

# Implementing Robust Fault Tolerance in Pegasus Scientific Workflows

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

## Compound of Interest

Compound Name: Pegasus

Cat. No.: B039198

[Get Quote](#)

Application Notes and Protocols for Researchers, Scientists, and Drug Development Professionals

## Introduction

In complex, long-running scientific workflows, particularly in computationally intensive fields like drug development, the ability to handle failures gracefully is paramount. The **Pegasus** Workflow Management System provides a suite of robust fault tolerance mechanisms designed to ensure that workflows can recover from transient and permanent failures, thereby saving valuable time and computational resources. These mechanisms are critical for maintaining data integrity and ensuring the reproducibility of scientific results.

This document provides detailed application notes and protocols for implementing and evaluating key fault tolerance features within **Pegasus**. It is intended for researchers, scientists, and drug development professionals who utilize **Pegasus** to orchestrate their computational pipelines.

## Core Fault Tolerance Mechanisms in Pegasus

**Pegasus** employs a multi-layered approach to fault tolerance, addressing potential failures at different stages of workflow execution. The primary mechanisms include:

- Job Retries: Automatically re-submitting jobs that fail due to transient errors.[\[1\]](#)[\[2\]](#)[\[3\]](#)

- Rescue DAGs (Workflow Checkpointing): Saving the state of a workflow upon failure, allowing it to be resumed from the point of failure.[\[1\]](#)[\[2\]](#)[\[4\]](#)
- Data Integrity Checking: Verifying the integrity of data products throughout the workflow to prevent corruption.[\[5\]](#)[\[6\]](#)[\[7\]](#)
- Monitoring and Debugging: Tools to monitor workflow progress and diagnose failures.[\[8\]](#)[\[9\]](#)[\[10\]](#)

These features collectively contribute to the reliability and robustness of scientific workflows, minimizing manual intervention and maximizing computational throughput.

## Quantitative Data Summary

The following tables summarize quantitative data related to the performance of **Pegasus** fault tolerance mechanisms. This data is derived from simulation studies and real-world application benchmarks.

Table 1: Simulated Performance of Fault-Tolerant Clustering

This table presents simulation results for a Montage astronomy workflow, demonstrating the impact of different fault-tolerant clustering strategies on workflow makespan (total execution time) under a fixed task failure rate.

Clustering Strategy	Average Workflow Makespan (seconds)	Standard Deviation
Dynamic Clustering (DC)	2450	120
Selective Reclustering (SR)	2300	110
Dynamic Reclustering (DR)	2250	105

Data derived from a simulation-based evaluation of fault-tolerant clustering methods in **Pegasus**.[\[11\]](#)

Table 2: Overhead of Integrity Checking in Real-World Workflows

This table shows the computational overhead of enabling integrity checking in two different real-world scientific workflows.

Workflow	Number of Jobs	Total Wall Time (CPU hours)	Checksum Verification Time (CPU hours)	Overhead Percentage
OSG-KINC	50,606	61,800	42	0.068%
Dark Energy Survey	131	1.5	0.0009	0.062%

These results demonstrate that the overhead of ensuring data integrity is minimal in practice.[\[7\]](#)

## Experimental Protocols

The following protocols provide detailed, step-by-step methodologies for implementing and evaluating fault tolerance mechanisms in your **Pegasus** workflows.

### Protocol 1: Configuring and Evaluating Automatic Job Retries

Objective: To configure a **Pegasus** workflow to automatically retry failed jobs and to evaluate the effectiveness of this mechanism.

Materials:

- A **Pegasus** workflow definition (DAX file).
- Access to a compute cluster where **Pegasus** is installed.
- A script or method to induce transient failures in a workflow task.

Methodology:

- Define Job Retry Count:

- In your **Pegasus** properties file (e.g., **pegasus.properties**), specify the number of times a job should be retried upon failure. This is typically done by setting the `dagman.retry` property.
- This configuration instructs the underlying DAGMan engine to re-submit a failed job up to three times.[\[1\]](#)
- Induce a Transient Failure:
  - For testing purposes, introduce a transient error in one of your workflow's executable scripts. For example, have the script exit with a non-zero status code based on a random condition or an external trigger file.
- Plan and Execute the Workflow:
  - Plan the workflow using **pegasus-plan**.
  - Execute the workflow using **pegasus-run**.
- Monitor Workflow Execution:
  - Use **pegasus-status** to monitor the progress of the workflow. You will observe the failed job being re-submitted.
  - After the workflow completes, use **pegasus-analyzer** to inspect the workflow's execution logs. The analyzer will report the number of retries for the failed job.[\[8\]](#)[\[9\]](#)
- Analyze the Results:
  - Examine the output of **pegasus-analyzer** to confirm that the job was retried the configured number of times and eventually succeeded.
  - Use **pegasus-statistics** to gather detailed performance metrics, including the cumulative wall time of the workflow, which will reflect the time taken for the retries.[\[3\]](#)[\[12\]](#)

## Protocol 2: Utilizing Rescue DAGs for Workflow Recovery

Objective: To demonstrate the use of Rescue DAGs to recover a failed workflow from the point of failure.

Materials:

- A multi-stage **Pegasus** workflow definition (DAX file).
- A method to induce a non-recoverable failure in a mid-workflow task.

Methodology:

- Induce a Persistent Failure:
  - Modify an executable in your workflow to consistently fail (e.g., by exiting with a non-zero status code unconditionally).
- Execute the Workflow:
  - Plan and run the workflow as you normally would. The workflow will execute until it reaches the failing job and then halt.
- Identify the Failure:
  - Use **pegasus-status** and **pegasus-analyzer** to identify the failed job and the reason for failure.
- Generate and Submit the Rescue DAG:
  - Upon failure, **Pegasus** automatically generates a "rescue DAG" in the workflow's submit directory.<sup>[4]</sup> This DAG contains only the portions of the workflow that did not complete successfully.
  - Correct the issue that caused the failure (e.g., fix the failing script).
  - Re-submit the workflow using the same **pegasus-run** command. **Pegasus** will detect the rescue DAG and resume the execution from where it left off.
- Verify Recovery:

- Monitor the resumed workflow to ensure it completes successfully.
- Use **pegasus**-statistics to analyze the total workflow runtime, which will be the sum of the initial run and the resumed run. The cumulative workflow runtime reported by **pegasus**-statistics will include the time from both executions.[13]

## Protocol 3: Ensuring Data Integrity with Checksumming

Objective: To configure and verify the use of checksums to ensure the integrity of data products within a workflow.

Materials:

- A **Pegasus** workflow with input files.
- A replica catalog file.

Methodology:

- Enable Integrity Checking:
  - In your **pegasus.properties** file, enable integrity checking. The full setting enables checksumming for all data transfers.[14]
- Provide Input File Checksums (Optional but Recommended):
  - In your replica catalog, you can provide the checksums for your raw input files. **Pegasus** will use these to verify the integrity of the input data before a job starts.[5]
- Plan and Execute the Workflow:
  - Run **pegasus-plan** and **pegasus-run**. **Pegasus** will automatically generate and track checksums for all intermediate and output files.[6]
- Simulate Data Corruption:
  - To test the mechanism, manually corrupt an intermediate file in the workflow's execution directory while the workflow is running (this may require pausing the workflow or being

quick).

- Monitor and Analyze:
  - The job that uses the corrupted file as input will fail its integrity check.
  - **pegasus**-analyzer will report an integrity error for the failed job.
  - The workflow will attempt to retry the job, which may involve re-transferring the file.

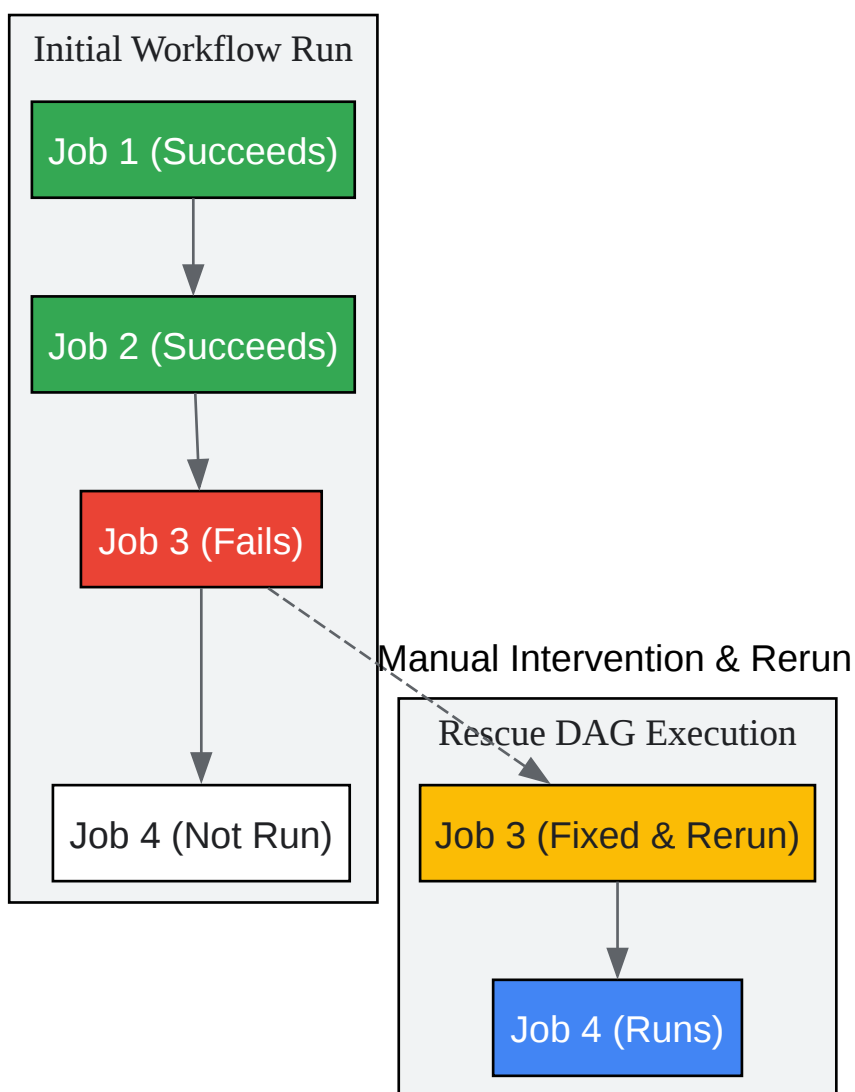
## Visualizing Fault Tolerance Workflows

The following diagrams, generated using the Graphviz DOT language, illustrate key fault tolerance concepts in **Pegasus**.




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

Caption: Automatic job retry mechanism in **Pegasus**.



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

Caption: Workflow recovery using a Rescue DAG.



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

Caption: Data integrity checking workflow in **Pegasus**.



## Conclusion

The fault tolerance capabilities of **Pegasus** are essential for the successful execution of complex scientific workflows. By implementing job retries, utilizing rescue DAGs, and ensuring data integrity, researchers can significantly improve the reliability and efficiency of their computational experiments. The protocols and information provided in this document serve as a guide for leveraging these powerful features to their full potential. For more advanced scenarios and detailed configuration options, users are encouraged to consult the official **Pegasus** documentation.

### Need Custom Synthesis?

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

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

## References

- 1. GitHub - pegasus-isi/pegasus: Pegasus Workflow Management System - Automate, recover, and debug scientific computations. [github.com]
- 2. Pegasus Workflows with Application Containers — CyVerse Container Camp: Container Technology for Scientific Research 0.1.0 documentation [cyverse-container-camp-workshop-2018.readthedocs-hosted.com]
- 3. research.cs.wisc.edu [research.cs.wisc.edu]
- 4. arokem.github.io [arokem.github.io]
- 5. pegasus.isi.edu [pegasus.isi.edu]
- 6. scitech.group [scitech.group]
- 7. agenda.hep.wisc.edu [agenda.hep.wisc.edu]
- 8. 9. Monitoring, Debugging and Statistics — Pegasus WMS 5.1.2-dev.0 documentation [pegasus.isi.edu]
- 9. GitHub - pegasus-isi/ACCESS-Pegasus-Examples: Pegasus Workflows examples including the Pegasus tutorial, to run on ACCESS resources. [github.com]
- 10. pegasus.isi.edu [pegasus.isi.edu]
- 11. pegasus.isi.edu [pegasus.isi.edu]

- 12. pegasus.isi.edu [pegasus.isi.edu]
- 13. Documentation – Pegasus WMS [pegasus.isi.edu]
- 14. 5. Data Management — Pegasus WMS 5.1.2-dev.0 documentation [pegasus.isi.edu]
- To cite this document: BenchChem. [Implementing Robust Fault Tolerance in Pegasus Scientific Workflows]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b039198#implementing-fault-tolerance-in-pegasus-workflows]

---

### 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](mailto:info@benchchem.com)