, R&D IT, clinical systems	Anthropic API, Amazon Bedrock, Google Vertex AI, Microsoft Foundry	Integrate Claude into our ELN to draft experiment summaries from structured results with source citations.
Claude Managed Agents	Running custom agents as hosted cloud services	Platform and scientific computing teams	Claude Platform (hosted by Anthropic)	Deploy our Literature Surveillance agent as a managed service with scoped database access and audit tracing.

[Chapter] Claude Science product overview

Claude Science is an AI daily driver that takes the toil out of science so researchers spend their time on the work only they can do. It's an agentic research app for the digital workflows core to scientific progress, from planning research and running analysis to explaining the results, all through natural language with Claude. It covers the six workflow archetypes that fill a biologist's day—literature review, experiment planning, data management, data analysis, figure generation, and scientific writing—and its starting thesis is to remove the toil between them: the database schemas, format conversions, and package interop that sit between a question and its answer, so the scientist keeps the design decisions and the interpretation.

The product is built first for computational and dry-lab biologists, and Anthropic's early-access program found that wet-lab scientists and PIs adopt it close behind—it extends their computational reach the same way Claude Code did for non-engineers. The starter catalog skews to biology, but the core extends to chemistry, neuroscience, and ML research.

How it works

Claude Science is a standalone application for macOS and Linux that runs a local daemon with its UI in the browser—the same model as a Jupyter notebook, but the agent is driving. The scientist installs it wherever the data lives: a laptop, a lab Linux box, an HPC login node, or a cloud VM. Data, compute environments, and agents stay on that machine, while the scientist connects from a laptop browser over an SSH tunnel when the daemon is running remotely. When a job needs bigger hardware, Claude Science dispatches it from the same session to the lab's own GPU box, SSH host, SLURM cluster (it auto-detects SLURM and writes the batch directives), or serverless GPU account.

Once installed, agents can be pointed at any local folder, including FASTQ files, AnnData objects, Seurat objects, and connect natively to S3, GCS, GitHub, and institutional literature access. Conda and pip environments are managed per specialist, and sessions, kernels, and artifacts persist across reboots.

Domain expertise, ready on day one

Claude Science ships with configurable capabilities for common scientific workflows, backed by optional connections to more than sixty scientific databases and roughly 150 curated skills. When a project spans domains, it plans and routes across them automatically. Public databases and other third-party resources are governed by their own licenses and use terms; organizations are responsible for ensuring their use — including commercial use — complies with those terms and their own entitlements. Information about network domains, connectors, and skills can be found in Settings.

Claude Sciences ships a library of pre-built Python skills that, if enabled, allow Claude to access certain major open scientific databases, so a scientist can ask a question in plain language and Claude writes and executes the query against the source of record, returning results with provenance.

Because these skills run code rather than retrieve documents, Claude can chain them: pull a gene's variants from one database, cross-reference drug interactions in another, and check expression across cell types in a third, inside a single analysis. Each skill is open source, so computational teams can inspect the query logic, pin versions, or extend a skill with their organization's own filters and output formats.

The table below groups them by the kind of question they answer.

Genes, variants, and annotation

Skill	Source	Use when
gene-database	NCBI Gene / Datasets	Looking up a gene by symbol or ID – RefSeqs, GO terms, genomic location, associated phenotypes, batch retrieval.
biothings-database	BioThings.io (MyGene, MyVariant, MyChem, MyDisease)	Resolving identifiers across databases or pulling aggregated annotation for a gene, variant, compound, or disease in one call.
ena-database	European Nucleotide Archive	Retrieving raw sequence data – FASTQ reads, assemblies, sample metadata – by accession.
harmonizome-database	Harmonizome	Asking what a gene is associated with across 170+ functional genomics resources at once.

Expression and single-cell

Skill	Source	Use when
cellxgene-census	CZ CELLxGENE Census	Filtering 125M+ cells by type, tissue, or disease and pulling expression matrices straight into scanpy or PyTorch.
immgen-database	ImmGen	Checking expression of a gene across mouse immune cell populations.

allen-brain-database	Allen Brain Atlas	Querying gene expression, connectivity, or spatial transcriptomics across mouse and human brain regions.
----------------------	-------------------	--

Oncology

Skill	Source	Use when
cosmic-database	COSMIC	Searching somatic mutations, the Cancer Gene Census, mutational signatures, or fusion events. Requires institutional authentication.
tcia-database	The Cancer Imaging Archive	Pulling DICOM imaging series from public cancer imaging collections for radiomics or model training.

Pharmacology and safety

Skill	Source	Use when
clinpgx-database	ClinPGx (PharmGKB)	Checking gene–drug interactions, CPIC dosing guidelines, or allele function for a pharmacogenomic question.
fda-database	openFDA	Searching adverse event reports, recalls, drug labels, device 510(k)/PMA records, or UNII identifiers.

Metabolomics

Skill	Source	Use when
hmdb-database	Human Metabolome Database	Looking up any of 220K+ human metabolites – properties, biomarker associations, NMR/MS spectra, pathway membership.
metabolomics-workbench-database	Metabolomics Workbench	Searching 4,200+ public metabolomics studies, RefMet nomenclature, or running m/z lookups.

Neuroscience

Skill	Source	Use when
neuromorpho-database	NeuroMorpho.Org	Retrieving neuron morphology reconstructions and morphometrics by species, brain region, or cell type.
openneuro-database	OpenNeuro	Finding and pulling BIDS-formatted MRI, fMRI, EEG, MEG, or PET datasets.

Multi-database toolkits

For exploratory work that crosses several sources, Claude can also load broader Python toolkits as skills:

Skill	Coverage	Use when
-------	----------	----------

bioservices	40+ services — UniProt, ChEMBL, PubChem, Reactome, QuickGO, KEGG, Ensembl, BioMart, and more	Running pathway analysis or ID mapping across many providers through one unified interface.
gget	20+ databases — Ensembl, UniProt, NCBI, ARCHS4, Enrichr, OpenTargets, PDB, AlphaFold, BLAST	Quick one-liner lookups: gene info, enrichment, reference genome download, a fast BLAST.
biopython (Bio.Entrez)	All NCBI Entrez databases — PubMed, Gene, Nucleotide, Protein, SRA, Taxonomy, Assembly, BioProject	Anything that lives behind NCBI E-utilities and isn't covered by a more specific skill above.
hmmer	Pfam, UniProtKB, Reference Proteomes	Profile and sequence searches — phmmer, hmmscan, hmmsearch, jackhmmer — for domain identification and remote homology.
foldseek	AlphaFold DB, PDB100, ESMAtlas, CATH	Searching by 3D structure rather than sequence to find structural homologs the sequence would miss.

The catalog is opinionated but open: a lab can save any working pipeline as a reusable skill, build a custom specialist for its own methods from inside any session, or wrap an internal API once so every future session inherits it.

Analysis that holds up under review

Five design choices make Claude Science suitable for scientific work.

Persistent kernels

Agents load a dataset into a persistent Python or R kernel once; from then on they're exploring it, not re-loading it. Variables, dataframes, and loaded models stay in memory across the whole analysis. Agents see their own plots—every figure generated is fed back into the agent's context, so it runs QC on its own output, spots the outlier cluster in its own UMAP, and filters it before moving on.

Full provenance on every artifact

Figures, tables, reports, and notebooks are first-class objects, not files to dig up afterward. Each ships with four layers of history: a human-readable description of what was done, the exact reproducible code that produced it, the conversation itself and the reasoning that led there, and a snapshot of every package and version used.

A background reviewer

A separate reviewer agent reads each session's transcript while the primary agent works and flags any claim it cannot trace to evidence. Findings surface inline at the suspect sentence and the agent fixes them before finishing. The reviewer runs every session by default and can be triggered manually at any point.

Plans before actions, permissions you can see

Claude Science drafts each task as a step-by-step plan and waits for approval before executing; the plan stays visible as a checklist and can be edited or scoped down while work runs. Whenever an agent needs to reach a new website, open a folder, or run code, the same approval card appears—allow it once, for this project, or always—and every decision can be reviewed or revoked from a single permissions screen. An OS-level sandbox with deny-by-default network egress sits underneath, and a human-approval broker gates code execution, network grants, host file access, deletions, MCP tools, remote compute, and skill persistence.

Built in biosecurity safeguards

Biology is a dual-use domain, and Claude Science ships with safeguards specific to it: biosecurity rules are written unconditionally into every agent's system prompt, a per-turn bio trajectory classifier runs in the binary and cannot be disabled, and the product completed external red-teaming and Anthropic Safeguards review before public release.

These sit on top of the runtime sandbox and approval broker as an additional layer for organizations whose risk and compliance teams will ask exactly this question.

Working with results

Scientists iterate on outputs in plain language. Click directly on a figure to drop an annotation, or just ask: drop the gridlines, make the axis log scale, switch to a colorblind-safe palette, or run it again without the outliers. The agent reads the exact code that produced the artifact and edits surgically rather than regenerating something that looks roughly right. Every version is saved with its provenance intact. From any point a scientist can fork the session to explore two approaches in parallel without losing the original thread.

Built-in scientific renderers display protein structures, DNA sequences, multiple sequence alignments, genome tracks, and small molecules natively, alongside spreadsheets, interactive HTML, PDFs, notebooks, and live LaTeX and Markdown for manuscript editing.

MCP connectors and extensions

Claude Science is MCP-native: any MCP server—data warehouses, ELNs, ticketing, and internal services—connects, and tools the rest of the organization already uses work here, too.

Scientific databases:



Basecamp Research



Phylo



protocols.io



ToolUniverse



Open Targets



Wiley



Clarivate



Owkin

Lab and partner tools:



Benchling



Modal



SandboxAQ



Revvity Signals



Boltz



Inductive Bio



Latch Bio



Synthesize Bio



Helix



10x Genomics



Medidata



BioRender

When to use a skill vs. a connector

Use a connector when the answer lives in the organization's own systems, such as an ELN, a CTMS, or a regulatory document repository, and entitlements matter. Use a scientific data skill when the answer lives in the public record and the value is in querying it precisely, reproducibly, and in combination with other sources. Most real questions use both: Claude pulls the internal context through a connector, grounds it against public reference data through a skill, and returns an analysis the scientist can verify line by line.

[Chapter] Building an AI-first life sciences organization: real-world examples

Claude is changing how life sciences teams work, from early discovery through regulatory submission and into medical affairs. The organizations adopting it across functions are reporting meaningful shifts in where scientist and writer hours go, and in how quickly results and documents move through review.

[Novo Nordisk](#) built NovoScribe, a generative AI platform powered by Claude that automates the creation of clinical study reports, device verification protocols, and patient materials. Clinical documentation that previously took more than ten weeks now reaches a reviewable first draft in roughly ten minutes.

"Claude has helped us cut writing times on CSRs by 90% so we can get documentation directly into human hands for review and approval." — Waheed Jowiya, Digitalization Strategy Director, Novo Nordisk

[The Garvan Institute of Medical Research](#) has adopted Claude across research, operations, and administration, with more than twenty software engineers and data scientists now using agentic development tools for drug discovery research, rare disease diagnosis through multi-agent systems that interpret genetic variants, and data analysis across genomics projects.

"Claude Code has completely transformed the way that I work as a scientist and as a leader." — Daniel MacArthur, Professor, Garvan Institute

[Biomni](#) uses Claude as the agent architecture behind its biomedical research platform, connecting to more than 150 tools and databases to automate literature review, bioinformatics analysis, and experimental protocol design. In benchmark testing, the

platform completed a bioinformatics analysis in 35 minutes that took human experts three weeks — roughly 800 times faster.

"Claude demonstrated the best performance across our benchmarks, particularly in scientific and biological knowledge, coding ability, and agentic workflows." — [TK — confirm full attribution for Huang, Biomni, with PMM]

[Sanofi](#) has deployed Claude enterprise-wide inside its internal Concierge application, where it is now used by the majority of Sanofi employees daily across the value chain.

"Claude, paired with internal knowledge libraries, has become integral to Sanofi's AI-driven transformation and is used by the majority of Sanofians on a daily basis within our internal Concierge app. We are seeing significant efficiency gains across the value-chain by optimizing our processes, while our enterprise-wide deployment has enhanced how our teams work. This collaboration with Anthropic augments human expertise to deliver life-changing medicines faster and more effectively to patients worldwide." — Emmanuel Frenehard, Chief Digital Officer, Sanofi

[Chapter] Claude Science adoption roadmap

Most successful Claude Science rollouts will follow this sequence: the team lays the foundation, runs a pilot inside one or two computational groups, and scales out across R&D—with the document-facing Claude surfaces (Claude Cowork, Microsoft 365) coming online alongside as adjacent functions pull for them.

Phase 1: Lay the foundation

Claude Science runs locally, so the foundation phase is about getting it next to the right data and compute rather than standing up cloud tenancy. For most research organizations this means deciding where the daemon will be hosted and confirming that scientists can reach that host from their browser. Research IT will want to review the OS-level sandbox, the deny-by-default network egress allowlist, and the human-approval broker that gates code execution, file access, and remote compute.

Account setup is the same SSO and SCIM work as any other Claude surface, with one extra step: an admin must enable Claude Science before any user in the org can download or sign in. On a Team plan, the admin turns it on under *admin settings* → *capabilities*. On an Enterprise plan, the admin creates or edits a role that includes the Claude Science permission, assigns that role to the pilot group, and then enables the capability—so access is scoped to the pilot from day one rather than org-wide. Claude Science is available in beta on first-party Claude paid plans and Claude for Enterprise, including enterprise customers who procure through AWS Marketplace.

Scope governance in parallel. For groups working with NIH-controlled data, patient-level data, or sponsor IP that cannot leave the lab's machines, the local-daemon model is the design that makes Claude Science deployable where a SaaS product would not be—but quality, IT security, and data privacy should review the install footprint, the network allowlist, and the compute-dispatch targets before the first scientist points it at a controlled-data folder.

Pick pilot teams with a motivated lead who is already pushing on AI and whose work is analysis-heavy and standard-shape. Computational biology and bioinformatics groups are the natural starting point; the early-access pattern at Anthropic is that wet-lab scientists and PIs follow quickly once they see a colleague run an analysis they could not have run themselves. Pick work where the value is obvious in the first session: a single-cell dataset that would otherwise take three weeks, a CRISPR screen analysis, a literature synthesis ahead of a program review.

[Callout] **Pro-tip:** the first session matters. A scientist who opens Claude Science, points it

at a folder of FASTQ files, approves the plan, and gets a clustered UMAP with the code and environment captured underneath will come back. A scientist who opens it without data in reach will close it. Make sure the install lands next to real data.

Phase 2: Pilot

At this stage, champions are running real analyses on real lab data and measuring against criteria defined upfront. Cycle time is the most common metric, in other words, how long the pilot analysis took before Claude Science and how long it takes after, on the same dataset class. The second is how often a scientist or PI trusts the result without re-running it by hand. The third, specific to this surface, is reproducibility: take an artifact produced in week one, hand its provenance bundle to a different scientist in week four, and confirm they can re-run it cold.

A strong signal that a pilot is working is when champions start saving their own skills. A bioinformatician takes the lab's internal normalization pipeline and wraps it as a skill so every future session inherits it, while a group lead wraps the lab's LIMS API. Those skills become the lab's catalog, and they can be shared across the organization.

Adjacent surfaces typically come online during this phase. Medical writing and regulatory groups watching the pilot ask for Claude Cowork and the Microsoft 365 add-ins for their document work; scientific computing asks for Claude Code for the production pipelines that sit downstream of the analysis. Run those as parallel tracks rather than gating them on the Claude Science pilot.

[Call-out] **Pro-tip:** schedule weekly check-ins with pilot teams. Edge cases surface fast — a database schema the connector does not handle, a cluster scheduler quirk, a renderer that does not cover a niche file format — and the catalog is designed to be extended in response.

Phase 3: Scale

At this stage, skills and specialists that worked during the pilot are rolled out to additional groups, and Research IT moves from per-lab installs to a managed deployment pattern: a standard daemon host per group, a vetted network allowlist, a curated skill catalog seeded from the pilot, and a defined set of compute-dispatch targets. Skills compound across groups because so much computational biology shares structure: a single-cell skill built in oncology is most of the way to one built in immunology. New hires start on day one with the lab's encoded pipelines rather than building them from scratch for their own workflows.

Governance should be settled before scale, not after. Decide who owns each skill in the catalog, how it is QC'd before it is shared beyond its author's group, how provenance

bundles are retained for analyses that feed regulatory or publication outputs, and how the network allowlist is reviewed when a group requests a new external database.

[Call-out] **Pro-tip:** for analyses that will feed a regulatory submission or a publication, treat the four-layer provenance bundle as a controlled record. The description, code, conversation, and environment snapshot together are what a reviewer, an auditor, or a journal will want to see — agree on where those bundles are stored and for how long.

Phase	Actions	What you'll see
Foundation	IT and data-governance review of local install, sandbox, and network allowlist. Decide daemon host pattern. Identify 2–3 champion groups in computational biology or bioinformatics. Confirm SSO/SCIM and Team/Enterprise plan.	Champions reporting back use cases. First "this would have taken me three weeks" moments.
Pilot	Champions run real analyses on real lab data. Weekly check-ins. Measure cycle time, keep-rate, and cold-reproduce rate. Stand up Cowork and M365 for adjacent document functions in parallel.	Measurable time savings. Champions saving custom skills and specialists. Wet-lab scientists and PIs joining behind the computational leads.
Scale	Managed daemon host pattern. Curated org skill catalog. Vetted network allowlist and compute-dispatch targets. Agreed provenance-retention policy for regulated and publication-bound analyses. Onboard the next wave of groups.	Skills shared across therapeutic areas. New hires ramping on encoded pipelines. Declining "can someone help me run this" requests to the bioinformatics core.

[Chapter] Function and workflow use cases

Claude Science was built to handle a scientist's most analysis-heavy work and repetitive workstreams. Claude Science handles the pipeline assembly, the database wrangling, and the figure iteration so your time goes to the design decisions and results interpretation. The sections below show what that looks like across the research lifecycle, with the adjacent Claude surfaces noted where the work crosses into document and submission territory.

Discover and plan

Early-stage work is reading, reconciling, and designing—across literature, public databases, and the lab's prior results. Claude Science compresses that cycle with:

- Literature review and synthesis across primary sources, with citation verification by the background reviewer
- Target and indication assessment from PubMed, ChEMBL, ClinicalTrials.gov, Open Targets, and internal program files
- Experiment and protocol design, construct design, and assembly strategy for ordering
- Competitive trial scans cross-referenced with preclinical model availability
- Hypothesis generation against pathway and ontology references

Analyze

Analysis is where most of the toil lives. Claude Science runs end-to-end pipelines with kernels that persist, agents that QC their own outputs, and branching for alternative approaches:

- Single-cell RNA-seq clustering, marker identification, and treatment-response analysis
- Genomics from FASTQ through alignment, QC, and variant calling
- CRISPR screen design and analysis
- Proteomics quantitation and interpretation, protein structure prediction, and homolog search
- Cheminformatics: ADMET prediction, similarity search, structural alert flagging
- Phylogenetic and evolutionary analysis
- Parameter sweeps and ML modeling dispatched to SLURM or remote GPU

Polish and publish

Results become defensible outputs: figures with provenance, methods sections that match the code, and manuscripts edited in place. Claude Science covers:

- Surgical figure iteration in plain language with every version saved
- Methods-section drafting from the artifact's provenance bundle
- Manuscript drafting and editing with live LaTeX and Markdown rendering
- Progress reports and PI briefings synthesized from session history and lab notes

Adjacent document and regulatory work (*Cowork*, *Microsoft 365*)

Once results leave the lab, the surrounding Claude surfaces pick up the document load with qualified human review before anything enters a regulated record:

- Protocol and synopsis drafting against the sponsor shell, CSR section drafting from TLFs
- CTD module authoring and gap analysis against current guidance
- SAE narrative drafting and QC, medical information response drafting
- SOP drafting, deviation write-ups, and inspection-readiness summaries

[Chapter] Frequently asked questions for CIOs and IT leaders

Note: The questions below cover Claude Science primarily, with Claude.ai and Claude Cowork addressed where the answer differs. Custom applications built on the Claude Platform have additional configuration options handled directly with the account team.

Where does Claude Science run, and where does our data go?

Claude Science is a local application for macOS and Linux. It runs a daemon on the machine where you install it — a scientist's laptop, a lab Linux box, an HPC login node, or a cloud VM in your tenancy — with its UI in the browser. Files remain on the host and are read in place; content the agent reads as part of an analysis is sent to Anthropic's API as context, subject to your plan's data-use and retention commitments. An OS-level sandbox with deny-by-default network egress controls all other traffic leaving the host. Claude Science is not currently available through Amazon Bedrock, Google Vertex AI, or Microsoft Foundry.

What does Claude Science install on user endpoints?

A signed binary that runs as a user-space daemon plus a browser-based UI. No kernel-level components. On Linux it can run headless and be reached over an SSH tunnel from the scientist's laptop. Windows is not currently supported.

How is code execution and file access controlled?

A human-approval broker gates thirteen action kinds, including code execution, network grants, host file access, deletions, MCP tool calls, remote compute dispatch, and skill persistence. Every request surfaces as a single approval card—allow once, for this project, or always—and every decision can be reviewed or revoked from one permission screen. An OS-level sandbox with allowlist-proxy network egress, SSRF and DNS-rebind defenses, and seccomp hardening sits underneath. Plan mode is on by default, so the agent drafts a step-by-step plan and waits for approval before executing.

How does Claude Science reach our HPC cluster or GPU hosts?

From inside a session the agent can dispatch jobs to an SSH host, a SLURM cluster (batch directives are written automatically), or a serverless GPU account the user supplies. Dispatch targets are gated by the same approval broker, so Research IT can restrict which hosts a group's install is permitted to reach.

Can Claude Science be used with NIH controlled-access or patient-level data?

Claude Science installs next to the data so files do not have to leave the lab's machines, which is the design that makes it deployable in environments where uploading is not an option. That said, NIH controlled-access datasets (for example, dbGaP) typically require the

analysis environment itself to meet NIST SP 800-171 controls, and Claude Science has not yet been assessed against that standard—formal NIH controlled-access data compliance is on the roadmap. Organizations working with controlled-access or patient-level data should pair the local-daemon model with their own data-governance review of the network allowlist and dispatch targets, and with internal policy on which datasets may be analyzed on which hosts. Claude Science is not a validated system; analyses that feed a regulated record go through the organization's existing qualified review.

Is Claude Science HIPAA-ready?

Not at launch. HIPAA readiness is on the post-launch roadmap. Until then, Claude Science should not be used to process protected health information. Your Anthropic account team can share current timing.

How is Claude Science priced, and how heavy is usage?

Claude Science draws down from the usage limits of the user's existing Claude plan—Pro, Max, Team, or Enterprise—with no separate license and no free tier. It is token-intensive: scientists routinely run several long agentic analyses in parallel, and heavy users consume at a rate comparable to heavy Claude Code use, so organizations should plan usage limits accordingly and expect heavy individual users to need Max-tier limits. Users can monitor consumption from inside the application under Settings → Usage. A subsidized Team plan is available for academic and nonprofit research labs; contact your Anthropic account team or visit the life sciences solutions page for eligibility.

What biosecurity controls are in place?

Biosecurity rules ship unconditionally in every agent's system prompt, a per-turn bio trajectory classifier runs in the binary and cannot be disabled by the user or admin, and authentication is OAuth-only with no anonymous or API-key access against the product. Claude Science completed external red-teaming and Anthropic Safeguards review against CBRNE risk before public release. These bio-specific controls sit on top of the OS-level sandbox, deny-by-default network egress, and human-approval broker described above.

What retention controls are available?

Team and Enterprise plans support custom data retention. Claude Science is a stateful product—sessions, artifacts, and provenance bundles require storage to function — so Zero Data Retention does not apply. ZDR is available on the Claude Platform (API) and Claude Code for approved customers.

What identity and access controls do you support?

Team and Enterprise plans include SSO via SAML, SCIM for user provisioning, role-based access controls, and admin-managed plugin and skill marketplaces. Claude Science is governed from the same admin console: a Team admin enables it under capabilities settings, and an Enterprise admin scopes it to specific groups via a role with the Claude Science permission before enabling the capability.

How does Claude Science integrate with our ELN or internal services?

Claude Science is MCP-native. Any MCP server — Benchling, an internal LIMS, a data warehouse, a ticketing system — connects, and skills can wrap internal APIs so the lab's own tools become first-class. Connectors authenticate as the end user and respect entitlements at the project and folder level, so the agent only sees what the scientist already has access to.

Can Claude Science be used in GxP-regulated workflows?

Claude Science is not a validated system. Organizations typically deploy it in research, analysis, and draft-support roles, with a qualified scientist or reviewer approving every output before it enters a validated record, a regulatory submission, or a publication. The four-layer provenance on every artifact—description, code, conversation, and environment snapshot—gives the reviewer a complete record of how each result was produced. Pair the deployment with your own CSV/CSA assessment of the surrounding process; your Anthropic account team can share how other sponsors and research organizations have approached this.

Where are Claude.ai and Claude Cowork hosted?

Both are SaaS products hosted by Anthropic. Organizations that need workloads to run inside their own cloud perimeter typically build on the Claude Platform via Amazon Bedrock, Google Vertex AI, or Microsoft Foundry. Cowork is a signed desktop application for macOS and Windows that reads only the local folders the user explicitly grants access to.

Who are the subprocessors?

A current list of Anthropic subprocessors is published at trust.anthropic.com and updated as the list changes.

Learn more in our [getting started guide](#).

[Chapter] Resources

- [Claude for life sciences solutions page](#): how Claude helps research, clinical, regulatory, and medical affairs teams across the development lifecycle.
- [Life sciences tutorials on Claude.com](#): step-by-step guides covering literature review, single-cell analysis, protocol drafting, and connector setup.
- [Claude skills catalog](#): Anthropic's public catalog of community-contributed skills, useful for finding pre-built scientific skills before authoring your own.

- [**The Enterprise AI Guide for Life Sciences**](#): Anthropic's enterprise transformation guide for life sciences IT and digital leaders.