

Pegasus: Unshackling Scientific Workflows from Execution Environments

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: Pegasus

Cat. No.: B039198

[Get Quote](#)


A comparative analysis of **Pegasus** workflow portability and performance against leading alternatives for researchers and drug development professionals.

In the complex landscape of scientific research and drug development, the ability to execute complex computational workflows across diverse environments—from local clusters to high-performance computing (HPC) grids and the cloud—is paramount. The portability of these workflows is not merely a matter of convenience; it is a cornerstone of reproducible, scalable, and collaborative science. This guide provides an in-depth comparison of the **Pegasus** Workflow Management System with other popular alternatives, focusing on the critical aspect of workflow portability and performance, supported by experimental data.

At its core, **Pegasus** is engineered to decouple the logical description of a workflow from the physical resources where it will be executed.^{[1][2][3]} This is achieved by defining workflows in an abstract, resource-independent format, which **Pegasus** then maps to a concrete, executable plan tailored to the target environment.^[4] This "just-in-time" planning allows the same abstract workflow to be executed, without modification, on a researcher's laptop, a campus cluster, a national supercomputing facility, or a commercial cloud platform.^{[2][5]}

The Portability Paradigm: A Logical Overview

Pegasus's portability is rooted in its architecture, which separates the definition of what needs to be done from how and where it is done. The following diagram illustrates the logical relationship of a **Pegasus** workflow's journey from an abstract description to execution across varied environments.



[Click to download full resolution via product page](#)

Figure 1: Pegasus workflow portability across execution environments.

Performance Across Diverse Infrastructures: A Comparative Look

While portability is a key strength of **Pegasus**, performance is equally critical. A 2021 study in the journal *Bioinformatics* evaluated several workflow management systems, including **Pegasus-mpi-cluster**, Snakemake, and Nextflow, for a bioinformatics use case. The results, summarized in the table below, highlight the performance characteristics of each system in a local execution environment.

Metric	Pegasus-mpi-cluster	Snakemake	Nextflow
Execution Time (minutes)	4.0	4.5	4.0
CPU Consumption (%)	Lowest	Higher	Higher
Memory Footprint (MB)	Lowest	Higher	Higher
Containerization Support	Yes (Singularity, Docker)	Yes (Conda, Singularity, Docker)	Yes (Conda, Singularity, Docker)
CWL Support	No	Yes	Yes

Table 1: Performance Comparison of Workflow Management Systems (Local Execution) Data synthesized from "Evaluating Workflow Management Systems: A Bioinformatics Use Case".[\[5\]](#)

The study concluded that for their specific bioinformatics workflow, **Pegasus-mpi-cluster** demonstrated the best overall performance concerning the usage of computing resources.[\[5\]](#) It is important to note that performance can vary significantly based on the nature of the workflow (e.g., I/O-intensive vs. CPU-intensive), the scale of the data, and the configuration of the execution environment.

Another study, "On the Use of Cloud Computing for Scientific Workflows," explored the performance of **Pegasus** across different environments for an astronomy application. While not a direct comparison with other workflow managers, the study provided insights into the overhead associated with cloud execution. The experiments showed that while there is a performance overhead when moving from a local grid to a cloud environment, the flexibility and scalability of the cloud can offset this for many applications.

Experimental Protocols: A Glimpse into the Methodology

To ensure the objectivity of the presented data, it is crucial to understand the experimental setup. The following provides a summary of the methodologies employed in the cited studies.

Bioinformatics Workflow Comparison (Larsonneur et al., 2021)

- Workflow: A bioinformatics pipeline for biological knowledge discovery, involving multiple steps of data processing and analysis.
- Execution Environment: A single computation node (local execution).
- Metrics: The study proposed and measured ten metrics, including execution time, CPU consumption, and memory footprint.
- Workflow Managers: **Pegasus**-mpi-cluster, Snakemake, Nextflow, and others were evaluated.
- Data: The study utilized real-world biological data and metadata.

Astronomy Workflow on Different Infrastructures (Juve et al., 2009)

- Workflow: The Montage application, which creates image mosaics of the sky.
- Execution Environments:
 - Local machine
 - Local Grid cluster
 - Virtual cluster on a science cloud
 - Single virtual machine on a science cloud
- Metrics: The primary metric was the overall workflow execution time.
- Data: The workflows were run with varying sizes to test scalability. Input data was staged by **Pegasus** for the virtual environments.

The Power of Abstraction and Automation

Pegasus's approach to portability is not just about running the same script in different places. It involves a sophisticated set of features that automate many of the tedious and error-prone

tasks associated with distributed computing:

- **Data Management:** **Pegasus** automatically handles the staging of input data to the execution sites and the transfer of output data to desired locations.[3]
- **Provenance Tracking:** Detailed provenance information is captured for every job, including the software used, parameters, and data consumed and produced. This is crucial for reproducibility and debugging.[2]
- **Error Recovery:** **Pegasus** can automatically retry failed jobs and even entire workflows, enhancing the reliability of complex, long-running computations.[5]
- **Optimization:** The **Pegasus** planner can reorder, group, and prioritize tasks to improve overall workflow performance and efficiency.[3]

Conclusion: Choosing the Right Tool for the Job

The choice of a workflow management system is a critical decision for any research or development team. While alternatives like Snakemake and Nextflow offer strong features, particularly within the bioinformatics community, **Pegasus** distinguishes itself through its robust and proven approach to workflow portability across a wide range of execution environments.[5]

For research and drug development professionals who require the flexibility to seamlessly move their computational pipelines from local development to large-scale production on HPC and cloud resources, **Pegasus** provides a powerful and automated solution. Its emphasis on abstracting the workflow from the execution environment not only simplifies the user experience but also promotes the principles of reproducible and scalable science. By handling the complexities of data management, provenance tracking, and execution optimization, **Pegasus** allows scientists and researchers to focus on what they do best: pushing the boundaries of knowledge.

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. researchgate.net [researchgate.net]
- 2. GitHub - pegasus-isi/pegasus: Pegasus Workflow Management System - Automate, recover, and debug scientific computations. [github.com]
- 3. arokem.github.io [arokem.github.io]
- 4. researchgate.net [researchgate.net]
- 5. researchgate.net [researchgate.net]
- To cite this document: BenchChem. [Pegasus: Unshackling Scientific Workflows from Execution Environments]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b039198#pegasus-workflow-portability-across-different-execution-environments]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com