

Computational Systems & Infrastructure

REVIEW | COMMISSIONED SYNTHESIS

Resilience Patterns for Data-Intensive Scientific Workflows

John Thompson - Corresponding author

Abstract

This review examines resilience patterns for data-intensive scientific workflows. The organizing question is how workflow engines should preserve scientific meaning through retries, partial failure, and infrastructure change. Ten related scholarly sources are synthesized through a decision-centered framework spanning problem definition, mechanism, measurement, evaluation, implementation, and governance. The review does not invent experiments, pooled estimates, or unreported quantitative results. It instead evaluates the strength and transferability of the available evidence, with particular attention to declaring a workflow successful because every process eventually exited. The resulting framework links technical or empirical performance to explicit use conditions and identifies tests that should precede wider adoption in long-running computational experiments on shared infrastructure.

Keywords: scientific workflows, resilience, fault tolerance, provenance, reproducibility

Synthesis lens	Principal question
Mechanism	the dependencies among task graphs, data products, checkpoints, schedulers, provenance, and storage
Evaluation	recovery point, recovery time, silent corruption, nondeterminism, and scientific equivalence after restart
Primary risk	declaring a workflow successful because every process eventually exited

1. Introduction

Research on resilience patterns for data-intensive scientific workflows sits at the boundary between a promising component and a consequential decision. The core question is how workflow engines should preserve scientific meaning through retries, partial failure, and infrastructure change. Answering it requires more than assembling favorable results. It requires a chain of evidence that makes the problem boundary, mechanism, measurement, comparator, uncertainty, and accountable decision visible to the reader.

This article presents a structured narrative review of ten related publications. Sources were selected for topical relevance, traceable publication metadata, and complementary coverage of technical, empirical, and governance questions. No meta-analytic pooling is attempted because the included studies differ in designs, populations, system boundaries, and outcomes. The purpose is decision-oriented synthesis: to identify what each source can support, where transfer remains uncertain, and which evidence would justify the next stage of use.

The central risk is declaring a workflow successful because every process eventually exited. A top-line metric or broad policy label can appear decisive while concealing the conditions that produced it. Accordingly, the synthesis separates observation from interpretation and implementation from validation. This separation is especially important when the same system creates benefits for one user or institution while transferring cost, risk, or administrative burden to another.

2. Evidence Base and Problem Boundary

A defensible review begins by defining the unit of analysis. For the present topic, the relevant boundary includes inputs, institutional or technical context, the mechanism producing an output, the person or system acting on that output, and the downstream outcome. Evidence falls outside the boundary when it measures only an intermediate proxy yet is used to claim an end-to-end benefit.

Guedes et al. (2019) addresses "Provenance-based fault tolerance technique recommendation for cloud-based scientific workflows: a practical approach." In this synthesis, that work is used as a source on problem definition within resilience patterns for data-intensive scientific workflows. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Salloum et al. (2016) addresses "Big data analytics on Apache Spark." In this synthesis, that work is used as a source on conceptual boundary within resilience patterns for data-intensive scientific workflows. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Pan et al. (2023) addresses "Data Provenance in Security and Privacy." In this synthesis, that work is used as a source on measurement within resilience patterns for data-intensive scientific workflows. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Together, these works provide complementary entry points, but their conclusions should not be flattened into a single ranking. Differences

in data generation, setting, temporal horizon, and outcome definition may be substantively meaningful. A review should therefore preserve those differences and state where analogy, rather than direct replication, connects one domain to another.

3. Mechanism and System Design

The operative mechanism is the dependencies among task graphs, data products, checkpoints, schedulers, provenance, and storage. This formulation discourages component-only reasoning. A model, platform, policy, assay, or infrastructure service acquires practical meaning only when its interfaces and use conditions are specified. The evidence chain should describe what information is available, what transformation occurs, which assumptions make that transformation credible, and who is authorized to act.

Versluis and Iosup (2021) addresses "A survey of domains in workflow scheduling in computing infrastructures: Community and keyword analysis, emerging trends, and taxonomies." In this synthesis, that work is used as a source on system mechanism within resilience patterns for data-intensive scientific workflows. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Akkaoui et al. (2020) addresses "EdgeMediChain: A Hybrid Edge Blockchain-Based Framework for Health Data Exchange." In this synthesis, that work is used as a source on representation choice within resilience patterns for data-intensive scientific workflows. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Senger and Geyer (2016) addresses "Parallel and distributed computing for Big Data applications." In this synthesis, that work is used as a source on failure analysis within resilience patterns for data-intensive scientific workflows. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

These studies also show why modular claims require interface tests. A component may perform well while the complete pathway fails because inputs drift, users interpret outputs differently, recovery semantics are ambiguous, or institutional incentives change. Design documentation should therefore include not only the intended pathway but also foreseeable deviations, degraded modes, and conditions under which the system should stop or defer.

4. Evaluation and Uncertainty

Evaluation should center on recovery point, recovery time, silent corruption, nondeterminism, and scientific equivalence after restart. Average performance remains useful, but it should be accompanied by distributions, error categories, relevant subgroups or operating regimes, and uncertainty at the point where a decision changes. Comparators must isolate the value of the proposed addition rather than compare a mature intervention with an intentionally weak baseline.

McKee et al. (2018) addresses "Survey of advances and challenges in intelligent autonomy for distributed cyber-physical systems." In this synthesis, that work is used as a source on comparative evaluation within resilience patterns for data-intensive scientific workflows. The

connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Pouya et al. (2016) addresses "Copernicus, a hybrid dataflow and peer-to-peer scientific computing platform for efficient large-scale ensemble sampling." In this synthesis, that work is used as a source on uncertainty within resilience patterns for data-intensive scientific workflows. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

External validation should introduce meaningful shifts in setting, time, population, workload, market structure, or environmental conditions. A nominally independent sample can still be too similar to test transfer. Conversely, a failed transfer is scientifically informative when the changed conditions are measured and the mechanism of failure is examined rather than hidden in an aggregate score.

5. Translation and Governance

Translation to long-running computational experiments on shared infrastructure depends on more than technical readiness. Decision rights, documentation, maintenance, monitoring, appeal, and recovery are part of the intervention. The governance requirement is failure injection, provenance validation, portable environments, and explicit completion criteria. Each control should be connected to a named failure mode and should identify an owner, evidence source, review interval, and escalation path.

Dai et al. (2025) addresses "State of the Art in Parallel and Distributed Systems: Emerging Trends and Challenges." In this synthesis, that work is used as a source on implementation within resilience patterns for data-intensive scientific workflows. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Zhang et al. (2016) addresses "Kira: Processing Astronomy Imagery Using Big Data Technology." In this synthesis, that work is used as a source on governance within resilience patterns for data-intensive scientific workflows. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for long-running computational experiments on shared infrastructure. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

A credible implementation plan also defines sunset and revalidation conditions. New data, software updates, institutional restructuring, population change, or shifts in incentives can invalidate previously reasonable assumptions. Monitoring should therefore test the validity of the evidence chain, not merely whether a service remains online or a headline metric remains stable.

6. Research Agenda

Future studies should preregister or otherwise predefine the intended decision, principal outcomes, exclusions, and analysis plan when feasible. They should report missingness, sensitivity analyses, negative or null findings, and the cost of errors to differently situated stakeholders. Reproducible artifacts should include sufficient metadata to reconstruct data provenance, model or analytical versions, parameter choices, and the sequence from evidence to conclusion.

Cross-disciplinary work needs a shared object of explanation rather than a collection of adjacent vocabularies. For resilience patterns for data-intensive scientific workflows, that shared object is the complete decision pathway. Technical, clinical, social, environmental, or economic expertise should each change the design or interpretation of the study. When evidence remains incomplete, the useful output is a bounded hypothesis, a transparent uncertainty statement, or a request for additional measurement.

7. Conclusion

The literature supports a disciplined approach to resilience patterns for data-intensive scientific workflows: define the decision, preserve system boundaries, test consequential failure modes, connect uncertainty to action, and assign accountable control. The strongest conclusion is not that one method is universally superior. It is that credible use in long-running computational experiments on shared infrastructure requires an auditable link from source evidence to design choice, evaluation result, and governed decision.

References

- Akkaoui, R., Hei, X., & Cheng, W. (2020). EdgeMediChain: A Hybrid Edge Blockchain-Based Framework for Health Data Exchange. *IEEE Access*, 8, 113467-113486. <https://doi.org/10.1109/access.2020.3003575>
- Dai, F., Hossain, M. A., & Wang, Y. (2025). State of the Art in Parallel and Distributed Systems: Emerging Trends and Challenges. *Electronics*, 14(4), 677. <https://doi.org/10.3390/electronics14040677>
- Guedes, T., Jesus, L. A., Ocaña, K. A. C. S., Drummond, L. M. A., & Oliveira, D. D. (2019). Provenance-based fault tolerance technique recommendation for cloud-based scientific workflows: a practical approach. *Cluster Computing*, 23(1), 123-148. <https://doi.org/10.1007/s10586-019-02920-6>
- McKee, D. W., Clement, S. J., Almutairi, J., & Xu, J. (2018). Survey of advances and challenges in intelligent autonomy for distributed cyber-physical systems. *CAAI Transactions on Intelligence Technology*, 3(2), 75-82. <https://doi.org/10.1049/trit.2018.0010>
- Pan, B., Stakhanova, N., & Ray, S. (2023). Data Provenance in Security and Privacy. *ACM Computing Surveys*, 55(14s), 1-35. <https://doi.org/10.1145/3593294>
- Pouya, I., Pronk, S., Lundborg, M., & Lindahl, E. (2016). Copernicus, a hybrid dataflow and peer-to-peer scientific computing platform for efficient large-scale ensemble sampling. *Future Generation Computer Systems*, 71, 18-31. <https://doi.org/10.1016/j.future.2016.11.004>
- Salloum, S., Dautov, R., Chen, X., Peng, P. X., & Huang, J. Z. (2016). Big data analytics on Apache Spark. *International Journal of Data Science and Analytics*, 1(3-4), 145-164. <https://doi.org/10.1007/s41060-016-0027-9>
- Senger, H., & Geyer, C. (2016). Parallel and distributed computing for Big Data applications. *Concurrency and Computation Practice and Experience*, 28(8), 2412-2415. <https://doi.org/10.1002/cpe.3813>
- Versluis, L., & Iosup, A. (2021). A survey of domains in workflow scheduling in computing infrastructures: Community and keyword analysis, emerging trends, and taxonomies. *Future Generation Computer Systems*, 123, 156-177. <https://doi.org/10.1016/j.future.2021.04.009>
- Zhang, Z., Barbary, K., Nothaft, F. A., Sparks, E. R., Zahn, O., Franklin, M. J., Patterson, D. A., & Perlmutter, S. (2016). Kira: Processing Astronomy Imagery Using Big Data Technology. *IEEE Transactions on Big Data*, 6(2), 369-381. <https://doi.org/10.1109/tbdata.2016.2599926>