

DO PERFORMANCE BENCHMARKS AND COMPARISONS MATTER?

A GUIDE TO ASSESSING APPLICATION SERVER PERFORMANCE RESULTS

OVERVIEW

Application server performance benchmarks and comparison studies are used throughout the technology industry to promote one vendor's application server over the other's. While a standard, industry-approved benchmark exists (SPECjAppServer2004), it is largely outdated and no longer representative of both the technologies and profiles used today in building and maintaining modern business applications.

Many vendors still produce results for this outdated benchmark for the sole purpose of marketing their application servers. Unfortunately, due to the nature of the benchmark's age, results often misrepresent how a more modern application would perform on that vendor's product. As a result, some vendors have initiated their own performance comparison studies, often commissioning a third party to perform a vendor-defined test suite.

Given these dynamics, how should end users review application server performance comparisons? This guide provides a frame of reference for assessing and interpreting performance results from the primary Java application server market players: IBM, Oracle, and Red Hat.

Good performance tests and comparisons are typically those that can be easily repeated by end users in their own environments. Whether you choose to repeat the tests or not, Red Hat recommends that you seek out all of the appropriate information (configuration, settings, etc.) that would enable you to do so. This level of transparency allows you to more accurately assess application server performance, taking into account how your specific environment, setup, and configurations may differ, as any of those can have an impact on any performance test.

In order to provide perspective on performance results from any vendor, Red Hat recommends evaluating application server performance studies across the following areas:

PERFORMANCE PRIORITIES

Every organization wants high-performing applications. Superior performance often comes from throughout your application stack, not just the application server. In many cases, the way the application is designed, as well as how it connects to databases and other software components, can have a large impact on overall application performance. From a tuning perspective, many organizations actually spend more time tuning their custom-built applications and databases than their underlying application servers. While performance studies and comparisons may still provide helpful information, depending on your application's profile and design, a superiorly performing application server may have only a minor impact on the overall performance of your applications. To help assess your performance priorities and determine the impact a high-performing application server would have, it may be helpful to keep the following in mind:

- Where does your organization spend most of its performance tuning time: the application, application server, Java virtual machine, database, or another component?
- When performance issues are identified, which of these components is typically fixed to remove the bottleneck?

TECHNOLOGY STACK AND TEST SETUP

Performance tests are frequently run on the latest software and hardware with the latest available microchips. These configurations may include performance features not yet available to the general market. For any comparison tests, hardware and software configurations should be as similar as possible to ensure fair comparisons. When reviewing a benchmark or performance comparison, consider the following questions

- What software versions were used?
- What hardware configurations were used?
- How does the test setup compare to my hardware and software configuration?



TEST ENVIRONMENT

Applications interact with numerous other components in the technology stack: web servers, databases, operating systems, etc. The overall test environment can have a significant impact on performance results. The type of network and related configuration may play a large part as well; results for performance tests across wide-area networks (WAN) can differ greatly from the same tests run on local area networks. Performance on local servers is typically higher when compared to tests run over distributed servers. Depending on the configuration, applications that share the same server as a web server or database may outperform applications that make calls to separate web or database servers. When assessing performance reports, answers to the following questions will help verify the test environment and provide perspective as to what parameters may have influenced test results:

- What was the physical and virtual server setup across each infrastructure component?
- Were the networks local or distributed across the WAN?
- How similar was the test environment to the way my application(s) are deployed?

TEST APPLICATION PROFILE

The type and profile of the application is a critical element of any performance study. Since each business application typically has unique requirements, it's important to understand how an application used for a performance test compares to your application(s). Applications often utilize a variety of technologies and frameworks and, depending on the use, each technology can positively or negatively impact performance. The type of application, number of end users, peak workload, and the functionality provided by the application should all be considered as well. As an example, performance requirements for read-intensive applications may be very different than requirements for write-intensive applications. Unless a benchmark or comparison measures performance of an application whose profile closely matches one of your applications, the test

results may not provide a good comparison. Keep in mind the following questions to gain perspective on the performance results and how it relates to your application's requirements:

- How does the test application compare to my application(s)? Were the same technologies and frameworks used?
- Was the application profile and workload similar to my requirements?

INFRASTRUCTURE SETTINGS AND CONFIGURATIONS

Configurations and tuning play a large part in achieving superior application performance. In many tests, settings throughout each infrastructure component (test application, operating system, application server, Java virtual machine, web server, database server, etc.) and configurations such as clustering, security, and logging will have an impact on the overall performance result. When reviewing performance studies, it's important to understand the types of configurations and settings that were made as part of the test and whether those settings would be acceptable for your environment. In comparison studies, to ensure a fair assessment, it's important to note which optimizations were made to both application servers and their surrounding components. Use the following questions to determine what role infrastructure settings and configurations played in the performance test:

- What type of settings and configurations were used (network, application server, database, etc.)?
- What specific optimizations were conducted as part of the performance test?
- Can I replicate those same configurations in my environment?



STUDY SPONSORSHIP AND PRODUCT EXPERTISE

Since many studies and comparisons are used primarily for vendor marketing purposes, it is important to evaluate whether the study was commissioned or if it was truly an independent review. Independent studies—where vendor influence hasn't played a role in the results—should be viewed in higher regard than vendor-sponsored studies. Studies can appear to be independent, listing vendor commissioning information only on the last page. Vendor-commissioned studies may still provide valuable information, but readers should be cautioned that these funded studies are typically intended for competitive marketing purposes that may only benefit the sponsor. You should also verify the level of product expertise at the firm used to conduct the study to ensure that results can be easily replicated with your staff. When reviewing comparison studies, ensure the testing firm was well-versed and trained in both vendors' products to ensure a fair and level exercise was conducted. Use the following questions to ascertain the independence and expertise of the firm used for the performance study:

- Who conducted the study?
- Was the study truly independent or was it funded by a third party?
- If a third party was used, does a business relationship exist with one of the application server vendors?
- Did the firm have sufficient expertise and familiarity with both application server products to perform a fair assessment?

REVIEWING AND ASSESSING RESULTS

The price/performance ratio of the solution is more important than standalone application server performance. Price/performance evaluates a given product's performance while taking into consideration the total costs required to achieve that performance. For example, a higher-cost application server that delivers superior performance may be significantly more expensive to

acquire, deploy, and maintain than an application server that delivers more than acceptable performance at a more affordable price point. For studies that contain price/performance results, review product pricing information to ensure your deployments would achieve comparable pricing. For studies that do not include price/performance figures, conduct a cost analysis on the solution to ensure performance costs are properly put into perspective. Answers to the following questions will provide greater insight into the price/performance results of any study:

- What cost elements were included in the price/performance figure?
- Are the price/performance results in an acceptable range?
- How do the pricing parameters compare to my deployment?

ESTABLISHING APPLICATION SERVER PERFORMANCE ASSESSMENT CRITERIA

Best practice recommendations

Due to all of the variables that impact applications and application server performance, it becomes very difficult to assess how a performance study result will translate into real performance benefits for a specific environment. Given this, when evaluating application server performance, Red Hat recommends that you use the following guidelines to establish your own application server performance assessment criteria. These best practice recommendations should be used in conjunction with your own application server performance criteria:

- **Develop a performance impact perspective.** Numerous technology components (network, load balancers, databases, etc.) have an impact on application performance. Based on experience, you should establish a weighting system based on where the most performance tuning time is spent (e.g. 50% application, 20% application server, 15% JVM, 5% database, etc.). These weights

will help you determine how much of an impact application server performance will have in your environment. For example, if the majority of your time is spent tuning specific application code or JVM settings, then a superiorly performing application server may not help you resolve the majority of your performance issues. Establishing performance impact weights provides you with a perspective to more accurately assess the impact and criticality of application server performance in your environment.

- **Establish your own performance tests.** Develop and maintain performance test suites that match the types of applications and workloads typically deployed in your environment. There are numerous open source projects and commercially available products that can be leveraged as a basis for any performance test suite. Performance tests should be run in your own environment in order to account for the interaction of other dependent technologies as well as the impact of any environment-specific settings and configurations. Test suites should leverage peak workloads with sufficient headroom and be based on usage profiles and performance scenarios (clustering, failover etc.) to ensure performance under maximum load is exercised. By using your own performance tests, you'll be able to achieve an independent analysis of application server performance that is specific to your environment. In many instances, as a way to ensure you're not introducing any performance degradations into your production environment, these tests can be re-used as part of any application server release qualification and acceptance process.
- **Test with your own application(s).** Since performance is directly impacted by the usage profile and design of the application, you should use sample application(s) to verify application server performance. These test applications should closely match the profiles of applications deployed in your environment. They can be copies of existing applications that are already deployed in a production setting. Alternatively, if it matches typical application performance profiles, you can reuse the generic test applications used to validate new application server releases received from your existing vendor. If they aren't already, these sample applications should be augmented so that they can run on any standard Java application server.
- **Establish acceptable price/performance ratios.** The costs associated with achieving acceptable application server performance must be considered as part of any review. Comparing application servers on a price-for-performance perspective provides a much more realistic view than simply assessing performance by itself. You should consider both the capital and operating expenses associated with the software and hardware required to achieve performance goals. Measuring performance by a standard metric, such as Java operations per second (JOPS), provides a simple and consistent way to evaluate any Java application server. Once performance metrics and costs are gathered, determine the acceptable price/performance tolerance for your organization and use this ratio to more accurately assess one application server's performance over the other.

JBoss SALES AND INQUIRIES NORTH AMERICA

1-888-REDHAT1
www.jboss.com

