

Performance comparison of Pegasus on different distributed computing infrastructures.

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

Compound of Interest

Compound Name: Pegasus

Cat. No.: B039198

[Get Quote](#)

Performance Showdown: Apache Pegasus Across Distributed Computing Landscapes

A Comparative Guide for Researchers and Drug Development Professionals

In the realm of large-scale data management, Apache **Pegasus**, a distributed key-value storage system, has emerged as a noteworthy contender, aiming to bridge the gap between in-memory solutions like Redis and disk-based systems like HBase.^[1] For researchers, scientists, and drug development professionals grappling with massive datasets, understanding how **Pegasus** performs across different distributed computing infrastructures is paramount for making informed architectural decisions. This guide provides a comprehensive performance comparison of Apache **Pegasus** on bare metal, with qualitative insights into its expected performance on virtualized and containerized environments, supported by experimental data and detailed methodologies.

At a Glance: Pegasus Performance Metrics

The following tables summarize the performance of Apache **Pegasus** based on publicly available benchmark data. It is important to note that direct, officially published comparative benchmarks of **Pegasus** across different infrastructures are not readily available. The data presented here is from a benchmark conducted on a bare-metal cluster, which serves as a baseline for understanding its capabilities.

Table 1: Write-Only Workload Performance on Bare Metal

Threads (Clients * Threads)	Read/Write Ratio	Write QPS (Queries Per Second)	Average Write Latency (µs)	P99 Write Latency (µs)
3 * 15	0:1	56,953	787	1,786

Source: Apache **Pegasus** Benchmark[2]

Table 2: Read-Only Workload Performance on Bare Metal

Threads (Clients * Threads)	Read/Write Ratio	Read QPS (Queries Per Second)	Average Read Latency (µs)	P99 Read Latency (µs)
3 * 50	1:0	360,642	413	984

Source: Apache **Pegasus** Benchmark[2]

Table 3: Mixed Read/Write Workload Performance on Bare Metal

Threads (Clients * Threads)	Read/W rite Ratio	Read QPS	Avg. Read Latency (µs)	P99 Read Latency (µs)	Write QPS	Avg. Write Latency (µs)	P99 Write Latency (µs)
3 * 30	1:1	62,572	464	5,274	62,561	985	3,764
3 * 15	1:3	16,844	372	3,980	50,527	762	1,551

Source: Apache **Pegasus** Benchmark[2]

Experimental Protocols: The Bare Metal Benchmark

The performance data presented above was obtained from a benchmark conducted by the Apache **Pegasus** community. The methodology employed provides a transparent and reproducible framework for performance evaluation.

Hardware Specifications:

- CPU: Intel® Xeon® Silver 4210 @ 2.20 GHz (2 sockets)
- Memory: 128 GB
- Disk: 8 x 480 GB SSD
- Network: 10 Gbps

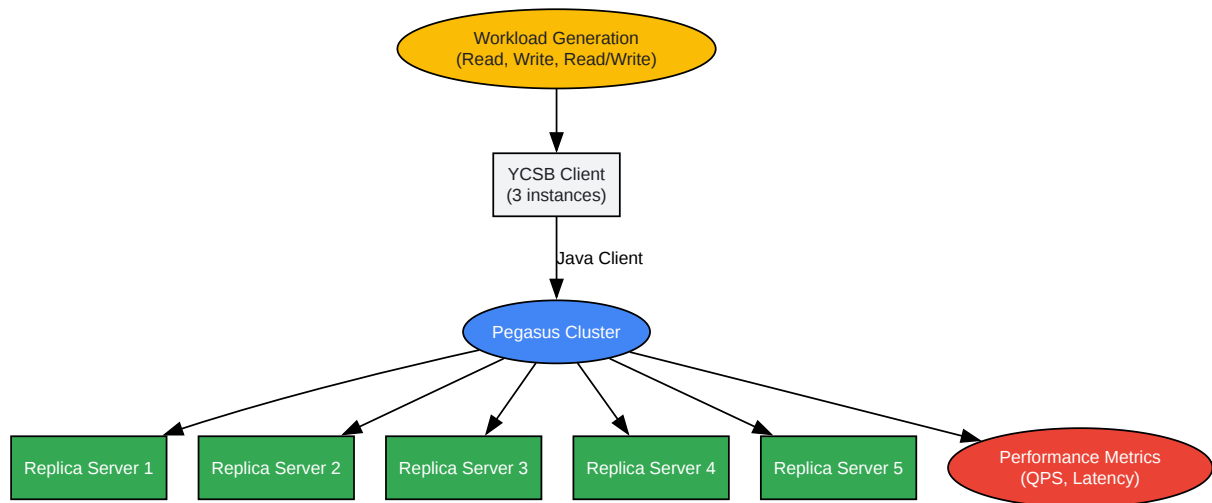
Cluster Configuration:

- Replica Server Nodes: 5
- Test Table Partitions: 64

Benchmarking Tool:

The Yahoo! Cloud Serving Benchmark (YCSB) was used to generate the workloads, utilizing the **Pegasus** Java client.^[2] The request distribution was set to Zipfian, which models a more realistic scenario where some data is accessed more frequently than others.^[2]

The experimental workflow for this benchmark can be visualized as follows:



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

A diagram illustrating the experimental workflow for the Apache **Pegasus** bare-metal benchmark.

Performance on Other Distributed Infrastructures: A Qualitative Analysis

While direct, quantitative benchmarks of Apache **Pegasus** on virtualized and containerized platforms are not readily available in official documentation, we can infer its likely performance characteristics based on the behavior of its underlying storage engine, RocksDB, and general industry observations of these infrastructures.

Virtual Machines (VMs)

Deploying **Pegasus** on a virtualized infrastructure would likely introduce a performance overhead compared to the bare-metal baseline. This is due to the hypervisor layer, which manages access to the physical hardware. For a high-performance, I/O-intensive application like **Pegasus**, which relies on the speed of its underlying SSDs, this virtualization layer can introduce latency.

However, modern virtualization technologies have significantly reduced this overhead. The performance impact would largely depend on the specific hypervisor, the configuration of the virtual machines (e.g., dedicated vs. shared resources), and the underlying hardware. For many use cases, the flexibility, scalability, and resource management benefits of VMs may outweigh the modest performance trade-off.

Kubernetes and Containers


Running **Pegasus** within containers, orchestrated by a platform like Kubernetes, presents an interesting performance profile. Containers offer a more lightweight form of virtualization than traditional VMs, sharing the host operating system's kernel. This generally results in lower overhead and near-native performance.

The performance of **Pegasus** on Kubernetes would be influenced by several factors:

- **Networking:** The choice of Container Network Interface (CNI) plugin in Kubernetes can impact network latency and throughput, which are critical for a distributed database.
- **Storage:** The performance of persistent storage in Kubernetes, managed through Container Storage Interface (CSI) drivers, would directly affect **Pegasus**'s I/O performance. Utilizing high-performance storage classes backed by SSDs is crucial.
- **Resource Management:** Kubernetes' resource allocation and scheduling capabilities can impact the consistent performance of **Pegasus** nodes. Properly configured resource requests and limits are essential to avoid contention and ensure predictable performance.

Given that **Pegasus** is designed for horizontal scalability, Kubernetes, with its robust scaling and management features, could be a compelling platform for deploying and operating a **Pegasus** cluster, especially in dynamic and large-scale environments. The performance is expected to be very close to that of a bare-metal deployment, provided the underlying infrastructure and Kubernetes configurations are optimized for high-performance stateful workloads.

The logical relationship for deploying **Pegasus** on different infrastructures can be visualized as follows:



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

A diagram showing the deployment relationship of Apache **Pegasus** across different infrastructures.

Conclusion

The benchmark data from the bare-metal deployment demonstrates that Apache **Pegasus** can achieve high throughput and low latency for both read- and write-intensive workloads. While direct comparative data on other infrastructures is not yet available, a qualitative analysis suggests that:

- Bare Metal offers the highest potential performance by providing direct access to hardware resources.
- Virtual Machines provide flexibility and manageability with a potential for a slight performance overhead.
- Kubernetes and Containers offer a compelling balance of near-native performance, scalability, and operational efficiency for managing distributed **Pegasus** clusters.

For researchers and professionals in drug development, the choice of infrastructure will depend on the specific requirements of their pipelines, balancing the need for raw performance with considerations of scalability, ease of management, and cost-effectiveness. As more community-driven benchmarks become available, a more granular, quantitative comparison will

be possible. For now, the strong performance of **Pegasus** on bare metal provides a solid indication of its potential across a variety of modern distributed computing environments.

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. GitHub - apache/incubator-pegasus: Apache Pegasus - A horizontally scalable, strongly consistent and high-performance key-value store [github.com]
- 2. Pegasus | Benchmark [pegasus.apache.org]
- To cite this document: BenchChem. [Performance comparison of Pegasus on different distributed computing infrastructures.]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b039198#performance-comparison-of-pegasus-on-different-distributed-computing-infrastructures]

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