



Scaling Nextflow pipelines and AI agents on Verily Workbench



Built on  by **verily**



Scalable pipelines, scalable operations

To ensure researchers can easily incorporate advanced analytics tools into their work, Verily Health recently added Nextflow into Verily Workbench. By bringing this powerful workflow framework directly into Workbench's enterprise-grade Trusted Research Environment (TRE), researchers can now seamlessly scale both massive genomic pipeline and high-throughput AI agent evaluations with secure co-analysis and data governance.

Nextflow has become a go-to tool for scalable, reproducible scientific workflows due to its ability to handle complex dependencies and make pipelines portable across clouds and HPC systems. However, running Nextflow as a separate deployment can require significant operational overhead requiring teams to manage infrastructure and ensure proper governance and data access for users.


Without a centralized platform like Workbench, running Nextflow pipelines in production settings surfaces a consistent set of challenges:

- **Queue waits and resource contention.** Shared High-Performance Computing (HPC) clusters force researchers to wait hours or days for jobs to start. High-memory steps (256GB+) may not run at all when the cluster is under load.
- **Reproducibility gaps.** Tracking which pipeline version ran against which inputs, with which configuration, on which infrastructure requires disciplined manual bookkeeping or custom tooling.
- **Data access friction.** Connecting pipelines to governed or restricted datasets means navigating access management processes that sit outside the workflow tool.
- **Fragmented observability.** Monitoring a pipeline means checking cloud console logs, Nextflow's own reporting, and sometimes separate dashboards.
- **Accessibility barriers.** Running a pipeline end-to-end requires navigating CLI commands, configuration files, and infrastructure decisions that sit outside most researchers' core work.



Nextflow on Verily Workbench

Verily Workbench brings Nextflow into a unified research-focused environment, adding clean workflow support to your complete research stack. This integration targets the core pain points of workflow management by:



- **Consolidating the workflow lifecycle into a single UI:** Researchers can register pipelines, configure inputs, submit runs, and track progress in real time. All workflows, runs, inputs, outputs, and logs are accessible in the Workbench UI and via the CLI.
- **Simplifying pipeline registration and configuration:** Pipelines can be registered directly from cloud storage paths or GitHub repositories (both public and private). Researchers can easily specify Nextflow profiles and upload sample sheets for batch runs.
- **Ensuring reproducibility and auditability:** Every run record captures critical metadata, including the submitter, inputs, profiles applied, and outputs.
- **Providing scalable job monitoring:** A centralized dashboard provides real-time tracking of runs and tasks, making it easy to monitor status and access logs and output folders.

Note: Nextflow execution via Workbench is available now on Google Cloud Platform (GCP) and will be coming soon on Amazon Web Services (AWS).



Governed health data, not just fast compute

Several platforms today are equipped to run Nextflow leveraging cloud computing. What distinguishes Workbench is what happens before and after the pipeline runs, and the data governance that applies throughout.

Data Collections: governed inputs and outputs

Workbench pipelines don't just read from cloud buckets. They consume Data Collections, which are datasets packaged with access controls and governance policies enforced at the platform level. A Data Collection carrying an egress restriction, a region constraint, or an audit requirement doesn't lose those protections when a Nextflow pipeline reads from it. The policies are attached to the data, not to the environment, and they are enforced automatically.

This means a data steward can publish a dataset with confidence that every downstream use, whether a Nextflow pipeline, a Jupyter notebook, or a collaborator's workspace, inherits the governance controls they configured. No manual review of each pipeline. No trusting that the researcher is abiding by the data use agreement they read.

Co-analysis across governance boundaries

Workbench can combine multiple Data Collections with different governance policies in a single workspace. Their policies compose automatically, where the most restrictive constraint wins. This enables a class of analysis that walled-garden TREs and standalone workflow managers can't support: running a single pipeline against datasets from different institutions, each under its own data sharing agreement, with governance enforced by the platform rather than by contracts alone.

Cloud-native cost transparency

Workbench runs on the customer's own cloud project. When a Nextflow pipeline provisions a high-memory instance for a computationally intensive step, the researcher pays the cloud provider's price directly, with no platform markup on compute or storage.

The full research lifecycle

Running a Nextflow pipeline is rarely an isolated task. Data must be found and accessed beforehand, and results must be analyzed, shared, and governed afterward.

Workbench is designed around this reality. Once a pipeline completes, you can manage outputs directly and continue analysis in JupyterLab or VS Code environments within the same workspace. Results can be shared securely with collaborators by updating Data Collections or inviting them to the workspace.



Nextflow on Workbench in action

The following illustrates how Workbench simplifies launching **nf-core** and custom pipelines, and monitoring job executions.

Launching pipelines from Github or cloud storage

Workbench supports running Nextflow pipelines by referencing a **main.nf** file directly from a Git repository, including private repositories via Workbench's [Github integration](#), or from a cloud storage bucket. Researchers can specify profiles, reference sample sheets, and configure additional parameters, keeping the full run lifecycle organized in one workflow configuration flow

Running an nf-core pipeline is as straightforward as pasting the repository URL (e.g. <https://github.com/nf-core/rnaseq>), selecting a branch or tag, and pointing to the **main.nf** file.

Adding workflow

1 Provide source 2 Review details

Where's the workflow source located?

☐ Workspace bucket ☒ Git repository

Select workflow type

☐ WDL ☒ Nextflow

Select Repository Branch or Tag
Enter a Git repository URL. If the URL contains a specific branch or tag, it will be automatically detected and locked. Plain repository URLs allow manual selection.

Repository URL (Required)
`https://github.com/nf-core/rnaseq/tree/3.26.0`
Must be an HTTPS URL to a Git repository. Exclude the .git extension.

Select reference type

☐ Branch ☒ Tag

Select a tag
3.26.0
Auto-detected from URL

`https://github.com/nf-core/rnaseq/releases/tag/3.26.0`
107 workflows found

Looking for a specific file?
You can enter the full path to a folder or file in the search box below to quickly navigate to it.



Monitoring job execution

Once a workflow is running or complete, the Job Details page provides a real-time breakdown of task execution, input/output provenance, and direct access to engine logs.

← RNASEQ-TESTFULL

Delete

Overall status

JOB ID
5e02e7c0-ac1d-46b7-9a9d-622e75fcbb49

OVERALL JOB STATUS

Completed

Execution details

STARTED ON
May 21, 2026, 03:10 PM

COMPLETED ON
May 21, 2026, 08:49 PM

DURATION
05:39:02

Tasks (303)

Search...

Status

Name	Status	Created	Output	Logs
nf-76647e21-1779410649336	Completed	May 21, 2026, 08:44 PM	View on GCP	View
nf-2a9a6461-1779407111957	Completed	May 21, 2026, 07:45 PM	View on GCP	View
nf-7db61eef-1779406941967	Completed	May 21, 2026, 07:42 PM	View on GCP	View
nf-c19ab24e-1779406912508	Completed	May 21, 2026, 07:41 PM	View on GCP	--
nf-4b809725-1779406912347	Completed	May 21, 2026, 07:41 PM	View on GCP	--

Case study: Detecting antimicrobial resistance with Nextflow on Workbench

Antimicrobial resistance (AMR) is [one of the top ten global public health threats identified by the World Health Organization \(WHO\)](#). In partnership with [Verily Sightline](#), researchers at Bangor University [use wastewater based epidemiology](#) to deliver high-quality public health data, including antimicrobial resistant gene (resistome) surveillance. To process the massive metagenomic datasets this resistome surveillance requires, the team adapted [FloRes \(Flow Resistome\)](#), a state-of-the-art Nextflow pipeline developed by Bangor and [Compass Bioinformatics](#), to run on Verily Workbench.

Pipeline overview

Unlike targeted methods such as PCR, metagenomics can simultaneously profile thousands of resistance genes across complex microbial communities. The FloRes pipeline efficiently processes raw sequencing data into comprehensive resistome and microbiome profiles. By supporting both short- and long-read sequencing and offering species-level resolution, FloRes enables researchers to not only identify which resistance genes are present, but pinpoint the specific organisms carrying them, delivering streamlined, analysis-ready data.




Figure 1: FloRes pipeline flow for the standard_AMR_wKraken_and_Bracken mode with SNP validation.



From HPC to cloud

Bangor’s primary bottleneck was securing the computing capacity needed to run FloRes at scale. While the university has access to Supercomputing Wales (a shared High-Performance Computing (HPC) cluster), competing academic demands trigger processing delays of weeks or months, severely undermining the data's utility for timely decision-making. This infrastructure strain is compounded by rigid hardware limitations: specific pipeline steps demand up to 256GB of RAM. As dataset sizes more than doubled and long-read data was integrated, certain critical analyses became entirely unfeasible. This barrier is further amplified for external stakeholders who lack both the necessary computing infrastructure and the command-line expertise required to execute the pipeline manually.

Workbench eliminates this bottleneck entirely, allowing users to launch FloRes with zero queue wait time. Additional benefits include per-step compute provisioning, such as enabling high-memory machines for host removal and lightweight instances for normalization.

In addition to performance improvements, Workbench makes FloRes accessible to researchers beyond dedicated bioinformaticians. Users can launch the pipeline directly from the UI and configure pipeline parameters – such as custom reference databases for host removal, AMR gene mapping, and taxonomic identification – without modifying any pipeline code.

	AMR++ on HPC	FloRes on Workbench
Queue wait	Up to 48h	None
Compute scaling	Shared cluster resources	Auto-provisioned per step
Long-read support	No	Yes (minimap2)
Species-level classification	No	Yes (Bracken)
Execution interface	CLI / SLURM	Workbench UI

Figure 2: Capability comparison of AMR++ on shared HPC vs. FloRes on Verily Workbench.

Extensibility and future directions

FloRes on Workbench is directly applicable to any research group studying AMR in clinical, agricultural, or environmental settings. The integration pattern – adding a Workbench profile, containerizing dependencies, and parameterizing cloud storage paths – serves as a reusable template for migrating any Nextflow pipeline from institutional HPC.

Verily and Bangor are already looking ahead to other workflows, including [nf-core/mag](#) for metagenome assembly. Through this partnership, the team aims to democratize modern metagenomics by providing researchers and clinicians with point-and-click, cloud-based bioinformatics pipelines that unlock critical biological insights and eliminate local infrastructure and command-line complexity.

Evaluating AI agents with Nextflow

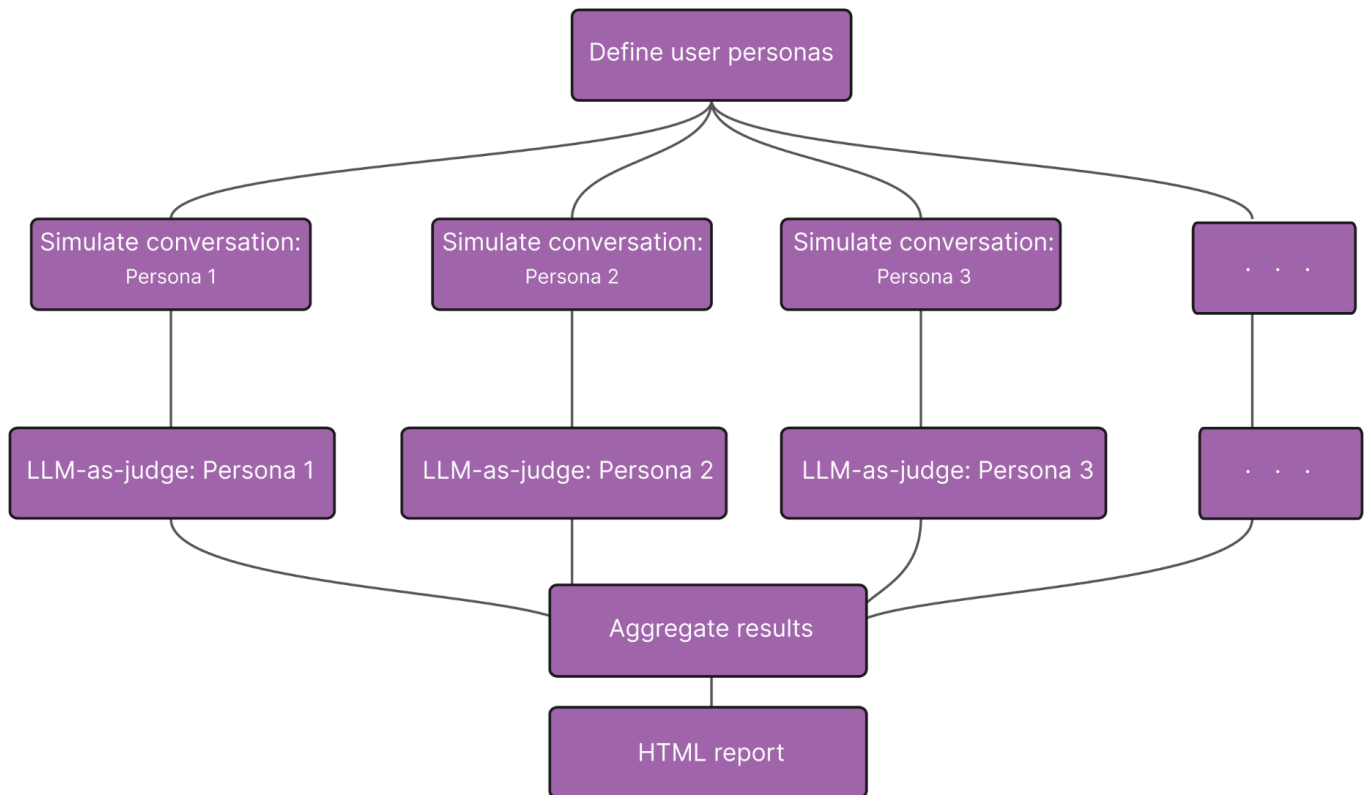
Nextflow's utility extends well beyond bioinformatics. To illustrate this, we walk through how an example use case could be run using Nextflow on Workbench. Specifically evaluating an AI agent's ability to coach a user who is interested in exercising.

Developing a robust AI agent requires evaluating its ability to respond properly in different user scenarios, handle edge cases, and handle off-topic or malicious queries. Testing all combinations of users, queries, and other variables can grow quickly, making parallelization a necessity in order to robustly evaluate agents. In our example, we run the following steps, where the number of simulations grows multiplicatively based on the number of user personalities and questions:

1. Create user personas with distinct personalities
2. Simulate a conversation between a user asking questions about exercising and their AI agent responding to those questions
3. Use an LLM-as-judge to assess how well the AI agent that is in-development helped each user
4. Summarize the results in an HTML report

The conversation simulation step relies on LangGraph, which we package into a Docker image. Each simulation runs independently, making the fan-out across personas a natural fit for parallel execution.





Scaling Workflows on Workbench

One of Nextflow's practical advantages is that the same pipeline runs locally and in the cloud without code changes. One can begin development on a single Workbench node or on a local machine to test the pipeline end-to-end immediately, ensuring fast and cost-effective iteration. Once validated, the same `main.nf` workflow can be scaled using Verily Workbench Workflows to execute tasks in parallel across multiple nodes, without modifying your pipeline code.

```
Nextflow 26.04.4 is available – Please consider updating your version to it

NEXTFLOW ~ version 25.10.5
Launching `main.nf` [prickly_lattes] DSL2 – revision: add5dd04cf

[4f/718369]create_patient_profiles_workflow:create_patient_profiles(2) [100%]
4 of 4, cached: 4✓
[cb/811de4]run_simulation_workflow:run_simulation(1) [100%]4 of 4, cached: 4✓
[ab/e2331a]run_simulation_workflow:llm_as_a_judge(1) [100%]4 of 4, cached: 4✓
[bd/818ea4]run_simulation_workflow:output_judgement_results(1) [100%]1 of 1, cached: 1✓
```



Results are available directly in the Workbench file explorer, where HTML outputs can be previewed without downloading. Below we preview the LLM-as-a-judge results from our pipeline:

nextflow

/simplified_barriers_eval_job_2026-06-18T043124437Z/09954da8-27a7-46fe-b2a1-f08b0b87dc6d/6f/0d186dfd16c1d3e572b5f5929924f2/judgement_results.html

HTML • 4.7 KB • Type text/html; charset=utf-8 • Last modified 7 minutes ago

- Cloud ID nextflow-outputs-vwb-dev-happy-apricot-410

[Copy preview link](#)

Summary:
 Total Records: 4
 Question Addressed (True): 2
 Question Not Addressed (False): 2
 Success Rate: 50.0%

Conversation	Question Addressed?	Reasoning
ai: Hi there, I am your AI Coach to help you exercise more. Please ask your question. human: Never lifted. What now? ai: Wonderful! Begin by learning basic movements with just your body weight or very light weights to master form safely. human: Exercises? ai: Start with squats, lunges, push-ups (on knees or against a wall), and planks to build foundational strength.	True	The AI successfully provided exercise suggestions as requested by the user.

Close

Close

Empowering science at scale

Nextflow is essential for modern, large-scale science, but managing infrastructure and waiting in long supercomputer queues often slows down discovery. Verily Workbench eliminates these barriers by bringing Nextflow into a secure, fully managed Trusted Research Environment.

As demonstrated by massive bioinformatics pipelines like FloRes and parallel AI simulations, Workbench automatically scales compute resources on the fly, entirely removing local hardware bottlenecks. By pairing this power with an intuitive user interface and governed datasets, it empowers everyone from domain scientists to senior bioinformaticians to easily launch and reproduce complex workflows.

Whether your team is tracking global public health threats or pioneering new AI tools, Verily Workbench stands ready to scale alongside you, transforming raw data into actionable insights.

Learn more at verily.com/workbench-request

