Benchmarking the supercomputing capabilities of Los Alamos against international standards.

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

Compound of Interest	
Compound Name: Alamos	
Cat. No.: B1666505	Get Quote

Los Alamos Supercomputing on the World Stage: A Comparative Analysis

For Immediate Release

LOS **ALAMOS**, N.M., December 14, 2025 – Los **Alamos** National Laboratory (LANL), a cornerstone of scientific and national security research, continues to advance its supercomputing capabilities with the recent deployment of innovative systems. This guide provides a comparative analysis of LANL's supercomputing power against top-ranked international systems, offering researchers, scientists, and drug development professionals a clear perspective on the current high-performance computing (HPC) landscape.

Los **Alamos** is actively deploying its next-generation supercomputer, Crossroads, which is anticipated to deliver a four- to eight-fold performance improvement on critical national security simulations compared to its predecessor, Trinity.[1][2][3] The laboratory has also introduced Venado, a system designed to accelerate artificial intelligence and machine learning workloads, leveraging the power of NVIDIA's Grace Hopper Superchips.[4][5] Furthermore, LANL has announced two new supercomputers, Mission and Vision, to be developed in partnership with HPE and NVIDIA, further solidifying its commitment to leadership in the exascale and AI computing era.[6][7]

This analysis benchmarks LANL's systems against the most powerful supercomputers globally, as ranked by the November 2024 TOP500 and HPCG (High-Performance Conjugate

Tech Support

Gradients) lists.

Global Supercomputing Landscape: A Quantitative Comparison

The following table summarizes the key performance metrics of leading international supercomputers and provides context for the capabilities of Los **Alamos**'s systems. While official benchmark results for Venado and Crossroads are not yet publicly available, their architectural details and expected performance offer valuable insights.

Rank (TOP500 Nov. 2024)	System Name	Institution	Country	Processor Architecture	Total Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	HPCG (PFlop/s)
1	El Capitan	Lawrence Livermore National Laboratory	USA	HPE Cray EX, AMD EPYC 4th Gen, AMD Instinct MI300A	11,039,616	1742.0	2746.4	N/A
2	Frontier	Oak Ridge National Laboratory	USA	HPE Cray EX, AMD EPYC 3rd Gen, AMD Instinct MI250X	8,699,904	1194.0	1679.8	14.05
3	Aurora	Argonne National Laboratory	USA	HPE Cray EX, Intel Xeon CPU Max, Intel Data Center	9,389,312	1012.0	1980.0	5.61

				GPU Max				
4	Eagle	Microsoft Azure	USA	NDv5, Intel Xeon Platinum 8480C, NVIDIA H100	1,128,960	561.2	846.8	N/A
5	Fugaku	RIKEN Center for Computational Science	Japan	Fujitsu, A64FX	7,630,848	442.0	537.2	16.0
-	Venado	Los Alamos National Laboratory	USA	HPE Cray EX, NVIDIA Grace Hopper Superchip	N/A	N/A	N/A	N/A
-	Crossroads	Los Alamos National Laboratory	USA	HPE Cray, Intel Sapphire Rapids	N/A	N/A	N/A	N/A

Note: Rmax and Rpeak values are based on the LINPACK benchmark. HPCG scores provide a different perspective on performance for a wider range of applications. "N/A" indicates that the data is not available or not yet submitted.

Experimental Protocols: Understanding the Benchmarks

The performance of supercomputers is evaluated using standardized benchmarks. The two primary benchmarks referenced in this guide are the High-Performance LINPACK (HPL) and the High-Performance Conjugate Gradients (HPCG) benchmarks.

High-Performance LINPACK (HPL) Benchmark

The LINPACK benchmark is the industry standard for ranking the world's fastest supercomputers on the TOP500 list.[8]

- Objective: To measure a system's floating-point computing power by solving a dense system of linear equations (Ax = b).[8][9]

- Methodology: The benchmark involves generating a large, random, dense matrix (A) and a corresponding vector (b) and then solving for the unknown vector (x). The primary computational task is the LU factorization of the matrix A with partial pivoting.[10] The size of the matrix is chosen to be as large as possible to fully utilize the system's memory and computational resources. The performance is measured in floating-point operations per second (FLOPS).[8]

- Significance: HPL is a measure of a system's ability to perform computationally intensive tasks and is a good indicator of its peak performance. However, it may not be representative of the performance of a broader range of scientific applications that involve more complex data access patterns.[10]

High-Performance Conjugate Gradients (HPCG) Benchmark
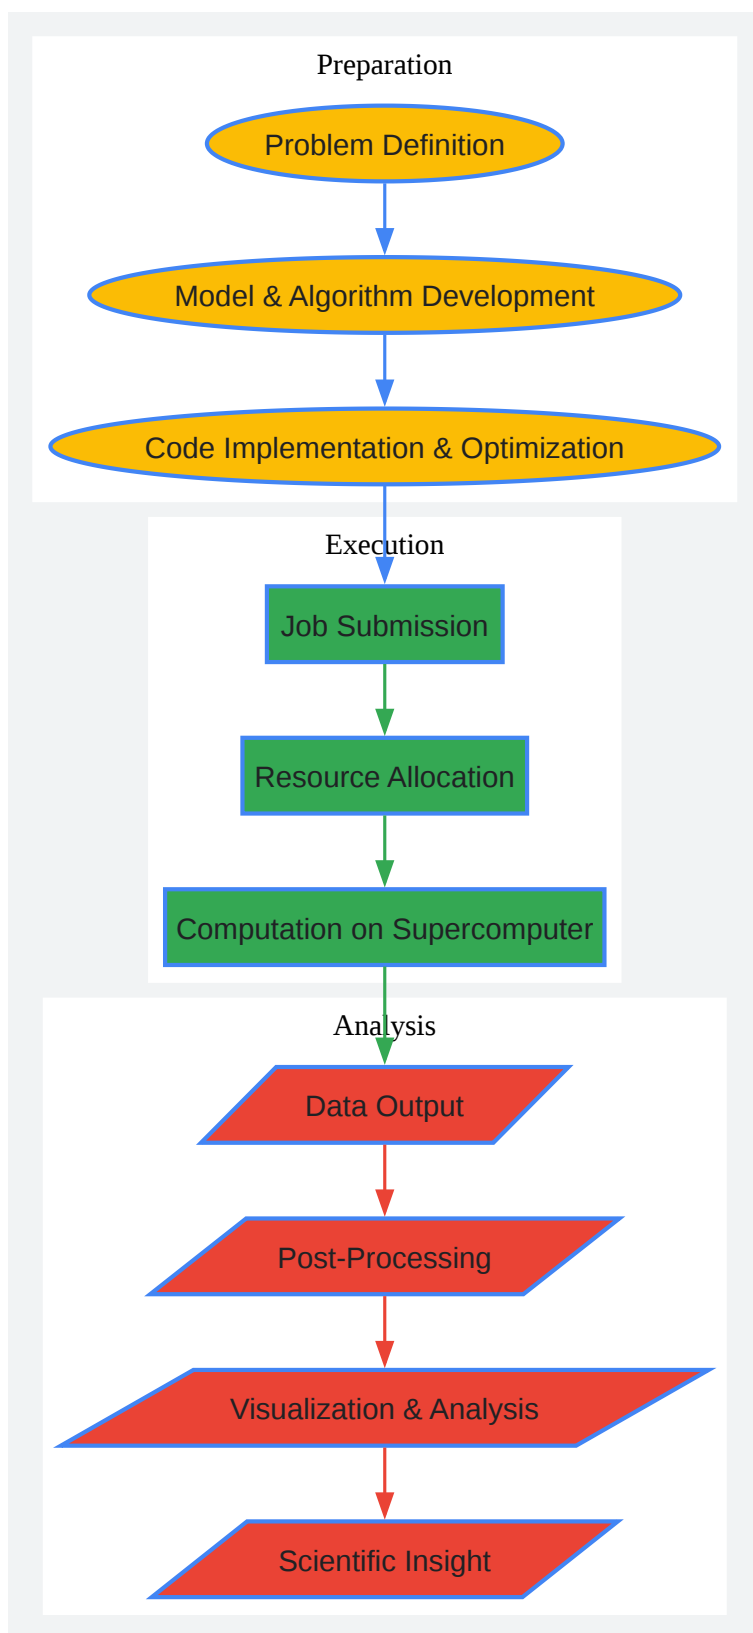
The HPCG benchmark is designed to complement the HPL benchmark by assessing the performance of a system on a different class of problems that are more representative of real-world scientific and engineering applications.[11]
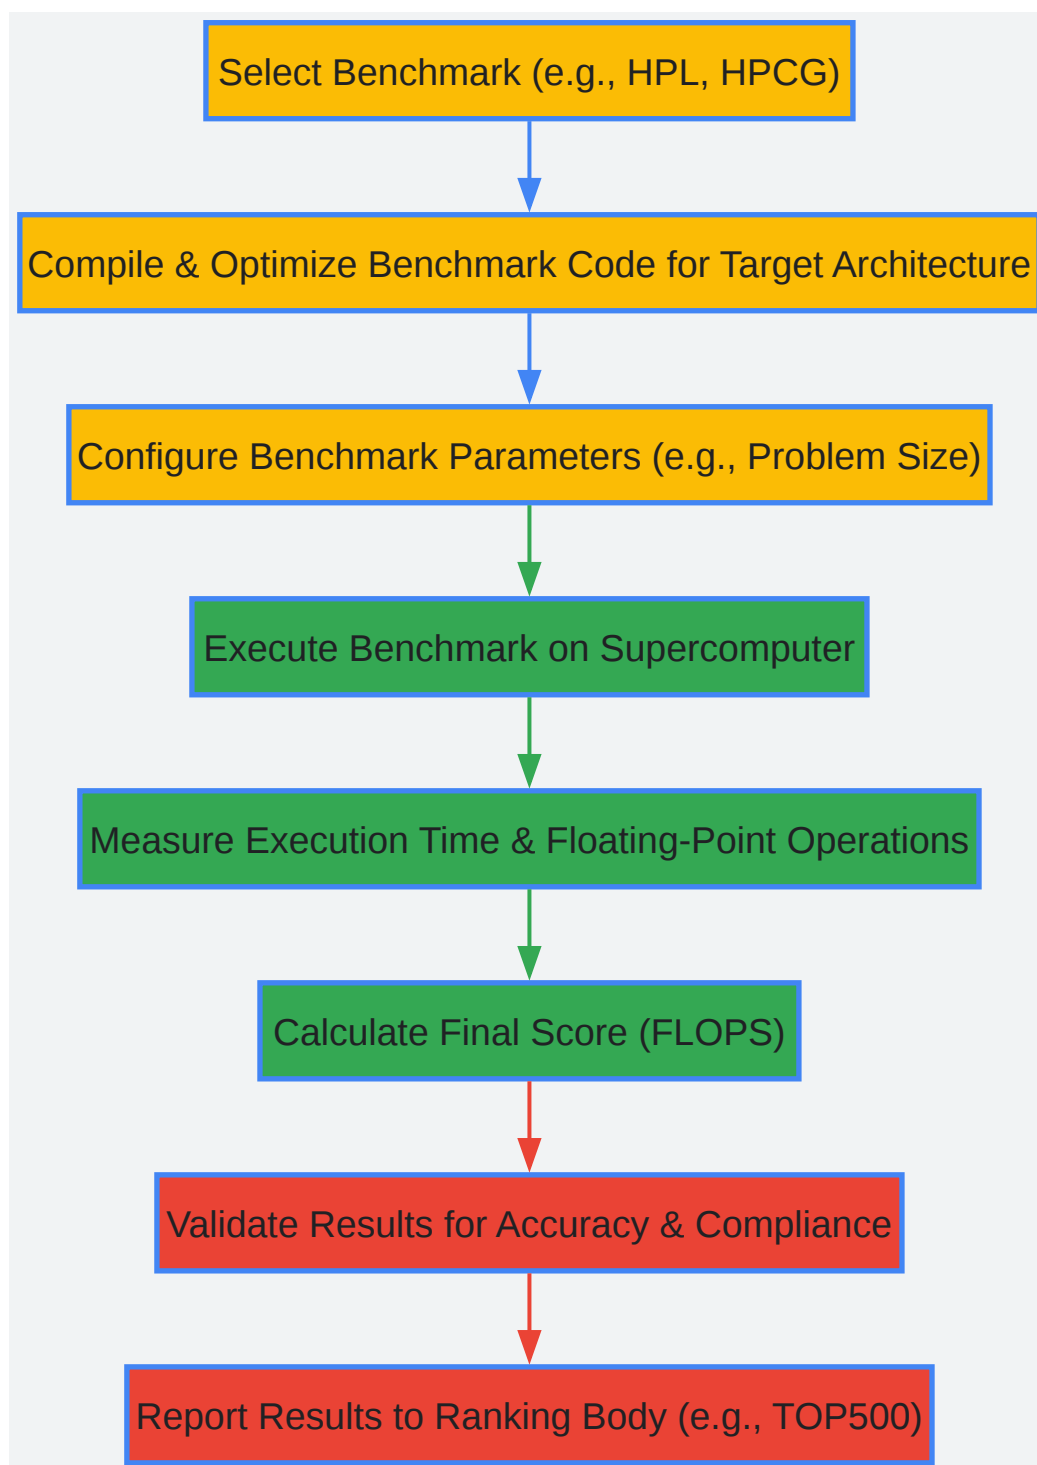
- Objective: To measure the performance of a system while solving a sparse system of linear equations using a preconditioned conjugate gradient (PCG) method.[11][12]

- Methodology: The benchmark generates a sparse linear system that mimics the computational kernels found in many scientific simulations, such as those involving finite element or finite volume methods.[12] The PCG algorithm involves a series of sparse matrix-vector multiplications, vector updates, and dot products, which stress the memory and interconnect subsystems of a supercomputer.[11][13] The performance is reported in FLOPS, but the values are typically a small fraction of the HPL results due to the memory-bound nature of the computation.

- Significance: HPCG provides a more realistic measure of a system's performance on a wider range of applications, particularly those that are limited by memory bandwidth and network latency rather than raw floating-point capability.[14]

Visualizing High-Performance Computing Workflows

To better understand the processes involved in utilizing these powerful machines, the following diagrams illustrate a typical experimental workflow and the logical flow of the benchmarking process.

Preparation

Problem Definition

Model & Algorithm Development

Code Implementation & Optimization

Execution

Job Submission

Resource Allocation

Computation on Supercomputer

Analysis

Data Output

Post-Processing

Visualization & Analysis

Scientific Insight

```
┌──────────────────────────────────────┐
│   Select Benchmark (e.g., HPL, HPCG)  │
└──────────────────────────────────────┘
                    │
                    ▼
┌──────────────────────────────────────┐
│ Compile & Optimize Benchmark Code for │
│          Target Architecture          │
└──────────────────────────────────────┘
                    │
                    ▼
┌──────────────────────────────────────┐
│ Configure Benchmark Parameters        │
│ (e.g., Problem Size)                  │
└──────────────────────────────────────┘
                    │
                    ▼
┌──────────────────────────────────────┐
│   Execute Benchmark on Supercomputer  │
└──────────────────────────────────────┘
                    │
                    ▼
┌──────────────────────────────────────┐
│ Measure Execution Time &              │
│ Floating-Point Operations             │
└──────────────────────────────────────┘
                    │
                    ▼
┌──────────────────────────────────────┐
│      Calculate Final Score (FLOPS)    │
└──────────────────────────────────────┘
                    │
                    ▼
┌──────────────────────────────────────┐
│ Validate Results for Accuracy &       │
│ Compliance                            │
└──────────────────────────────────────┘
                    │
                    ▼
┌──────────────────────────────────────┐
│ Report Results to Ranking Body        │
│ (e.g., TOP500)                        │
└──────────────────────────────────────┘
```

Click to download full resolution via product page

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopiclabeling.

Email: info@benchchem.com or Request Quote Online.

References

- 1. executivebiz.com [executivebiz.com]

- 2. losalamosreporter.com [losalamosreporter.com]

- 3. potomacofficersclub.com [potomacofficersclub.com]

- 4. hpcwire.com [hpcwire.com]

- 5. Venado installation opens the door for AI applications at Los Alamos | LANL [lanl.gov]

- 6. lanl.gov [lanl.gov]

- 7. NNSA Announces Two New Supercomputers at Los Alamos National Laboratory, Propelling America's Leadership in Science | Department of Energy [energy.gov]

- 8. LINPACK benchmarks - Wikipedia [en.wikipedia.org]

- 9. LINPACK Benchmarks | Encyclopedia MDPI [encyclopedia.pub]

- 10. netlib.org [netlib.org]

- 11. HPCG Benchmark [hpcg-benchmark.org]

- 12. upcommons.upc.edu [upcommons.upc.edu]

- 13. hpcf-files.umbc.edu [hpcf-files.umbc.edu]

- 14. insidehpc.com [insidehpc.com]

- To cite this document: BenchChem. [Benchmarking the supercomputing capabilities of Los Alamos against international standards.]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b1666505#benchmarking-the-supercomputing-capabilities-of-los-alamos-against-international-standards]

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