


# DNAexus vs Seven Bridges vs Seqera: 2026 Comparison

IntuitionLabs.ai · Adrien Laurent · 2026-09-19 · dnanexus vs seven bridges vs seqera · dnanexus · seven bridges · seqera · genomics workflow platforms · cloud bioinformatics · nextflow · cwl · enterprise genomics

---



## Executive Summary

**DNAexus**, **Seven Bridges**, and **Seqera Platform** can all run production omics pipelines, but they embody different operating models. DNAexus is the strongest fit when the desired center of gravity is a managed precision-health data environment that combines governed data, native applications, portable workflows, notebooks, and cohort analysis. Seven Bridges is the clearest fit for teams standardizing on the Common Workflow Language (CWL) while requiring cloud, local high-performance computing (HPC), or hybrid operation. Seqera is the most direct fit for organizations already committed to Nextflow and wanting the control plane separated from customer compute and data.

Workflow language should be the first shortlist filter. DNAexus documents compilation of **CWL 1.2** and WDL draft-2, 1.0, and 1.1, plus licensed Nextflow execution, but each route has platform-specific constraints ([documentation.dnanexus.com](https://documentation.dnanexus.com)). Seven Bridges supports CWL draft-2 through **CWL 1.2**, while its WDL wrapper

lacks memoization and its Nextflow path packages source for platform execution ([docs.sevenbridges.com](https://docs.sevenbridges.com) ([docs.sevenbridges.com](https://docs.sevenbridges.com))). Seqera provides version-controlled, access-controlled execution of Nextflow; its documented CWL and WDL route is conversion into Nextflow DSL2, not native execution ([seqera.io](https://seqera.io)).

Deployment and data gravity provide the second filter. DNAnexus runs managed compute across AWS, Azure, and Oracle Cloud Infrastructure, and fixes project files and jobs to the selected region ([documentation.dnanexus.com](https://documentation.dnanexus.com) ([documentation.dnanexus.com](https://documentation.dnanexus.com))). Seven Bridges documents virtual-private-cloud, local HPC, and hybrid choices, although buyers should confirm the exact contracted region and deployment package ([sevenbridges.com](https://sevenbridges.com)). Seqera Enterprise installs in customer infrastructure, with backends spanning the major clouds, Kubernetes, and common HPC schedulers.

No credible public evidence establishes a universal performance winner. A 2021 peer-reviewed comparison used a genomic variant-calling pipeline and a scalability framework, but it did not test these three current commercial configurations head to head ([nature.com](https://www.nature.com)). Public prices are also incomplete. Seqera alone publishes managed-compute units including **\$0.10 per CPU-hour** and **\$0.025 per GiB-hour** ([docs.seqera.io](https://docs.seqera.io) ([docs.seqera.io](https://docs.seqera.io))). The defensible decision is therefore a scored proof of concept using the same versioned pipeline, containers, inputs, regions, concurrency, and acceptance evidence on each platform.

## Introduction and Background

An enterprise genomics platform is not merely a workflow runner. It is an operating boundary around data movement, identity, compute procurement, software provenance, collaboration, and regulated evidence. A platform that looks convenient to a pipeline developer may be awkward for a data-governance lead if results must cross regions. Conversely, a highly governed data environment may be unnecessary overhead for a group whose main requirement is to dispatch existing Nextflow pipelines to several HPC and cloud backends.

This comparison treats the three products as different control-plane and data-plane designs. **DNAnexus** emphasizes a centralized platform for omics and phenotypic data management and workflow automation ([dnanexus.com](https://dnanexus.com)). **Seven Bridges** combines a cloud bioinformatics workspace with CWL-centered portability and programmatic integration. **Seqera Platform** operationalizes Nextflow across vendor-hosted or customer-operated infrastructure. These boundaries matter more than a raw feature count.

For readers searching for **genomics workflow management platforms**, this is both a **cloud bioinformatics platforms comparison** and an **enterprise genomics analysis platforms** decision framework. It also answers the practical **Nextflow vs CWL workflow platforms** question. A DNAnexus vs Seqera decision usually turns on integrated managed data versus a Nextflow control plane, while a Seven Bridges vs Seqera decision usually turns on CWL versus Nextflow as the authoritative workflow language.

The intended reader is selecting a production environment, not choosing a language in isolation. The analysis therefore asks six practical questions:

- **Workflow authority:** Which representation is canonical, native DNAnexus workflow, CWL, WDL, or Nextflow?
- **Compute ownership:** Does the vendor manage execution, or does the platform orchestrate customer-owned cloud and HPC?
- **Data authority:** Where do source data, work files, metadata, logs, and results physically remain?

- **Evidence authority:** Which records demonstrate configuration, execution, change history, and approval?
- **Integration authority:** Which system triggers work and which system remains the downstream system of record?
- **Commercial authority:** Which costs are platform fees, cloud consumption, data transfer, support, validation, and migration?

IntuitionLabs is an adjacent life-sciences implementation and data-engineering consultancy, not a genomics workflow-platform vendor. Its relevant perspective is architectural: the firm describes data pipelines, integration, warehousing, and business intelligence as a service area ([intuitionlabs.ai](https://intuitionlabs.ai)). Accordingly, it is not included as a product in the comparison tables.

## DNAexus

### Capabilities

DNAexus is best understood as a managed precision-health data and analysis environment. Its native workflows connect platform apps or applets into staged analyses. A documented architectural limit is that one native workflow cannot be used as a stage inside another ([documentation.dnanexus.com](https://documentation.dnanexus.com)). For portable languages, dxCompiler translates CWL 1.2 and multiple WDL versions, but imported task and workflow names must be globally unique across the import tree ([documentation.dnanexus.com](https://documentation.dnanexus.com)). Nextflow pipelines use a head job and process subjobs, and the feature requires a license ([documentation.dnanexus.com](https://documentation.dnanexus.com)).

The development surface is broad. The Apache-2.0 dx-toolkit includes the command-line interface (CLI), build and debugging tools, upload and download agents, and language bindings; official documentation specifically says it provides tools for building and debugging apps ([documentation.dnanexus.com](https://documentation.dnanexus.com)). Separately licensed JupyterLab environments keep notebooks inside projects and can be run noninteractively ([documentation.dnanexus.com](https://documentation.dnanexus.com)).

### Adoption and operating model

The operating model favors centralized projects and governed regional placement. Files and jobs stay in the project's chosen region, and platform cloning does not move projects across regions ([documentation.dnanexus.com](https://documentation.dnanexus.com)). That is useful where residency boundaries can be mapped to projects, but it makes cross-region migration a design problem rather than a transparent platform action.

DNAexus documents exchange with [electronic lab notebooks \(ELNs\)](#), [laboratory information management systems \(LIMS\)](#), and clinical data management systems ([dnanexus.com](https://dnanexus.com)). Enterprise identity services can be used for organization login, and job identity tokens can connect workloads to external services ([documentation.dnanexus.com](https://documentation.dnanexus.com)).

### Strengths and limitations

- **Strength, governed data workspace:** Omics, phenotype, notebooks, applications, and workflows share a managed project boundary.
- **Strength, multi-language ingestion:** Native workflows, CWL, WDL, and licensed Nextflow provide migration paths.

- **Strength, managed cloud choice:** The documented compute catalog includes AWS, Azure, and Oracle Cloud Infrastructure GPUs ([documentation.dnanexus.com](https://documentation.dnanexus.com)).
- **Limitation, translation semantics:** CWL and WDL support comes through dxCompiler, so teams must test unsupported language features rather than infer conformance from a version label.
- **Limitation, licensed modules:** Cohort Browser requires an Apollo license, while JupyterLab and Nextflow access have their own license requirements ([documentation.dnanexus.com](https://documentation.dnanexus.com)).
- **Limitation, public pricing:** Enterprise subscription, professional services, and optional modules require an order-specific commercial review. The agreement describes support as available 24 hours a day, every day, but that is not a price schedule ([dnanexus.com](https://dnanexus.com)).

## Seven Bridges

### Capabilities

Seven Bridges is the most CWL-centered contender. Its documentation lists sbg:draft-2, CWL 1.0, 1.1, and 1.2 support ([docs.sevenbridges.com](https://docs.sevenbridges.com)). This aligns with the CWL project's own description of CWL as portable across computing platforms ([commonwl.org](https://commonwl.org)). Portability is not equivalence, however. Seven Bridges states that CWL document preprocessing is unsupported, and its scheduler can apply platform-specific execution hints ([docs.sevenbridges.com](https://docs.sevenbridges.com)).

The product also accepts Nextflow source packaged as a ZIP file and WDL through sbpack\_wdl ([docs.sevenbridges.com](https://docs.sevenbridges.com) ([docs.sevenbridges.com](https://docs.sevenbridges.com))). The WDL route is materially different from native job decomposition: a WDL pipeline runs as one platform job and has no memoization ([docs.sevenbridges.com](https://docs.sevenbridges.com)).

### Adoption and operating model

Seven Bridges documents deployment in a virtual private cloud, local HPC environment, or hybrid combination ([sevenbridges.com](https://sevenbridges.com)). Current project-location documentation names AWS us-east-1 and Google Cloud us-west1, with Azure regions on request. The chosen location governs both uploaded files and execution ([docs.sevenbridges.com](https://docs.sevenbridges.com)). Because official pages differ in how they describe provider and region availability, procurement should make the required regions and deployment packaging explicit contract acceptance criteria.

The integration surface includes a REST application programming interface (API), SB CLI, SBFS, and Python, R, and Java clients ([docs.sevenbridges.com](https://docs.sevenbridges.com)). The vendor describes API integration with LIMS, sequencing machines, analysis packages, and diagnostic-reporting services ([sevenbridges.com](https://sevenbridges.com)). Data Studio provides JupyterLab or RStudio for Python, R, and Julia work ([docs.sevenbridges.com](https://docs.sevenbridges.com)).

### Strengths and limitations

- **Strength, CWL alignment:** The platform's native authoring and execution model maps directly to a published workflow standard.
- **Strength, integration clients:** Multiple language clients and file-system tooling complement the REST API.
- **Strength, public research ecosystem:** The National Cancer Institute describes its Seven Bridges Cancer Genomics Cloud as hosting more than **850** curated tools and workflows ([datacommons.cancer.gov](https://datacommons.cancer.gov)).

- **Limitation, uneven language semantics:** Nextflow and WDL are supported through packaging or wrappers, with different caching and job behavior from CWL.
- **Limitation, deployment specificity:** Public marketing states hybrid capability, but buyers still need technical architecture, responsibility, upgrade, and support boundaries for their proposed topology.
- **Limitation, nonpublic license rates:** Documentation says multiple license and tier choices are priced through Sales ([docs.sevenbridges.com](https://docs.sevenbridges.com)).

## Seqera

### Capabilities

Seqera Platform is a Nextflow control plane. Its pipeline model is explicitly version-controlled, access-controlled, and reproducible ([docs.seqera.io](https://docs.seqera.io)). That is the cleanest fit when Nextflow source and configuration are already authoritative. Nextflow itself describes its purpose as creating scalable, portable, reproducible workflows ([nextflow.io](https://nextflow.io)).

The portability mechanism spans executors rather than multiple workflow languages. Official training material names Slurm, LSF, SGE, and PBS among supported HPC schedulers ([training.nextflow.io](https://training.nextflow.io)). Seqera markets AI-assisted conversion of Bash, CWL, and WDL to DSL2, but conversion is not native execution and should be evaluated like a code migration ([seqera.io](https://seqera.io)).

Seqera's adjacent services extend the run environment. Data Explorer supports AWS S3, Azure Blob Storage, Google Cloud Storage, and S3-compatible endpoints ([docs.seqera.io](https://docs.seqera.io)). Studios provides containerized Jupyter, R, Visual Studio Code, and remote-desktop sessions ([docs.seqera.io](https://docs.seqera.io)). Wave can attach Software Package Data Exchange (SPDX) software bills of materials to builds ([docs.seqera.io](https://docs.seqera.io)).

### Adoption and operating model

Cloud is hosted by Seqera, while Enterprise is installed in customer infrastructure ([docs.seqera.io](https://docs.seqera.io)). The core data-location distinction is unusually explicit: Seqera says orchestration does not transfer customer data to the infrastructure where Platform runs ([docs.seqera.io](https://docs.seqera.io)). Enterprise can expose AWS, Azure, Google Cloud, Kubernetes, Slurm, LSF, Grid Engine, PBS Pro, Moab, and local backends, subject to configuration ([docs.seqera.io](https://docs.seqera.io)).

Automation uses an OpenAPI-described API, CLI, and `sekerakit`. Seqera cautions that it does not guarantee generated OpenAPI clients ([docs.seqera.io](https://docs.seqera.io)). Documentation gives pipeline execution as data arrives from a sequencer as an automation example ([docs.seqera.io](https://docs.seqera.io)).

### Strengths and limitations

- **Strength, Nextflow-native operations:** The control plane follows Nextflow's pipeline, executor, cache, and configuration concepts.
- **Strength, bring compute to data:** Customer cloud and HPC resources can remain the execution boundary while Platform supplies orchestration and observability.

- **Strength, explicit managed-compute units:** Published rates make one deployment path modelable before a sales call.
- **Limitation, language concentration:** CWL and WDL estates face conversion and regression testing rather than native parity.
- **Limitation, incomplete total-cost estimate:** Seqera says its run estimator excludes storage, networking, and the head job ([docs.seqera.io](https://docs.seqera.io)).

## Feature Comparison

Table 1 summarizes architectural fit, not a winner by checkbox count. A capability is useful only when its supported version, deployment boundary, and evidence meet the buyer's actual use case.

| Decision axis                   | DNAexus                                                                                                                                                                | Seven Bridges                                                                                                                          | Seqera Platform                                                                                                                                                               |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Primary operating model</b>  | Managed precision-health data and analysis environment, with projects as the central governance boundary ( <a href="https://dnanexus.com">dnanexus.com</a> ).          | Bioinformatics workspace spanning cloud, local HPC, and hybrid deployment ( <a href="https://sevenbridges.com">sevenbridges.com</a> ). | Nextflow control plane as vendor-hosted Cloud or customer-installed Enterprise.                                                                                               |
| <b>Best-aligned workflow</b>    | Native applet workflow, plus compiled CWL/WDL and licensed Nextflow.                                                                                                   | CWL, with documented Nextflow and WDL packaging paths.                                                                                 | Nextflow; CWL/WDL are conversion candidates, not documented native execution.                                                                                                 |
| <b>Language caveat</b>          | Imported names must be unique, and each portable route has unsupported features ( <a href="https://documentation.dnanexus.com">documentation.dnanexus.com</a> ).       | CWL preprocessing is unsupported; WDL lacks memoization ( <a href="https://docs.sevenbridges.com">docs.sevenbridges.com</a> ).         | Operational semantics are tied to Nextflow versions and executors. <b>Nextflow 26.04</b> was released April 29, 2026 ( <a href="https://docs.seqera.io">docs.seqera.io</a> ). |
| <b>Data location</b>            | Project region fixes storage and jobs; cross-region platform cloning is not supported ( <a href="https://documentation.dnanexus.com">documentation.dnanexus.com</a> ). | Project location governs uploaded data and analysis ( <a href="https://docs.sevenbridges.com">docs.sevenbridges.com</a> ).             | Data remains in customer infrastructure rather than moving to Platform hosting.                                                                                               |
| <b>Interactive analysis</b>     | Licensed JupyterLab and Spark JupyterLab ( <a href="https://documentation.dnanexus.com">documentation.dnanexus.com</a> ).                                              | Data Studio with JupyterLab or RStudio ( <a href="https://docs.sevenbridges.com">docs.sevenbridges.com</a> ).                          | Studios with Jupyter, R, VS Code, and Xpra.                                                                                                                                   |
| <b>Programmatic integration</b> | dx CLI, agents, API, and Python, Java, C++, and R bindings.                                                                                                            | REST API, CLI, SBFS, and Python, R, and Java clients.                                                                                  | OpenAPI API, CLI, and sekerakit, without guaranteed generated-client compatibility.                                                                                           |
| <b>Public commercial signal</b> | Quote and order form; optional modules require licensing.                                                                                                              | Multiple tiers quoted by Sales.                                                                                                        | Managed compute publishes resource-unit rates; broader plans require commercial confirmation.                                                                                 |

The matrix shows why “best genomics workflow platform” has no context-free answer. DNAexus reduces the number of separate systems when governed multimodal data is the center. Seven Bridges minimizes conceptual distance for a CWL estate. Seqera minimizes conceptual distance for a Nextflow estate and preserves customer infrastructure as the compute boundary.

Table 2 turns data gravity into an explicit deployment decision. Cloud charges vary by service and contract, so the table focuses on controllable placement rather than fabricated total prices.

| Situation                                            | Primary risk                                       | Architecture test                                                                                                                                                                                                                                                   | Decision implication                                                                                                                                                         |
|------------------------------------------------------|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Large source data already in one cloud region</b> | Repeated transfer, latency, and duplicated storage | Run the same pipeline with compute colocated and non-colocated; record bytes moved and all charge categories. AWS identifies compute, storage, and outbound transfer as fundamental cost drivers ( <a href="https://docs.aws.amazon.com">docs.aws.amazon.com</a> ). | Prefer a supported execution boundary in the data region. Seqera explicitly recommends same-region compute and data ( <a href="https://docs.seqera.io">docs.seqera.io</a> ). |
| <b>Multi-region regulated program</b>                | Inadvertent localization or result movement        | Trace source, work, logs, cache, result, backup, and support-access locations separately.                                                                                                                                                                           | DNAnexus project regions give a clear boundary, while Seven Bridges locations and Seqera customer storage require topology-specific confirmation.                            |
| <b>Azure storage and compute together</b>            | Misstating transfer assumptions                    | Verify that the selected services and paths qualify as intra-region traffic. Azure guidance says bandwidth within a single region is free ( <a href="https://learn.microsoft.com">learn.microsoft.com</a> ).                                                        | Same-region design can remove one bandwidth category, but not compute, storage, platform, or operations cost.                                                                |
| <b>Google Cloud ingestion</b>                        | Treating free ingress as free processing           | Separate inbound transfer from processing and later egress. Google Cloud states there is no charge for inbound data transfer ( <a href="https://cloud.google.com">cloud.google.com</a> ).                                                                           | Model every subsequent movement and colocate storage with accessing compute where practical ( <a href="https://docs.cloud.google.com">docs.cloud.google.com</a> ).           |
| <b>On-premises HPC is strategic</b>                  | Hidden gateway, cache, or operations burden        | Test scheduler integration, identity, containers, shared file systems, outbound connectivity, upgrades, and recovery.                                                                                                                                               | Seqera offers the broadest documented executor list; Seven Bridges states local HPC support; DNAnexus public material centers managed cloud execution.                       |

The practical lesson is to inventory every durable and transient copy. “Data stays in place” may refer to primary objects while metadata, logs, cached containers, work directories, or derived outputs follow different paths. The contract, architecture diagram, and observed network telemetry should agree.

## Performance and Benchmarks

There is no current, independent, apples-to-apples benchmark of DNAnexus, Seven Bridges, and Seqera running the same pipeline with identical infrastructure, versions, containers, cache state, and service configuration. Vendor throughput and savings claims therefore cannot identify a cross-vendor winner.

The available independent evidence is useful mainly for methodology. A peer-reviewed 2021 study compared workflow systems with both a variant-calling use case and a scalability framework ([nature.com](https://nature.com)). That study is not a purchasing benchmark for the current commercial products. Hardware, storage, executor, concurrency, caching, software version, and dataset composition determine whether the comparison transfers.

A fair proof of concept should collect at least these measurements:

- **Wall-clock duration:** From accepted trigger to verified result, including queue time and staging.
- **Compute consumption:** Requested and actual CPU, memory, accelerator, and runtime by task.

- **Data movement:** Bytes entering, leaving, and crossing regions or availability zones.
- **Cache behavior:** Cold run, warm repeat, partial invalidation, and failed-task resume.
- **Reliability:** Completed runs, retries, preemptions, and operator interventions.
- **Reproducibility:** Output hashes and metadata equivalence under a defined tolerance.
- **Operator effort:** Minutes to deploy, diagnose, approve, rerun, and export evidence.
- **Developer effort:** Changes needed to workflow source, configuration, containers, and tests.

Nextflow cache data is grouped under `.nextflow/cache` by session identifier, which makes resume-state handling a concrete test rather than a generic checkbox ([training.nextflow.io](https://training.nextflow.io)). CWL portability also needs conformance testing: the current authoritative standard is CWL 1.2.0, with 1.2.1 described as a corrective patch adding clarifications and conformance tests ([commonwl.org](https://commonwl.org) ([commonwl.org](https://commonwl.org))). For WDL, the current OpenWDL site identifies **WDL 1.3.0** as the latest stable specification, while each platform supports a particular subset through its execution engine ([openwdl.org](https://openwdl.org)).

## Data Analysis and Evidence

The measurable commercial evidence is asymmetric. DNAnexus and Seven Bridges do not publish complete enterprise price schedules in the reviewed sources. Seqera exposes managed-compute rates, but those rates are one component of a total deployment, not a comparable subscription price.

For Seqera Compute, one credit equals one US dollar in the documented pricing model. CPU is **0.1 credits, or \$0.10, per CPU-hour** ([docs.seqera.io](https://docs.seqera.io)); memory is **0.025 credits, or \$0.025, per GiB-hour** ([docs.seqera.io](https://docs.seqera.io)); storage is **0.025 credits, or \$0.025, per GB-month** ([docs.seqera.io](https://docs.seqera.io)). CPU and memory are billed from requested resources, not observed utilization ([docs.seqera.io](https://docs.seqera.io)). That makes right-sizing a measurable cost variable.

*Table 3* is a reader-supplied scorecard. Scores are observations from the buyer's own proof of concept, on a **0 to 5** scale. The weights are an illustrative starting point and must be disclosed and changed to match program risk.

| Criterion                      | Illustrative weight | Observed score, 0 to 5 | Required evidence                                             |
|--------------------------------|---------------------|------------------------|---------------------------------------------------------------|
| Workflow semantic fit          | 18                  | Buyer supplied         | Regression suite, unsupported-feature log, source changes     |
| Data locality and residency    | 15                  | Buyer supplied         | Architecture, object inventory, network telemetry             |
| Reproducibility and provenance | 12                  | Buyer supplied         | Hashes, manifests, container digest, workflow revision        |
| Identity and least privilege   | 10                  | Buyer supplied         | SSO, role matrix, service identity, revocation test           |
| Audit and validation evidence  | 12                  | Buyer supplied         | Audit export, change history, approvals, validation documents |
| Compute and scheduler fit      | 10                  | Buyer supplied         | Backend configuration, queue tests, failure recovery          |
| Integration effort             | 8                   | Buyer supplied         | Sequencer or LIMS trigger, API code, result handoff           |
| Performance and reliability    | 7                   | Buyer supplied         | Cold/warm timings, retries, interventions                     |

| Criterion                    | Illustrative weight | Observed score, 0 to 5 | Required evidence                                            |
|------------------------------|---------------------|------------------------|--------------------------------------------------------------|
| Three-year cost transparency | 5                   | Buyer supplied         | Quote, cloud bill, transfer, support, operations assumptions |
| Migration and exit           | 3                   | Buyer supplied         | Export test, portable assets, replacement runbook            |

The weighted score is **sum(weight × observed score) divided by sum(weights)**. Because these illustrative weights sum to **100**, the result remains on a 0 to 5 scale. This is not an industry benchmark. NIST says cloud procurement metrics should be representative, accurate, and reproducible ([nist.gov](https://nist.gov)). NIST also recognizes inspection, analysis, demonstration, and testing as general verification methods, all of which can be mapped to acceptance criteria ([csrc.nist.gov](https://csrc.nist.gov)).

For regulated use, product claims are inputs to validation, not substitutes for it. US electronic-record rules require closed systems to be validated for accuracy and reliability and require record changes not to obscure prior information ([ecfr.gov](https://ecfr.gov) ([ecfr.gov](https://ecfr.gov))). European Union Good Manufacturing Practice Annex 11 separately says the application should be validated and IT infrastructure qualified ([health.ec.europa.eu](https://health.ec.europa.eu)). Therefore, the validation package should include:

- **Intended use:** The exact research, clinical, or regulated process in scope.
- **Configuration baseline:** Platform edition, region, executor, workflow and container versions.
- **Supplier evidence:** Security reports, quality documents, release process, and support boundaries.
- **Requirements traceability:** Each user and system requirement linked to a test and result.
- **Audit evidence:** Who changed, approved, launched, accessed, and exported what, and when.
- **Retention:** Policy-aligned preservation of records, logs, manifests, and approvals.
- **Change control:** Impact assessment and regression testing for platform, workflow, and container updates.
- **Continuity:** Backup, recovery, degraded-mode, and vendor-exit tests.

The governing evidence is broader than a vendor's compliance label. FDA's Part 11 guidance says predicate-rule recordkeeping requirements remain enforceable ([fda.gov](https://fda.gov)). Its October 2024 clinical-investigations guidance recommends validating electronic systems before use ([fda.gov](https://fda.gov)) and keeps the regulated entity responsible for provider-hosted records ([fda.gov](https://fda.gov)). Annex 11 also calls supplier competence and reliability key selection factors ([health.ec.europa.eu](https://health.ec.europa.eu)).

Retention and tamper resistance should be tested directly. NIST describes hardware-enforced write-once media as an audit-protection enhancement ([nvlpubs.nist.gov](https://nvlpubs.nist.gov)), and NIST SP 800-171 requires retention aligned with the records-retention policy ([nvlpubs.nist.gov](https://nvlpubs.nist.gov)). HHS sets a six-year documentation period for the HIPAA Security Rule ([hhs.gov](https://hhs.gov)). Platform demonstrations should therefore cover export, immutable retention, search, review, and recovery, not only the presence of an audit screen.

Portability claims also need authoritative baselines. CWL calls version 1.2.0 its current authoritative standard ([commonwl.org](https://commonwl.org)) and describes the language as portable across platforms ([commonwl.org](https://commonwl.org)). OpenWDL identifies WDL 1.3.0 as the latest stable release ([openwdl.org](https://openwdl.org)) and ties real portability to the chosen execution engine ([openwdl.org](https://openwdl.org)). Nextflow's official training names Slurm, LSF, SGE, and PBS among its scheduler targets ([training.nextflow.io](https://training.nextflow.io)). These standards and runtime facts define the regression matrix for each platform adapter.

Cloud cost evidence should likewise remain separate from platform fees. AWS names compute, storage, and outbound transfer as basic cost drivers ([docs.aws.amazon.com](https://docs.aws.amazon.com)). Azure says same-region bandwidth is free ([learn.microsoft.com](https://learn.microsoft.com)), while Google Cloud lists inbound transfer at no charge ([cloud.google.com](https://cloud.google.com)). These statements explain what to measure, not which commercial platform is cheapest.

DNAexus documents a GxP license with a daily CSV audit trail ([documentation.dnanexus.com](https://documentation.dnanexus.com)). Seven Bridges says its data-access audit logs are retained for six years ([sevenbridges.com](https://sevenbridges.com)). Seqera exposes enhanced pre-change and post-change audit state with its GxP add-on ([docs.seqera.io](https://docs.seqera.io)). These are materially different evidence products and should be demonstrated with the buyer's own roles, retention policy, and export destination.

## Implications and Future Directions

The workflow-language decision has a longer half-life than the initial platform contract. CWL is explicitly designed for portability, but implementation profiles still differ. WDL can run wherever its selected engine is supported ([openwdl.org](https://openwdl.org)). Nextflow portability resides in its executor and configuration layers. Buyers should therefore define portability as a tested asset bundle, not as the presence of a language logo.

That bundle should contain:

- **Workflow source** pinned to a specification or engine version.
- **Container images** pinned by immutable digest, plus software bills of materials where available.
- **Configuration** separated by environment, with secrets excluded.
- **Test datasets** small enough for continuous regression and representative enough to expose edge cases.
- **Expected outputs** expressed as hashes, tolerances, schema checks, and biological quality controls.
- **Provenance manifest** covering source revision, parameters, images, reference data, executor, and region.
- **Exit runbook** showing how to export metadata and reproduce the workload outside the incumbent platform.

Container pinning is necessary but not sufficient. OpenWDL recommends specifying a task container to preserve portability ([openwdl.org](https://openwdl.org)). DNAexus can replace external registry references with platform-stored Docker tarballs to improve provenance and offline execution ([documentation.dnanexus.com](https://documentation.dnanexus.com)). Seqera Wave can attach an SPDX bill of materials to a build. The proof of concept should confirm whether those artifacts remain exportable and meaningful after platform exit.

Integration architecture deserves equal weight. A typical chain is sequencer or LIMS trigger, intake validation, workflow launch, compute execution, notebook review, result registration, and downstream analytics. The system of record at each handoff must be explicit. IntuitionLabs describes implementation and integration of complex enterprise systems as a service capability ([intuitionlabs.ai](https://intuitionlabs.ai)). In an adjacent-adviser role, that perspective supports requirements, integration design, and evidence planning, but it does not change which platform wins the buyer's measured scorecard.

# Frequently Asked Questions (FAQs)

## Which is the best genomics workflow platform?

There is no universal winner. DNAnexus fits a centralized, governed precision-health data environment; Seven Bridges fits a CWL-centered bioinformatics operating model; Seqera fits a Nextflow-first organization that wants orchestration across its own cloud, hybrid, or HPC resources. The best platform is the one that produces the highest disclosed proof-of-concept score under the buyer's own weights.

## How should a team decide between Nextflow and CWL platforms?

Choose based on the authoritative workflow estate and operating topology. CWL supplies a language-level portability standard, while Nextflow combines a dataflow language with a runtime and executor ecosystem. Test unsupported constructs, caching, retries, containers, file staging, and metadata export. Language conversion should be budgeted as migration work, not assumed to be lossless.

## Does DNAnexus support Nextflow, CWL, and WDL?

Yes, but through different mechanisms. DNAnexus documents licensed Nextflow pipeline execution and dxCompiler translation for CWL 1.2 and several WDL versions ([documentation.dnanexus.com](https://documentation.dnanexus.com)). Buyers should test the documented limitations against their own syntax and runtime behavior.

## Does Seven Bridges run Nextflow and WDL as well as CWL?

It supports all three, but not with identical semantics. CWL is the native center. Nextflow is added as a packaged code bundle, while WDL runs through a local-mode wrapper and lacks platform memoization ([docs.sevenbridges.com](https://docs.sevenbridges.com)).

## Does Seqera move genomic data into its own control plane?

Its Cloud documentation says customer data is not transferred to the infrastructure where Seqera Platform runs. Buyers should still map logs, metadata, work directories, caches, support access, backups, and derived outputs for the exact deployment.

## Can public list prices determine the cheapest option?

No. Seqera publishes some managed-compute unit rates, while full enterprise prices for all three require commercial confirmation. Cloud compute, storage, transfer, platform subscription, optional regulated-use modules, support, engineering, validation, and migration all belong in the cost model. A common pipeline and measured resource ledger provide a more defensible comparison than list-price fragments.

## Conclusion

**DNAexus vs Seven Bridges vs Seqera** is fundamentally a choice among operating models. DNAexus offers the most integrated managed data-and-analysis environment of the three. Seven Bridges offers the clearest CWL-centered environment with documented cloud, local HPC, and hybrid positioning. Seqera offers the clearest Nextflow-centered control plane and the strongest documented separation between platform hosting and customer data and compute.

The shortlist can be reduced quickly. A Nextflow-standardized organization with heterogeneous executors should begin with Seqera. A CWL-standardized organization that values visual bioinformatics tooling and packaged integration clients should begin with Seven Bridges. A program seeking a managed project boundary for multimodal data, cohorts, notebooks, applications, and several workflow ingestion routes should begin with DNAexus. “Begin with” is deliberate: none of these architectural matches replaces a proof of concept.

The final decision should use one representative, versioned pipeline and a disclosed scoring model. Keep infrastructure, region, container digests, input data, concurrency, cache state, and acceptance rules comparable. Measure runtime, consumption, movement, reliability, operator effort, reproducibility, evidence export, and migration effort. Obtain written confirmation of editions, supported versions, regions, support boundaries, regulated-use packages, and total commercial terms. This process converts a feature comparison into an auditable platform decision and prevents vendor claims from becoming unsupported cross-vendor benchmarks.

---

For informational purposes. Verify critical medical, legal, financial, regulatory, and technical information with qualified professionals and primary sources.