

# Designing a Software Performance Engineering Laboratory

Wasim Haque 

Woodstock, GA, USA

Email: haque.wasim@gmail.com

**How to cite this paper:** Haque, W. (2025) Designing a Software Performance Engineering Laboratory. *Journal of Software Engineering and Applications*, 18, 87-97. <https://doi.org/10.4236/jsea.2025.183006>

**Received:** February 1, 2025

**Accepted:** March 17, 2025

**Published:** March 20, 2025

Copyright © 2025 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0). <http://creativecommons.org/licenses/by/4.0/>



Open Access

---

## Abstract

In today's fast-paced digital landscape, software applications are under tremendous pressure to deliver smooth and lightning-quick user experiences. Any delay in response time, unexpected crashes, or scalability challenges can significantly impact user satisfaction and harm a company's reputation. This is where Software Performance Engineering (SPE) becomes crucial. SPE is a specialized field focused on optimizing an application's performance throughout its entire lifecycle. From initial design through development and testing, SPE experts apply structured methodologies to ensure the software consistently meets high-performance standards. A key element in a successful SPE is the creation of a dedicated performance engineering lab. This specialized environment enables engineers to thoroughly analyze and test applications under a variety of conditions, replicating real-world scenarios to identify potential bottlenecks. Within the lab, engineers leverage advanced tools and techniques to assess key performance metrics, troubleshoot issues, and enhance the application's efficiency. This article provides a detailed guide to designing and managing a performance engineering lab, including essential aspects such as hardware and software selection, network setup, and testing strategies. It also outlines best practices for lab operations, ensuring that the lab continues to be a vital resource for achieving optimal software performance.

## Keywords

Software Performance Engineering, Software Reliability, Scalability, Performance Engineering Lab, Cloud Computing, Microsoft Azure, User Experience, Load Testing, Stress Testing

---

## 1. Introduction

Performance Engineering Lab is a dedicated environment where software is systematically tested, analyzed, and optimized for performance metrics such as speed,

scalability, stability, and resource efficiency. The lab provides the tools, infrastructure, and processes necessary to conduct thorough performance evaluations, ensuring that software applications meet required performance benchmarks before being released into production for customers to use.

For this paper, we will use the following entities to architect a performance engineering lab. Please note that the design of this PE Lab is agnostic of the application under test and the various tools/technologies that build it. The only reason names of tools and technologies are mentioned here is for the readers' ease so that they can understand the ecosystem.

- An environment built in the cloud, say in MS Azure.
- A software requirement-gathering tool like JIRA.
- A performance testing tool like JMeter [1].
- An orchestration tool like Jenkins or any CI/CD (Continuous Integration/Continuous Deployment) tool.
- A source code control tool like GitHub or GitLab, etc.
- An application performance monitoring (APM) like Datadog or Dynatrace.
- Any reporting or communication tool like Slack or MS Teams.

### **Need for a PE Lab**

First, determine the purpose of this Performance Engineering (PE) Lab.

- What specific performance or load tests do you plan to conduct?
- Which software and hardware metrics or Key Performance Indicators (KPIs) will you capture?
- How will you analyze the results to determine if your software product meets performance requirements and customer expectations?

Some of the common reasons organizations need a SPE Lab are [2]:

- Simulate real-world usage scenarios and stress-test the application.
- Identify bottlenecks that may not surface during functional testing.
- Measure throughput, latency, and resource usage (CPU, memory, disk I/O).
- Reducing the chances of poor user experience due to slow or unresponsive systems.
- Detecting issues related to concurrency, memory leaks, or resource contention that may degrade performance over time.
- Avoid over-provisioning or under-provisioning resources.
- Optimize configurations to achieve higher performance on existing infrastructure, reducing operational costs.
- Plan for future growth by understanding how the software scales with increasing load.
- Benchmark products against competitors or industry standards.
- Ensures high reliability and resilience, particularly for software that operates in high-demand or high-availability environments.
- Ensure that software meets contractual performance obligations (SLAs) for clients.
- Provide documentation and proof of performance for compliance and audits.

## 2. SPE Lab Ecosystem

A Software Performance Engineering (SPE) Laboratory focuses on analyzing, optimizing, and ensuring the performance of software systems. The technology ecosystem within an SPE lab includes tools, methodologies, and frameworks that support performance evaluation, monitoring, and tuning of software applications.

“Reference [3] Cloud computing introduces new challenges in performance management due to its dynamic, distributed, and multi-tenant nature”. A Software Performance Engineering (SPE) Lab for cloud systems is essential to ensure scalability, reliability, and efficiency in cloud environments. In this paper, we’ll focus on how we can create an effective SPE Lab for complex cloud systems. Let’s dive into the various components of a SPE Lab.

### 2.1. Cloud Application to Be Tested

A typical cloud-based software application that undergoes performance testing and engineering has a specific architecture. Microsoft, Amazon, and Google are some of the cloud services providers that help organizations build modern-day enterprise-level software by offering services for hosting software, data streaming, databases, caching, application security, and, last but not least, a secure network. An example of such an application using Microsoft Azure’s services is shown below [4] (Figure 1).

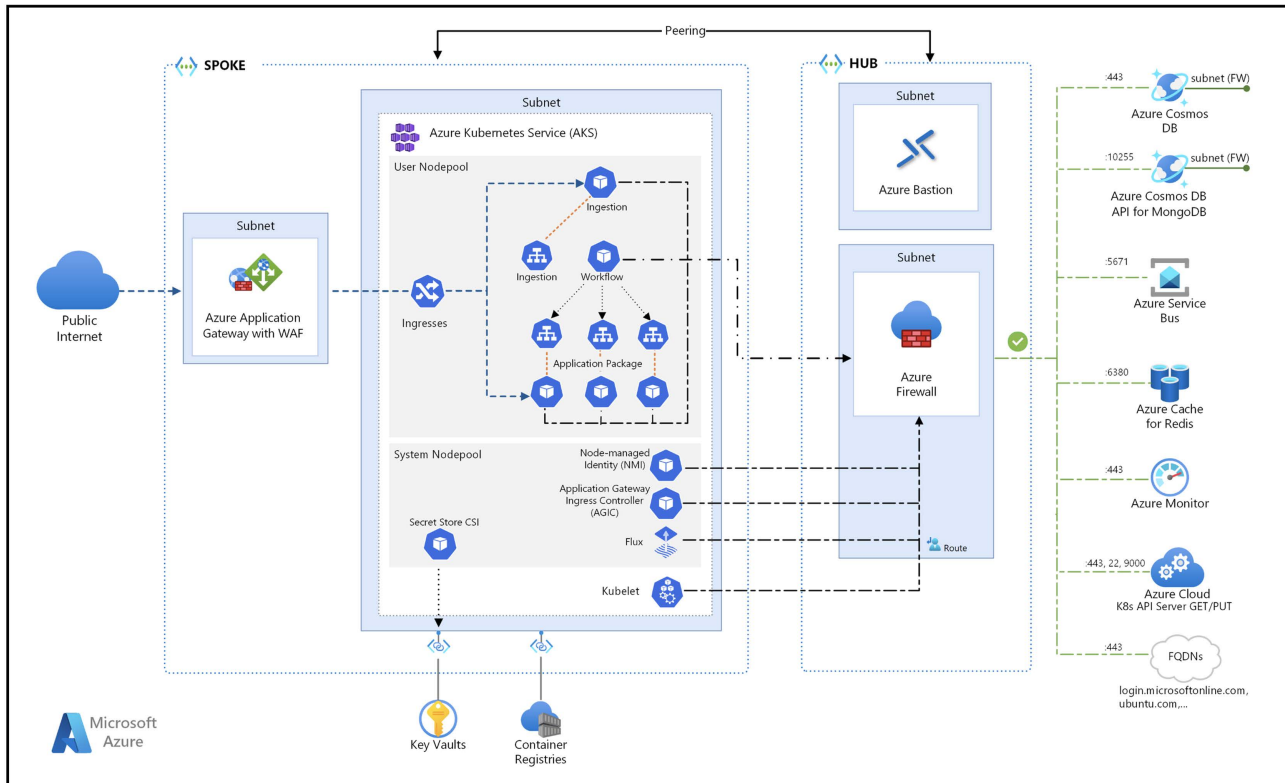


Figure 1. Example of an application using Microsoft Azure.

There are few important guidelines that should be followed while creating your performance engineering environment:

- This environment should be an exact replica of what your production environment already is or is going to be in the future so that it is easier to simulate real-life scenarios [5].
- Performance environments should have the same configurations and settings as production.
- Databases and microservices infrastructure should be scalable and sized either the same as production or scaled appropriately if cost is a concern.
- Make sure the lab environment has all the cloud services (caching, messaging, and streaming) built into it, the same as production, and are appropriately sized.
- The SPE Lab should have the same network topology and security standards as production. For example, if the production environment is behind Cloudflare, then the SPE Lab must be behind Cloudflare too.
- Tenant and data modelling are also important factors to consider when setting up the environment.
- Have at least 3 tenants, small, medium, and large, for varied performance use cases and model your tests around those.
- Pump test data, such as user and application-specific data, into these tenants. As a best practice, always have at least 25% more data than your biggest production tenant to prepare you for the long haul.

## 2.2. Requirement Management System

Think of a “requirement” as simply something your project must deliver. This could be a specific feature in a software product or a broader goal like increasing customer satisfaction. Before starting any project, it’s crucial to clearly define these requirements. This ensures everyone understands the goals and expectations. Usually, these requirements are managed in a central repository that’s easily accessible to everyone on the team. JIRA is such a tool from Atlassian. However, our SPE Lab design is system agnostic, and all it cares about is that the performance requirements are clearly articulated and have the following details:

- Scenario or user workflows.
- Expected user load or concurrency.
- Data pre-seeded into the application.
- Expected response times and throughput.
- N/W considerations.
- Geolocation considerations.
- Test duration.
- Failure rate thresholds.

## 2.3. Performance Testing Tool

Once the performance requirements have been baselined, a performance engineer or a developer should start translating them into source code by writing executa-

ble test scripts. There are several tools, both open-source and licensed, available in the market that help you with this step. JMeter is an extremely popular open-source tool that helps developers and engineers create performance scripts using its record & playback feature along with its standalone API call feature. Gatling is another popular tool amongst the developer community that lets you create your performance test scripts using popular programming languages like Java and Scala. On the other hand, there are licensed tools like LoadRunner, Blaze Meter, and K6. Although they come at a price, they do provide additional capabilities like:

- Provisioning test infrastructure that your test needs on the fly.
- Generating load from different geolocations.
- Simulating load under various network bandwidths like 3G, 4G, 5G, and various WIFI standards.
- Rich reporting capabilities with the ability to retain and compare results encompassing multiple software releases.
- Out-of-the-box integration with APM tools like Datadog and Dynatrace, etc.
- AI for faster Root Cause Analysis (RCA) to find bottlenecks in the system.

I always recommend using a licensed tool because:

- Of their out of box capabilities.
- Managing an in-house load-testing infrastructure is expensive, time-consuming, and cumbersome.

#### **2.4. Software Configuration Management & Test Orchestrator**

The next set of tools that you need in your SPE Lab ecosystem is a source code management system like GitHub, GitLab, or Bitbucket and a test orchestrator like Jenkins, Circle CI, or Azure DevOps. A test orchestrator is a tool that downloads source code containing performance tests from the source control systems onto an engine or a virtual machine and then executes them. Although an orchestrator is an optional tool because most of the licensed tools have built-in execution engines, if you're relying on open-source tools like Gatling, it is always advisable to have an orchestrator. An orchestrator also facilitates continuous performance testing by plugging your performance tests into existing CI/CD pipeline/s, thus creating a fully automated/autonomous system.

My recommendation here would be to go for tools that are being used in your organization, as all we need is a means to manage your test code and plug them into existing CI/CD pipelines.

#### **2.5. APM (Application Performance Management)**

An Application Performance Monitoring (APM) tool is essential for performance testing, offering real-time insights into an application's performance across its entire stack. It helps teams identify bottlenecks, diagnose the root causes of performance issues, and resolve them proactively, ensuring optimized application performance and a seamless user experience, especially during load testing.

APM is essentially the brain of the SPE lab. APM tools offer deep-dive capabil-

ities to drill down into performance issues, allowing teams to identify the root cause of problems by analyzing logs, traces, and other related data, facilitating faster troubleshooting. Datadog and Dynatrace are two such tools that are leaders in this space and are must-have tools for the SPE lab.

## 2.6. Fire Extinguishers

A science lab is incomplete without a fire extinguisher, and we very well know what could happen when a science experiment goes wrong, as it could bring the entire lab down. In the same way, a performance engineering experiment gone wrong can destabilize the entire lab, making it difficult to ask the engineering teams to analyze what caused the break and what it would take to fix those issues and bring the lab back up. However, engineering teams can come up with in-house tools or scripts that can be run to restore the lab back to its pristine form. Some examples of such tools/scripts are:

- Database backup/restoration scripts.
- Purging cache.
- Purging messages from data streams like Kafka and ActiveMQ.

## 2.7. Communication & Collaboration

It is desirable to have a centralized hub where high-level test results can be shared and communicated to the desired teams so that they can collaborate on the results. That's where Slack and Microsoft Teams come in with their out-of-the-box integration capabilities, using their APIs to post results in a group or channel dedicated to the cause. Threaded conversations keep discussions organized and context intact.

## 3. Design & E2E Workflow

Now that we've figured out all the building blocks and tools required to build the SPE Lab let's connect the dots.

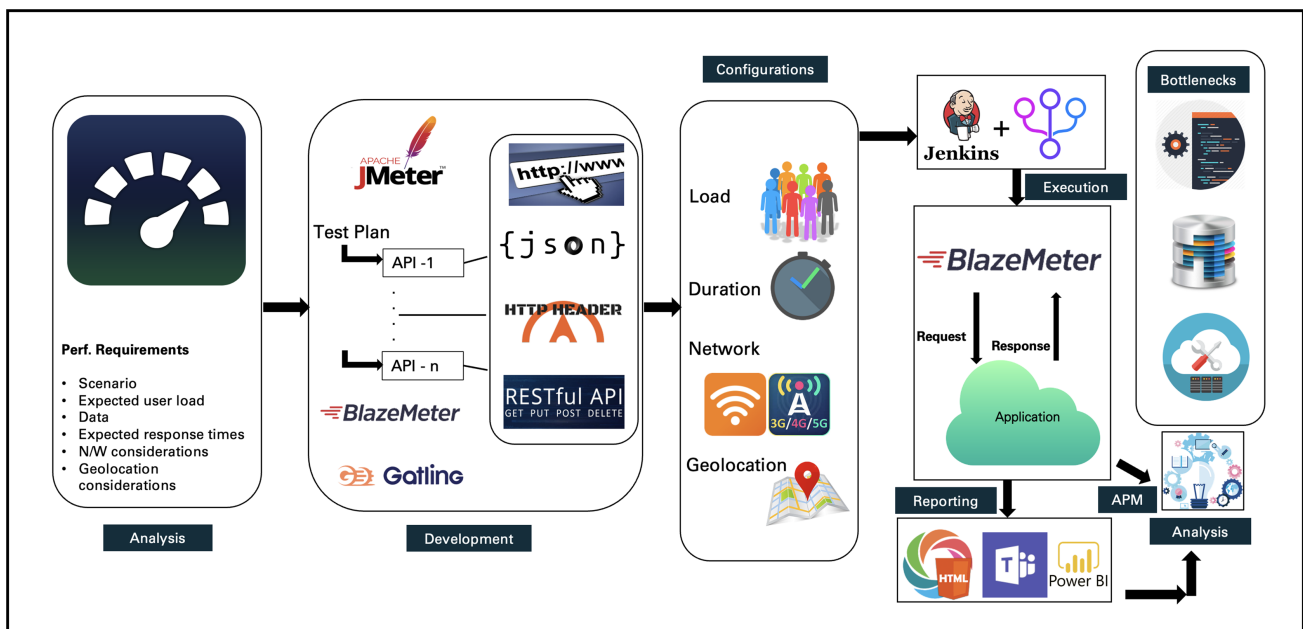
- Just like any software development process, developing a performance test starts with a requirement that must be clearly documented in a system like JIRA as explained in section 2.2 and, along with that, outlining goals of the performance experiments such as:
  - Response times (average, 90th, and 95th percentiles).
  - Throughput (Transactions per sec/min/hr.).
  - Error % thresholds.
  - Users ramp up schedule.
  - H/W (Hardware) resources like CPU and Memory consumption thresholds.
- An engineer then studies the requirement and starts developing performance scripts using a tool provided in the SPE lab, as described in section 2.3. After a few iterations of debugging and finetuning, the scripts are mostly ready. It is always a good idea to manage the source code of the script in a SCM tool, as discussed in detail in section 2.4.
- After baselining the scripts, we need to prepare our performance test envi-

environment hosting the application to be load/performance tested. Such an environment in the cloud looks like the one explained in section 2.1.

- Once the test environment is ready, the engineer needs to configure the performance scripts as per the requirements discussed in section 2.2 and then execute the script using tools provided in the lab, as explained in section 2.3.
- As the scripts are executed, the engineer can see real-time data and analytics in an APM tool, which is an integral part of the lab, as explained in section 2.5.
- After all the actions once results are generated, the team can start analyzing results, application logs in APM, and collaborate using a common communication channel as described in section 2.6 until they conclude whether the results are as expected or there are issues resulting in bottlenecks that are degrading the performance of the system.

The point to be noted here is we just accomplished our objective of conducting end-to-end performance testing & engineering in a SPE Lab by utilizing the tools and processes that it provided to its users.

As I mentioned earlier, the aim of this paper is to help create a tool/technology-agnostic SPE. However, if I were to create a pictorial representation of this SPE Lab using some of the tools discussed in sections 2 and 3, then this would look like the following (**Figure 2**):



**Figure 2.** Example of a SPE lab ecosystem.

#### 4. KPIs & Performance Metrics

Performance metrics and Key Performance Indicators (KPIs) are critical for the success of cloud applications. They provide actionable insights to ensure that applications perform reliably, meet user expectations, and align with business objectives. Examples of Key Cloud Application Metrics:

- System Metrics: CPU, memory, disk usage, and network bandwidth.
- Application Metrics: Latency, error rate, and transaction processing time.
- User Metrics: Active users, session duration, and bounce rate.

By systematically monitoring and analyzing these metrics and KPIs, organizations can ensure the scalability, reliability, and efficiency of their cloud applications while aligning with strategic objectives.

Continuing with the example of Microsoft Azure as the cloud services provider for an application deployed on a Kubernetes setup, some of the real-life metrics and industry benchmarks are given below (**Table 1**). The performance lab components, such as APM and the performance testing tools, use out-of-the-box features to provide these metrics [6].

**Table 1.** Performance benchmarks.

Component	Metrics & KPIs		
	Metric	Description	Benchmark
Kafka/AMQ	Messages queued/de-queued	The number of messages that are currently in the queue and how quickly they are being processed	Lag should not build up. All messages should get processed after the tests are complete
Kubernetes	CPU utilization	CPU utilized by individual nodes/pods	Must be within the limits specified in the resource definition of individual nodes/pods
Kubernetes	Memory	Memory utilized by individual nodes/pods	Must be within the limits specified in the resource definition of individual nodes/pods
Throughput	Transactions per second (TPS) or Transactions per minute (TPM)	The number of requests handled successfully by the service per unit of time	At least 1.5 times x requests/min x-value in production
Database	DTU %	A unit of measure representing a blended measure of CPU, memory, reads, and writes.	Average must be <= 75%
Cache	Server load	The Server Load metric represents the load on the Redis Server alone	Must be <= 80%
Cache	Memory	All the data in Redis is stored in memory and is often used for caching web pages and reducing the load on the server	Must be <= 80% of the total allocated memory
Backend	Response time	Average or 95th Percentile of the observed response times are equal to or less than the specified time	Must be <= 3000 milliseconds
Frontend	Response time	Average or 95th Percentile of the observed response times are equal to or less than the specified time	Must be <= 5000 milliseconds
Failures	Error %	Total number of errors observed during the test	Must be <= 0.05%

Any deviation/s from these benchmarks during testing should be investigated immediately. APMs provide in-depth analysis and help determine bottlenecks in the system. For example, if the response time of a service is on the higher side, it could lead us to the root cause, which, for example, could be due to a SQL query either taking more time to execute and return data or could be using more resources (DTUs) thus slowing down the system's performance.

## 5. Best Practices

To ensure accurate, reliable results and high-quality performance testing within the lab environment, it is very important that we follow some best practices.

### 5.1. Do's

- Always be mindful that you are not the only person using this performance tool and the environment.
- Have clear test requirements before jumping in. To have clear test case requirements, you should ensure they are specific, measurable, achievable, relevant, and time-bound (SMART).
- Check if the test data is already present in the tenant (if your test requires a certain amount of volume). Many of the tenants have created some basic test data.
- Check if there are preexisting test scripts already existing in the test suite. If not, you can create a similar one using the existing one as a sample.
- Use appropriate restoration jobs to bring back the databases (including Cosmos and SQL) to how it was before you started testing.
- Periodically upgrade all the tools that make the PE ecosystem and make sure they are all compatible with each other.
- Grant access and permissions as needed and revoke them when the job is done to avoid misuse of the lab.
- Implement alert and notification mechanism/s for continuous monitoring that could help to debug failures in the PE Lab.
- Make sure you've easy-to-use jobs/tools for bringing the PE Lab environment back to its pristine state either after the tests are complete or when there's a failure or breakdown of services in the PE Lab.

### 5.2. Don'ts

- Do not run tests if the PE Lab environment is unstable.
- Do not run the restoration jobs while there are ongoing tests running in the environment. While using jobs to clear AMQ and Kafka, remember that you are not the only person using this job. Running this job clears the data for all the users in the environment.
- Do not corrupt the tenants with incorrect, unusable data.

## 6. Discussion

The establishment of a Software Performance Engineering (SPE) Lab is vital for

ensuring that applications fulfill performance standards across diverse scenarios. A strategically designed lab allows teams to proactively detect bottlenecks, enhance resource efficiency, and uphold software dependability.

Historically, performance engineering has been a reactive endeavor; however, contemporary methodologies advocate for its early incorporation into the software development lifecycle. The integration of cloud-based platforms, automation tools, and AI-enhanced analytics has markedly increased the efficiency of performance testing. Nonetheless, challenges such as infrastructure expenses, tool compatibility, and the intricacies of workload characterization continue to exist.

Effective collaboration among development, QA, and operations teams is crucial to fully leverage the advantages of a performance lab. By cultivating a performance-oriented culture, organizations can improve software quality, mitigate deployment risks, and elevate the overall user experience.

## 7. Conclusions

Establishing a Software Performance Engineering Lab necessitates thorough attention to hardware, software, and methodologies. This paper has highlighted essential elements for creating a performance lab, covering everything from requirement definition to the implementation of monitoring techniques.

An effectively organized performance lab guarantees that applications can scale effectively, manage peak demands, and provide a smooth user experience. The incorporation of automated performance testing, ongoing monitoring, and predictive analytics will continue to enhance software performance engineering.

Future developments in AI, cloud technology, and serverless architectures will further influence the transformation of performance engineering labs, making them increasingly adaptive and intelligent. Organizations that view performance engineering as a fundamental discipline will secure a competitive edge in delivering superior software solutions.

## Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

## References

- [1] APACHE JMeter (2023) User's Manual. <https://jmeter.apache.org/usermanual/index.html>
- [2] Microsoft Azure Documentation (2023) Performance Testing in the Cloud. <https://docs.microsoft.com/en-us/azure/>
- [3] Liu, H.H. (2008) Software Performance and Scalability: A Quantitative Approach. IEEE Computer Society.
- [4] Microsoft Learn (2023) Advanced Azure Kubernetes Service (AKS) Microservices Architecture.

<https://learn.microsoft.com/en-us/azure/architecture/reference-architectures/containers/aks-microservices/aks-microservices-advanced>

- [5] Google (2023) Site Reliability Engineering. <https://sre.google/>
- [6] Bondi, A.B. (2014) Foundations of Software and System Performance Engineering: Process, Performance Modeling, Requirements, Testing, Scalability, and Practice. Addison-Wesley.