

# GEM-5 tutorial for beginners in academic research

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

## Compound of Interest

Compound Name: GEM-5

Cat. No.: B12410503

[Get Quote](#)

An In-Depth Guide to **GEM-5** for Academic Research

## Abstract

The gem5 simulator is a modular and flexible platform for computer architecture research, enabling detailed performance analysis of complex hardware systems.<sup>[1]</sup> This guide provides a comprehensive introduction for researchers and scientists new to gem5. It covers the fundamental concepts, initial setup, simulation modes, and a typical research workflow. Detailed protocols for key experiments are provided, along with visualizations of core concepts and workflows to facilitate understanding. The document aims to equip beginners with the necessary knowledge to start using gem5 effectively in their academic research.

## Introduction to GEM-5

gem5 is a modular, discrete-event driven computer system simulator platform.<sup>[2]</sup> Its key characteristics make it an invaluable tool for academic research:

- **Modularity:** gem5 is composed of interchangeable components, known as SimObjects, which can be configured, extended, or replaced to model novel architectures.<sup>[2][3]</sup>
- **Flexibility:** It supports multiple Instruction Set Architectures (ISAs) like X86, ARM, and RISC-V.<sup>[2]</sup>

- Dual Simulation Modes: gem5 offers two primary simulation modes: Syscall Emulation (SE) and Full System (FS), catering to different research needs.[\[2\]](#)[\[4\]](#)
- Collaborative and Open-Source: As a widely-used, open-source project, it benefits from a large community of academic and industrial contributors.[\[1\]](#)

The simulator's architecture is primarily based on C++ for performance-critical models and Python for configuration and control, a separation that allows researchers to easily define and modify complex systems.[\[2\]](#)[\[3\]](#)

## Getting Started: The Initial Setup

The first and often most challenging step for beginners is setting up the simulation environment.[\[5\]](#) This protocol outlines the standard procedure.

### Experimental Protocol 1: Environment Setup

This protocol details the steps to get a working build of gem5 on a Unix-like operating system.

- System Requirements: Ensure your host system is a Unix-like OS (Linux is recommended) with necessary dependencies installed, such as git, scons, g++, and Python development libraries.[\[6\]](#)[\[7\]](#)
- Clone the Repository: Download the gem5 source code from its official repository using git.[\[6\]](#)[\[7\]](#)
- Compile **GEM-5**: Use scons to build the simulator. The build target specifies the ISA and the optimization level. Building can be time-consuming and memory-intensive.[\[6\]](#) Using multiple threads with the -j flag is recommended on multi-core machines.[\[7\]](#)[\[8\]](#)

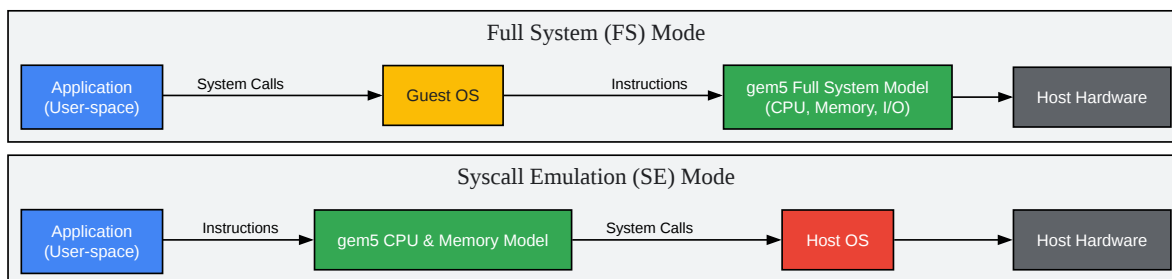
Table 1: Common **GEM-5** Build Targets

Target Suffix	Description	Use Case
.opt	An optimized build with debugging symbols.	General use, balancing performance and debuggability.[4]
.fast	A highly optimized build with no debugging symbols.	Maximum simulation speed for large-scale experiments.[8]
.debug	A build with full debugging symbols and no optimizations.	Development and debugging of new models.[9]

## Understanding GEM-5 Simulation Modes

gem5 provides two main modes of operation, each with distinct advantages and use cases.[4]

- **Syscall Emulation (SE) Mode:** In SE mode, gem5 simulates only the user-space instructions of an application.[10] System calls are trapped and handled by the host operating system.[8] This mode is faster and simpler to configure, making it ideal for research focused on CPU and memory subsystem performance without the complexity of a full OS.[4][9]
- **Full System (FS) Mode:** FS mode simulates a complete hardware system, including CPUs, caches, memory, and I/O devices.[11] This allows it to boot an unmodified operating system and run a full software stack.[11] While more complex to set up—requiring a compiled kernel and disk image—it is essential for research involving OS interactions or complex I/O behavior.[11]



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

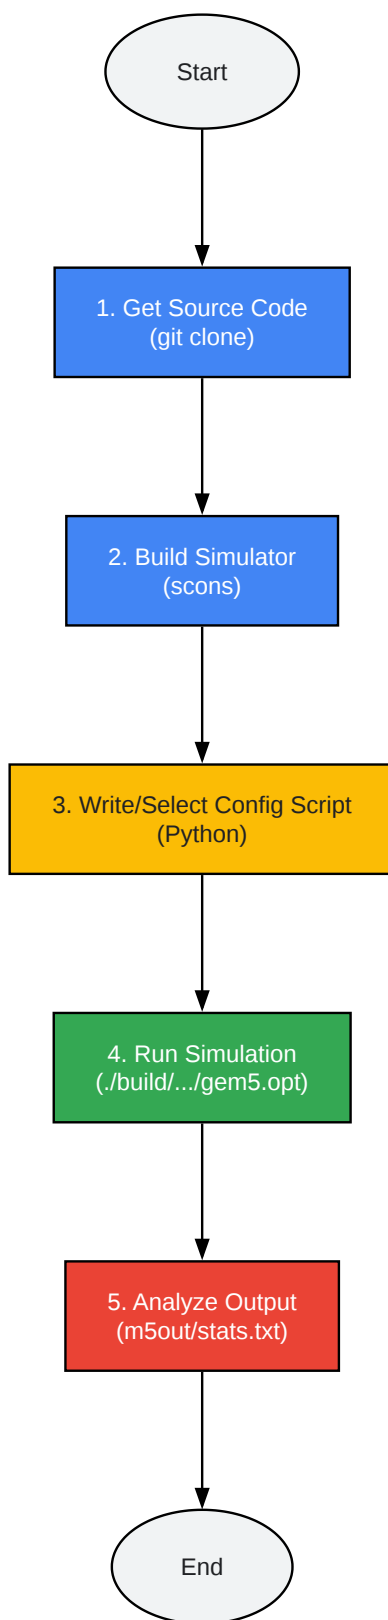
**Caption:** Comparison of SE and FS mode execution flows.

## Running Your First Simulation (SE Mode)

This protocol guides you through running a simple "Hello World" application in SE mode, which is the typical starting point for new users.<sup>[6]</sup>

### Experimental Protocol 2: Executing a "Hello World" Program

- **Identify the Configuration Script:** gem5 uses Python scripts for simulation configuration. For this test, we use a basic script provided with the source code: `configs/learning_gem5/part1/simple.py`.<sup>[6]</sup> This script defines a simple system with a CPU and memory.
- **Prepare the Command:** The gem5 executable takes the configuration script as an argument. The script, in turn, may take its own options. For this test, no additional options are needed.
- **Execute the Simulation:** Run the following command from the root of the gem5 directory.
- **Inspect the Output:** After execution, gem5 creates an output directory named `m5out/`.<sup>[4]</sup> This directory contains simulation statistics, configuration details, and any standard output from the simulated program. You should see "Hello world!" printed in the terminal output.<sup>[6]</sup>



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

**Caption:** A high-level overview of the standard gem5 simulation workflow.

## Analyzing Simulation Output

The primary source of quantitative data from a gem5 simulation is the stats.txt file located in the m5out directory.<sup>[4][8]</sup> This file contains a detailed breakdown of various metrics for every SimObject in the simulation.

Table 2: Example Performance Metrics from stats.txt

Statistic Name	Description	Common Use
simSeconds	The total simulated time in seconds.	Overall simulation runtime.
system.cpu.numCycles	The number of CPU cycles simulated.	Core performance measurement.
system.cpu.committedInsts	The number of instructions committed by the CPU.	Instructions Per Cycle (IPC) calculation.
system.cpu.dcache.overallMisses	The total number of misses in the L1 data cache.	Memory access pattern analysis.
system.mem_ctrls.num_reads	The number of read requests to the memory controller.	Memory bandwidth analysis.

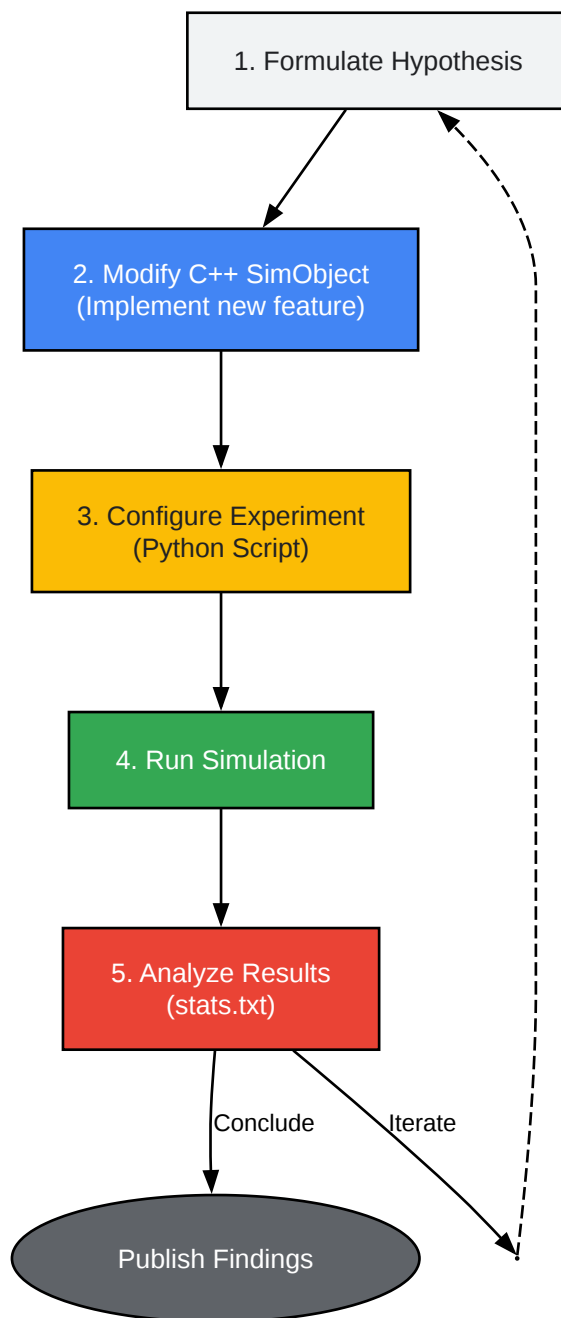
## A Typical Academic Research Workflow

Using gem5 in academic research is an iterative process that involves modifying the simulator to model a novel idea, running experiments, and analyzing the results.

The workflow typically involves:

- **Hypothesis Formulation:** Define a new architectural feature or optimization to be evaluated.
- **Model Development:** Modify the gem5 C++ source code to implement the new feature, often by creating or extending a SimObject.<sup>[2]</sup>
- **Experiment Configuration:** Write or adapt a Python configuration script to integrate and parameterize the new model within a larger system.

- Simulation Execution: Run a set of benchmarks or workloads on the modified simulator.
- Data Analysis: Analyze the stats.txt output to quantify the performance, power, or thermal impact of the proposed change.
- Iteration: Refine the model based on the analysis and repeat the process.



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

**Caption:** The iterative cycle of academic research using gem5.

## Conclusion

gem5 is a powerful and essential tool for modern computer architecture research. While it has a steep learning curve, its modularity and extensive capabilities provide unparalleled opportunities for exploration and innovation. By following a structured approach—starting with environment setup, understanding the core simulation modes, and running simple experiments—beginners can build the foundational knowledge required to tackle complex research questions. For continued learning, the official gem5 documentation and community resources are invaluable.<sup>[2]</sup><sup>[12]</sup>

### 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. gem5.org [gem5.org]
- 2. gem5: Learning gem5 [gem5.org]
- 3. youtube.com [youtube.com]
- 4. developer.arm.com [developer.arm.com]
- 5. gem5.org [gem5.org]
- 6. gem5: Getting Started with gem5 [gem5.org]
- 7. Gem5 and GS Gem5-Validate Tutorial [web-archive.southampton.ac.uk]
- 8. cs.pomona.edu [cs.pomona.edu]
- 9. scribd.com [scribd.com]
- 10. research.cs.wisc.edu [research.cs.wisc.edu]
- 11. gem5 Full System Simulation — gem5 Tutorial 0.1 documentation [lowepower.com]
- 12. gem5: gem5 documentation [gem5.org]
- To cite this document: BenchChem. [GEM-5 tutorial for beginners in academic research]. BenchChem, [2025]. [Online PDF]. Available at:



[<https://www.benchchem.com/product/b12410503#gem-5-tutorial-for-beginners-in-academic-research>]

---

**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)