SERENA SOFTWARE

Performance Improvements in SBM

Authors: Kevin Ahn and Pete Dohner

2011-03-17

Table of Contents

Who Should Read This Paper?

This paper is intended for customers who are considering upgrading to the latest version of SBM. It discusses the improvements in scalability and performance between TeamTrack and SBM and from earlier versions of SBM.

In addition, comparison test results for new features that enhance performance are provided.

Note: The tests discussed in this document were performed using an older version of TeamTrack or SBM and version 2009 R2. Similar or greater performance results can be expected with version 2009 R3.

Related Documents

For additional information related to SBM performance and scalability, refer to these documents located on http://www.serena.com:

- SBM: Scaling for the Enterprise
- Enterprise Performance and Scalability Test Results

Why Upgrade to SBM?

Enterprise performance and scalability are areas of continual focus for SBM. Product development standards adhere to improvements in performance, scalability, and system resource usage. Rigorous testing in Serena's Performance Lab is conducted with each release cycle and demonstrates continued improvement.

Testing results in this paper focused on two primary areas:

- Comparison tests between older versions and SBM 2009 R2
- Feature tests that demonstrate how you can improve your system's performance by implementing full-text searching and seat-based licensing

Upgrade Test Scenarios

To accurately reflect the improvements in SBM, three sets of scalability and performance tests were performed using TeamTrack or an earlier version of SBM. Comparison tests were then performed using SBM 2009 R2. The same test configuration and dataset were used in each comparison test.

Scalability test results comparing the following versions and scenarios are discussed in this paper:

Older Product Version	Test Focus	Test Architecture	Test Dataset
TeamTrack Comparison Tests [page 11]	Common end-user tasks	TeamTrack-to-SBM 2009 R2 Test Architecture [page 6]	Mid-sized Dataset [page 10]
Comparison Tests for End-User Tasks [page 14]	Common end-user tasks	SBM 2009 R1-to-2009 R2 Test Architecture [page 7]	Large Enterprise Dataset [page 10]
Comparison Test for Orchestrations [page 18]	System service interactions using orchestrations	SBM 2009 R1-to-2009 R2 Test Architecture [page 7]	Large Enterprise Dataset [page 10]

Feature Test Scenarios

New features introduced in SBM can dramatically improve your system's performance and resource usage, as shown by the following load tests. All tests were performed using SBM 2009 R2.

Feature	Benefit	Test Architecture	Test Dataset
Legacy Keyword Searching vs. Full-Text Searching [page 20]	Optimizes keyword searching performance for Basic, Global, and Advanced Search features; keyword searching options in report specifications; and Value Find searching on forms for certain Single Relational and Multi-Relational fields.	SBM 2009 R1-to-2009 R2 Test Architecture [page 7]	SQL Server 2005 Enterprise Edition; Large Enterprise Dataset [page 10]
			Oracle 10g (10.2.0); Large Enterprise Dataset [page 10]
Concurrent Licenses vs. Seat Licenses [page 23]	Seat licenses reduce the license verification process. Tests show large improvements over systems using concurrent licenses.	SBM 2009 R1-to-2009 R2 Test Architecture [page 7]	Mid-sized Dataset [page 10]

Performance Improvements in SBM

Dramatic performance improvements were made to the core runtime components in SBM 2009 R2 and later releaes. These changes benefit small and large organizations, but particular emphasis has been placed on improving performance, enabling scalability, and reducing system resource usage for large enterprises.

Improvements fall into three categories: those that benefit end-users as they work with applications, those that ensure reliable and stable system service interactions, and core infrastructure changes that enable the system to run smoothly and with limited impact to your network resources.

Performance Improvements for Common End-User Tasks

- **Keyword Searching Improvements** – You can optimize keyword searching performance by enabling full-text searching for large primary and auxiliary tables. For details, refer to S137570 in the Serena Knowledge Base.
- **Project Scalability Improvements** – Several improvements will benefit systems that contain over 5,000 projects in individual applications. For example, users can browse lists of **Multi-View** reports more quickly and when they click the **Reports** tab in the Navigation pane or the **Show Me My Reports** link, the system filters out projects that have no reports defined. In addition, the system evaluates role privileges more efficiently, resulting in noticeable improvements for systems using roles with a large number of projects.
- **Favorites Folders** – Favorites folders now load more quickly in systems that have a large number of users. This improvement also applies to Folder fields used on Submit, Transition, and Update forms.
- **Secondary Owner Improvements** – Users will experience increased performance as they submit and update items that include a Secondary Owner field that is tied to a *Multi-User* or *Multi-Group* field. Reports that query a Secondary Owner field will also return results quicker than in previous releases.
- **Trend Reports** – Label generation based on a user's locale has been optimized to ensure faster response times.
- **Log On as Another User** – A **Find** button has been added to the **Log On as Another User** option on the **Advanced** tab of the User profile page. This enables you to search for a specific user rather than returning a full list of every user in your system.

Improvements for System Service Interactions

- **Guidance on Creating Scalable Orchestrations** – Serena experts have prepared a white paper explaining how to design your orchestrations for optimal performance and scalability. Contact customer support for a copy of this white paper.
- **Curbing Database Growth** – A new command-line tool prevents database growth by enabling you to purge intermediate SBM Orchestration Engine data from some database tables.
- **Throttling** – The internal throttling mechanism now handles asynchronous and synchronous orchestrations with separate queues and includes a configurable resource pool. This change allows for increased throughput and better use of system resources.
- **Improved Diagnostics** – Diagnostics now run at server startup and detect whether the SBM Orchestration Engine is able to execute orchestration workflows. Results are reported in the `sbm_oe_health.log` file, the Log Viewer in SBM Composer and the Common Log in Application Administrator
- **Apache Axis 2 Upgrade** – The Event Manager was upgraded to use Axis 2 1.4, which solved several performance issues.

Core Infrastructure Improvements

- **Oracle Improvements** – Fewer calls are made to the database when users submit, transition, or update items. In addition, response times have improved for applications built using the SBM API.
- **ReadLock Improvements** – The number of required ReadLocks has been greatly reduced. This change is most noticeable on systems running under high loads.
- **SBM User Workspace Template Optimization** – To reduce the amount of template processing time and improve performance when users view items and lists of items, many JavaScript functions have been removed from the view.htm and list.htm templates and added to cacheable JavaScript files.
- **Deployment and Promotion** – Significant improvements have been made to auxiliary table promotion. In addition, response times for the Process App view in Application Administrator have greatly improved.

Test Methodology

Testing was conducted in a private enterprise performance testing lab. Tests measured application response time, throughput, and system resource usage under varying load conditions and dataset configurations.

To narrow the testing gap, enterprise customer configurations and test datasets are modeled for testing purposes.

Test Application: SBM 2009 R2

Test Driver: HP LoadRunner 8.1.4

LoadRunner generates load by creating virtual users to model real human user activity. Each virtual user performs scripted tasks and sends crafted HTTP requests to the target application. LoadRunner then uses a threading execution model to create multiple instances of unique virtual users to create load and concurrency on the target application.

Runtime Test Architecture

Two sets of server configurations were used to perform the tests described in this document:

- TeamTrack-to-SBM 2009 R2 Test Architecture [page 6]
- SBM 2009 R1-to-2009 R2 Test Architecture [page 7]

TeamTrack-to-SBM 2009 R2 Test Architecture

A standard two-server environment comprised of a Web server and a database server was used for the TeamTrack 6.6.1 to SBM 2009 R2 upgrade comparison test, as shown in the following figure. This architecture was used in the TeamTrack Comparison Tests [page 11].

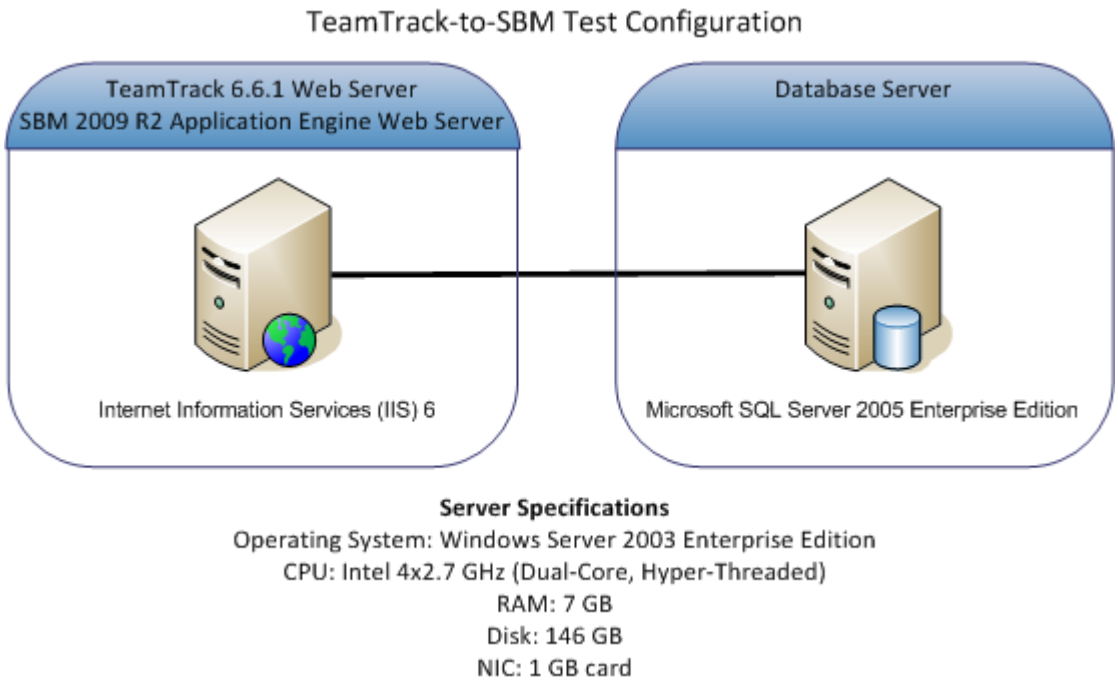

SBM 2009 R1-to-2009 R2 Test Architecture

The SBM 2009 R2 runtime test activities described in this document are executed on the SBM Server. There are two types of runtime activities:

- **Manual End-user Tasks** - Executed by the SBM Application Engine. Manual end-user tasks are performed in the SBM User Workspace using a standard Web browser. The SBM User Workspace is implemented using HTML, JavaScript, and Flash. The SBM Application Engine coordinates these "human" workflow, or "application," activities by executing as an in-process ISAPI plug-in to Internet Information Services (IIS). It receives processes and responds to HTTP and Web service requests from the Web server.
- **System Service Interactions** - Executed by the SBM Orchestration Engine. These system service interactions, or orchestrations, are kicked off by users as they perform workflow tasks. The SBM Orchestration Engine receives events from applications and executes associated orchestrations as BPEL processes. These BPEL processes can use Web services to execute a business process across multiple tools. The SBM Orchestration Engine is part of the Serena Common JBoss service, which is a J2EE container for various services provided in the SBM Server.

The following tests were performed using a three-server environment containing the SBM Application Engine Web server on one server, Serena Common JBoss and the Event Manager on a second server, and a third machine acting as a database server:

- Comparison Tests for End-User Tasks [page 14]
- Comparison Test for Orchestrations [page 18]
- Legacy Keyword Searching vs. Full-Text Searching [page 20]
- Concurrent Licenses vs. Seat Licenses [page 23]

The following figure illustrates the server configuration for 2009 R1 and 2009 R2.

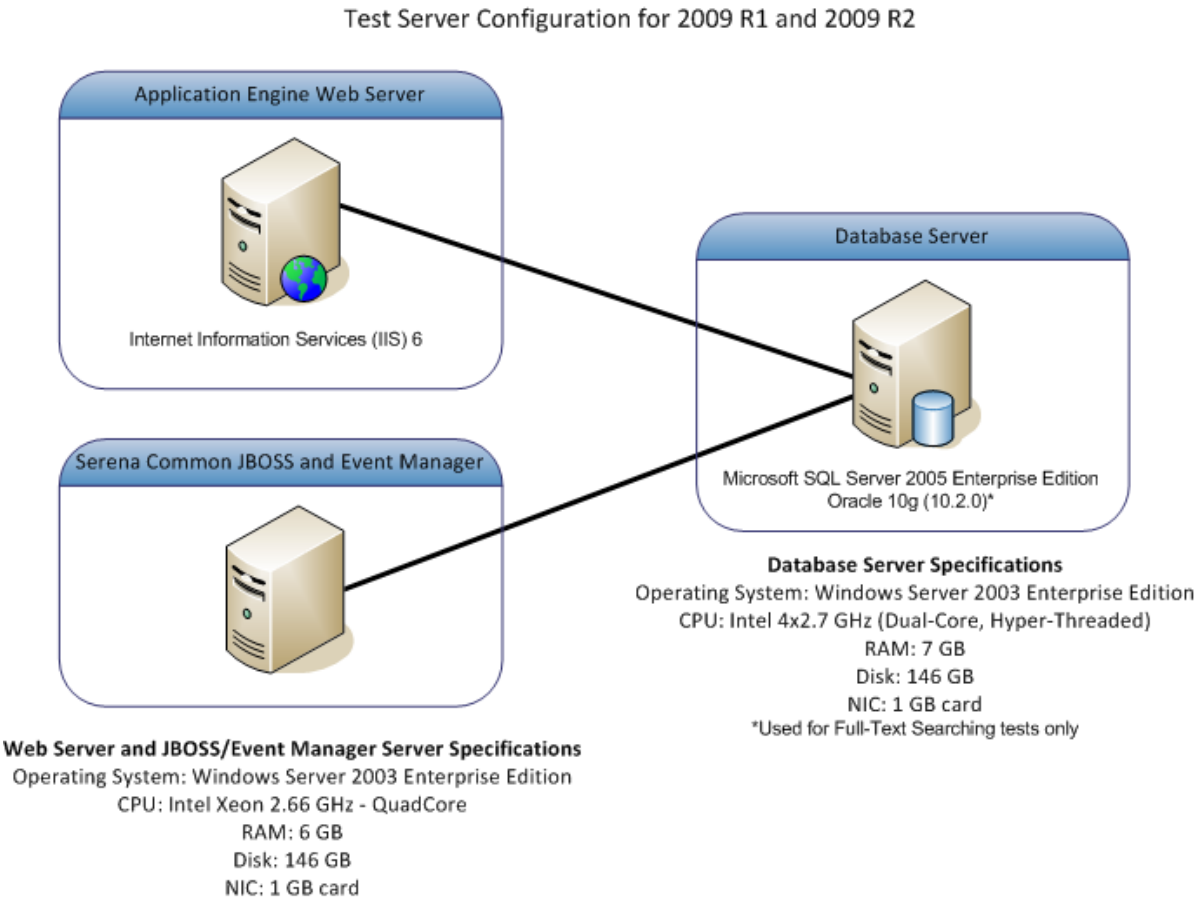

NOTE: The SQL Server database was configured to use the Read Committed Snapshot isolation level. For details, refer to Microsoft's configuration documentation.

TEST SCENARIOS

SBM Application Engine Test Scenarios

SBM Application Engine comparison tests simulate a scaling number of common, concurrent, repeated tasks completed by varying numbers of virtual users. Using a standard defect tracking application that enables a team to track product defects, the tests replicate the life cycle of defects from submittal, approval, and assignment. Notes and attachments are added to each item as it moves through the process. Keyword searches, ID searches, and various reports are also run against the system.

Scalability tests for end-user tasks were performed by a set number of unique virtual users over a 90-minute time period. “Think” time between virtual user clicks was 10 seconds. To test scalability, sets of 50, 100, and 200 unique virtual users log in to the SBM User Workspace once, perform a variety of tasks, and then log out at the end of the test.

Tasks include:

- Keyword search
- Submit item

- Assign item to an owner
- Reassign item
- ID search
- Add note to an item (2 kb in size)
- Add attachment to an item (12 kb in size)
- Run a Listing report
- Run a Multi-view report

SBM Orchestration Engine Test Scenarios

The SBM Orchestration Engine tests simulate a repeated, scaling number of synchronous and asynchronous events initiated by virtual users as they submit items into a “human” workflow application. These system service interactions are referred to as orchestrations.

Using a simple human process to initiate orchestrations, the tests replicate large, complex business processes that continually increase the amount and size of data processed by the SBM Orchestration Engine. Tests include repeated updates of items and automated searches using built-in Web services, representing complex Web service calls that mimic high-load business usage. This ensures that the system can concurrently process a small to large number of business transactions.

Tests were performed by a set number of unique virtual users over one hour. “Think” time between virtual user clicks was 10 seconds. To test scalability, sets of 20, 40, and 80 unique virtual users log in to the SBM User Workspace, submit an item into the system to initiate the orchestrations, perform additional tasks, and then log out of the system. Asynchronous and synchronous orchestration tests were run concurrently. For each scaling model, the total number of virtual users was split at 80 percent for the asynchronous calls and 20 percent for the synchronous calls.

Asynchronous Tasks

Virtual users:

- Login
- Submit Item
- ID Search
- Logout

Orchestrations (Initiated by virtual users submitting items):

- 7 Web service calls to update the item
- 2 Web service calls to perform a keyword search by title

Synchronous Tasks

Virtual users:

- Login
- Submit Item
- Logout

TEST DATASETS

Large Enterprise Dataset

The large enterprise dataset is shown in the table below. This dataset is used in the following tests:

- Comparison Tests for End-User Tasks [page 14]
- Comparison Test for Orchestrations [page 18]
- Legacy Keyword Searching vs. Full-Text Searching [page 20]

Application Item	Number of Items
Workflows (define the process items follow)	31
Projects (store process items, such as issues and incidents)	5,000
Users	21,000
Groups	138
Folders (store personal and group favorites)	147,570
Contacts (similar to address book entries)	20,053
Issues (process items that follow a workflow)	58,026

Mid-sized Dataset

The mid-sized enterprise dataset is shown in the table below. This dataset is used in the following tests:

- TeamTrack Comparison Tests [page 11]
- Concurrent Licenses vs. Seat Licenses [page 23]

Application Item	Number of Items
Workflows (define the process items follow)	31
Projects (store process items, such as issues and incidents)	235
Users	1,100
Groups	138
Folders (store personal and group favorites)	7,555
Contacts (similar to address book entries)	52
Issues (process items that follow a workflow)	57,984

TeamTrack Comparison Tests

Scalability test results between TeamTrack 6.6.1 and SBM 2009 R2 indicate significant improvements in transaction time and database server CPU usage. Notable improvements are also seen in Web server memory and CPU usage.

Comparison tests between TeamTrack 6.6.1 and SBM 2009 R2 show strong results in the following categories:

- Average Transaction Response Times [page 11]
- Database Server CPU Usage [page 12]
- Web Server CPU Usage [page 12]
- Web Server Memory Usage [page 13]

The following sections show side-by-side results for each category.

Average Transaction Response Times

The following chart reflects average transaction response time for load levels of 50, 100, and 200 virtual users. SBM demonstrates higher throughput and faster response times, with most transactions remaining below the sub-second level for all operations.

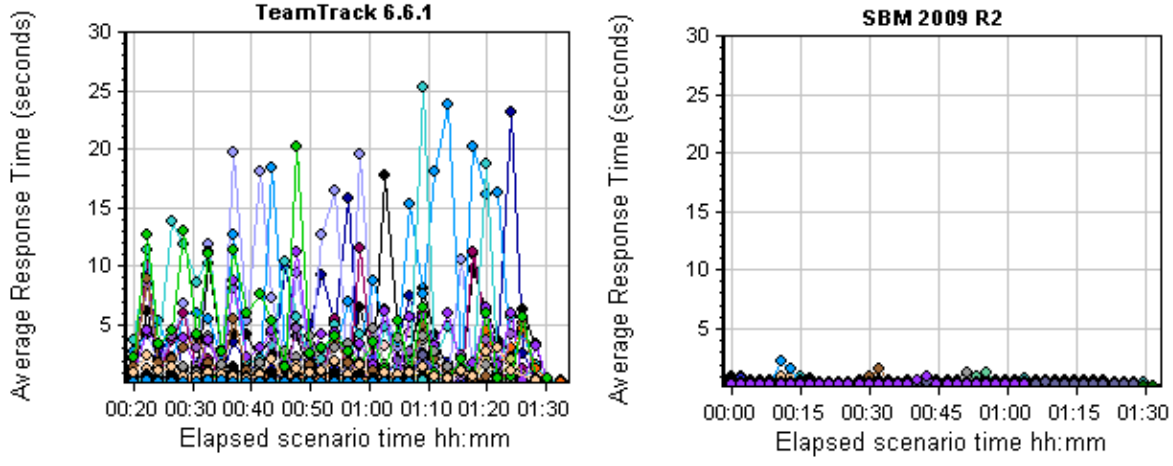

Database Server CPU Usage

Large improvements in database server CPU stability and resource consumption levels are seen with SBM, as reflected in the chart below. Note that the testing was performed using Microsoft SQL Server, but similar results can be expected for systems using Oracle.

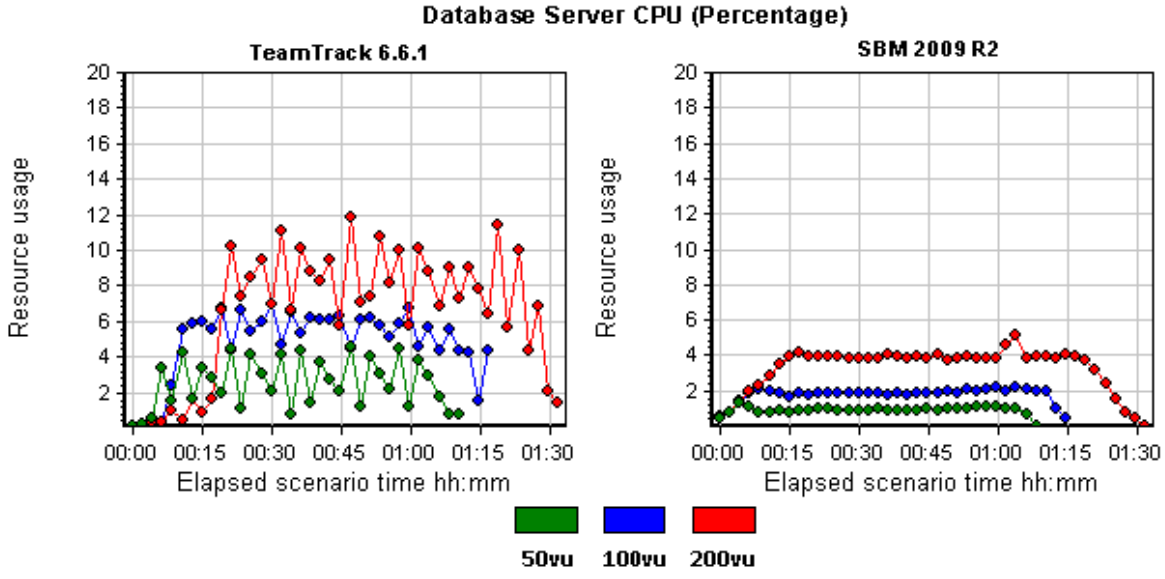

Web Server CPU Usage

The following chart reflects the Web server's CPU usage for all load levels. SBM demonstrates stable and improved levels of resource consumption as the workload increases.

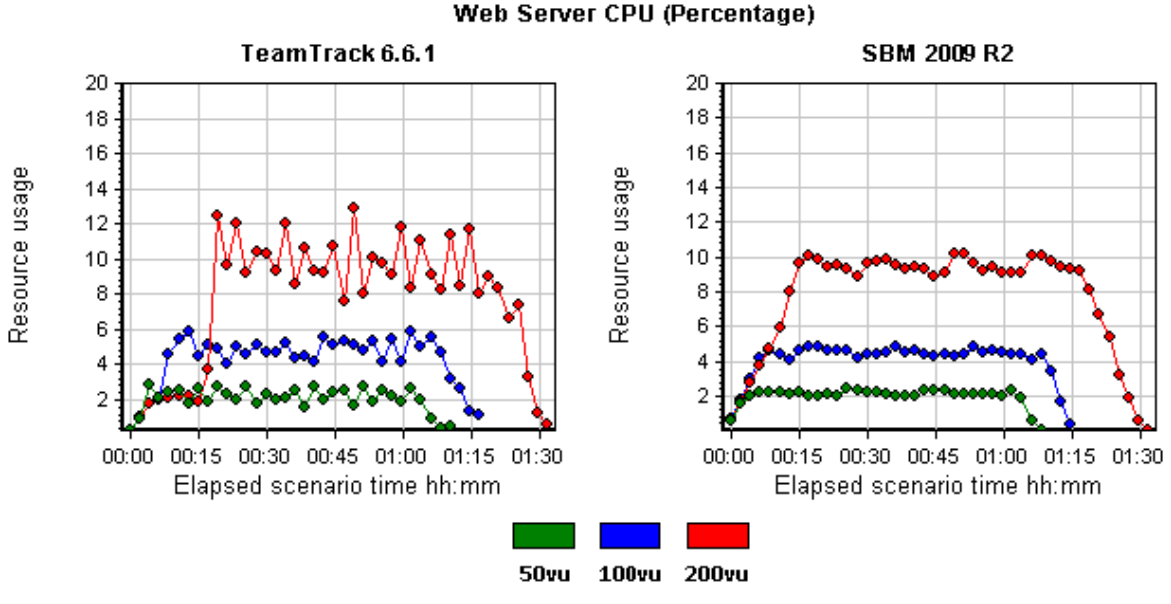

Web Server Memory Usage

The charts below reflect Web server physical and virtual memory usage for all load levels. As the workload increases, SBM consumes less physical memory on the Web server. Memory usage also remains stable throughout the test period.

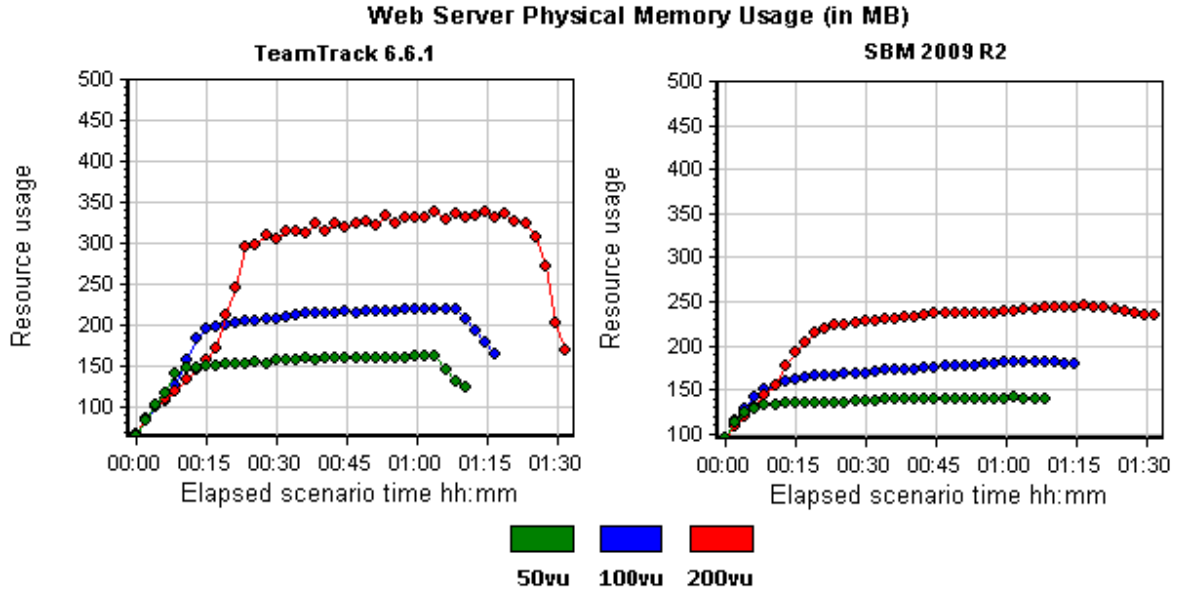

While initial virtual memory usage is greater with SBM, overall memory usage during the test period is lower than in TeamTrack. The increased memory usage can be attributed to native Unicode support in SBM. The graph below shows IIS virtual memory consumption under all three load levels. Though the initial footprint is larger, the runtime footprint remains measurably lower in SBM 2009 R2.

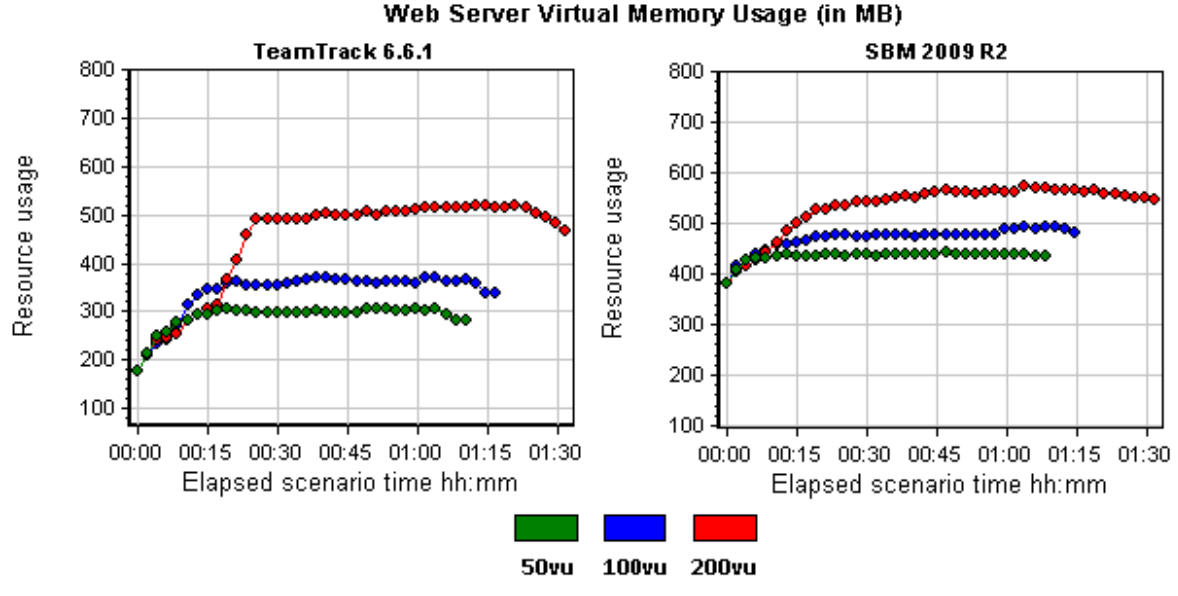

Comparison Tests for End-User Tasks

Based on a dataset with a large number of projects, users, and personal folders, scalability test results between SBM 2009 R1 and 2009 R2 indicate significant improvements to transaction time as user loads increase. Significant improvements are also noted for the Web server CPU usage, with the Web server memory usage showing no performance degradation from version 2009 R1. These gains in Web server performance resulted in slightly higher database server CPU usage.

Comparison tests results are provided for the following categories:

- Average Transaction Response Time [page 14]
- Web Server CPU Usage [page 16]
- Database Server CPU Usage [page 17]
- Database Server CPU Usage [page 17]

The following sections show side-by-side results for each category.

Average Transaction Response Time

At 50- and 100-virtual-user loads, average transaction response times are slightly improved in version 2009 R2. Dramatic improvements are seen at higher user loads, however, as indicated in the image depicting results from the 200-virtual-user test. This improvement is attributed to improvements in how Favorites folders are processed, as well as ReadLock changes that removed resource contention by reducing a Web server CPU bottleneck.

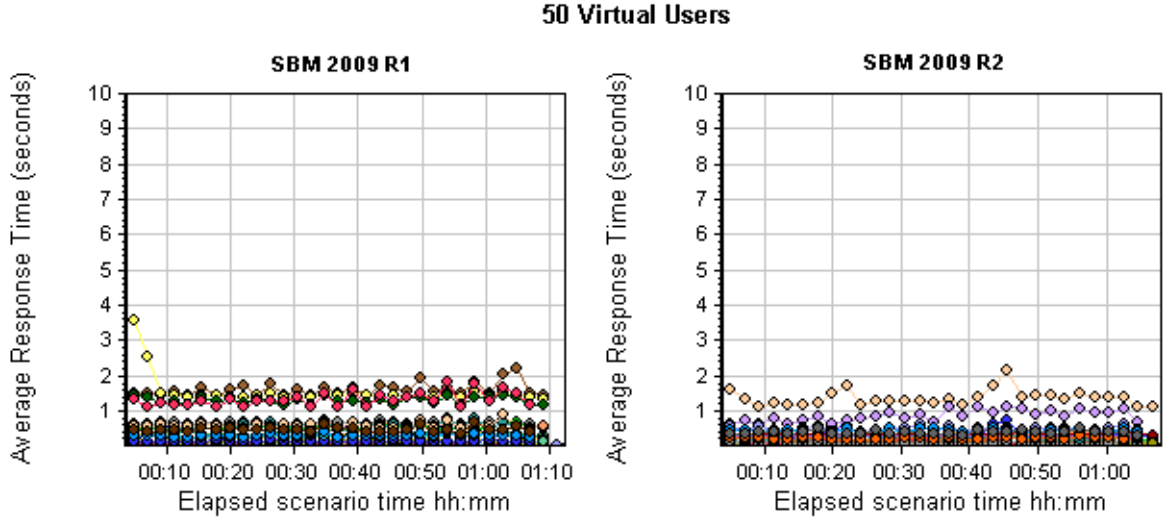

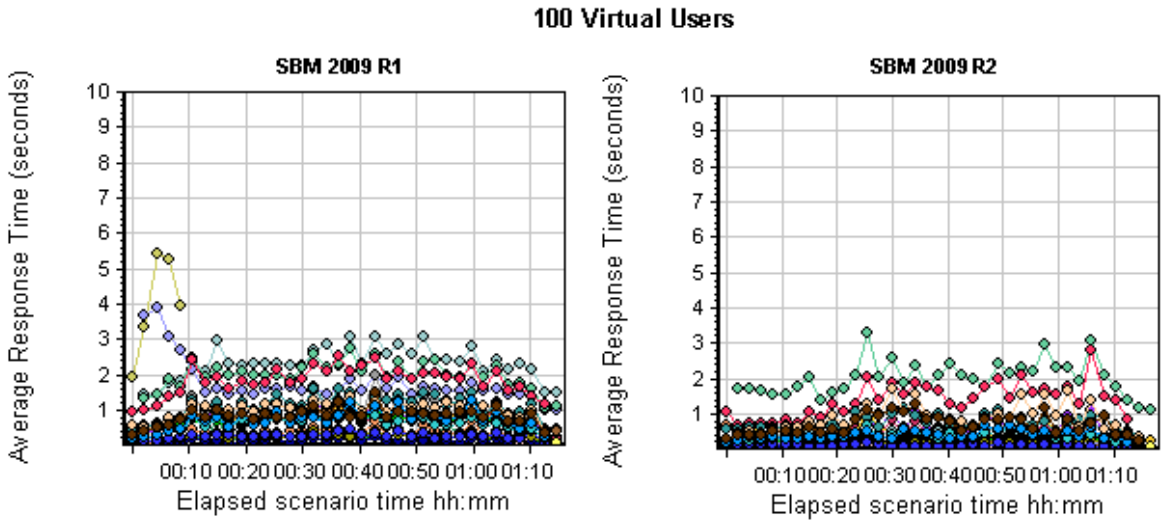
100 Virtual Users
SBM 2009 R1
Average Response Time (seconds)
10
9
8
7
6
5
4
3
2
1
00:10 00:20 00:30 00:40 00:50 01:00 01:10
Elapsed scenario time hh:mm
SBM 2009 R2
Average Response Time (seconds)
10
9
8
7
6
5
4
3
2
1
00:10 00:20 00:30 00:40 00:50 01:00 01:10
Elapsed scenario time hh:mm

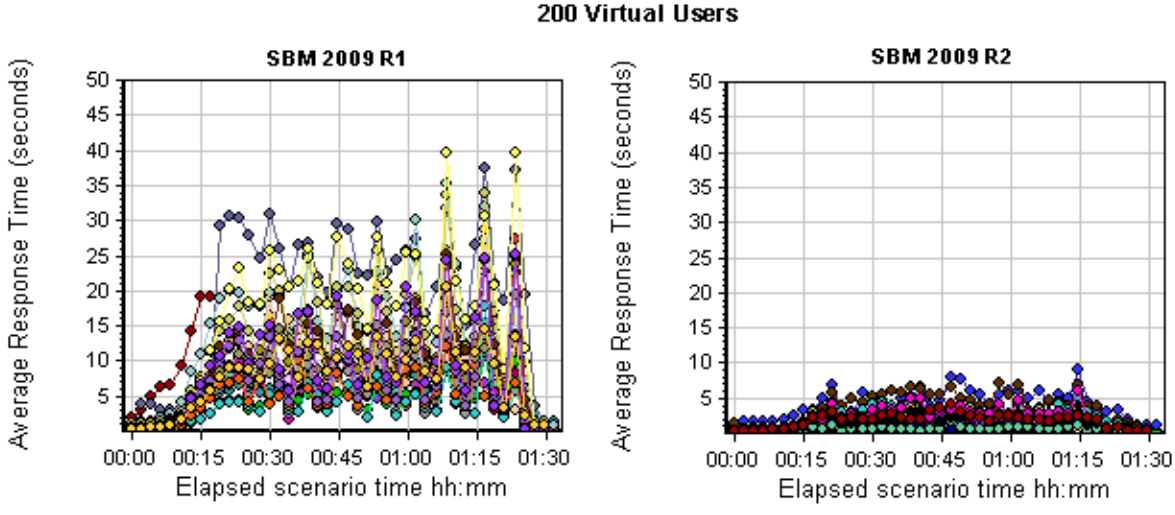
200 Virtual Users
SBM 2009 R1
Average Response Time (seconds)
50
45
40
35
30
25
20
15
10
5
00:00 00:15 00:30 00:45 01:00 01:15 01:30
Elapsed scenario time hh:mm
SBM 2009 R2
Average Response Time (seconds)
50
45
40
35
30
25
20
15
10
5
00:00 00:15 00:30 00:45 01:00 01:15 01:30
Elapsed scenario time hh:mm

Web Server CPU Usage

SBM 2009 R2 has been optimized to handle large dataset characteristics and reduce the workload on the Web server, as demonstrated in the following charts.

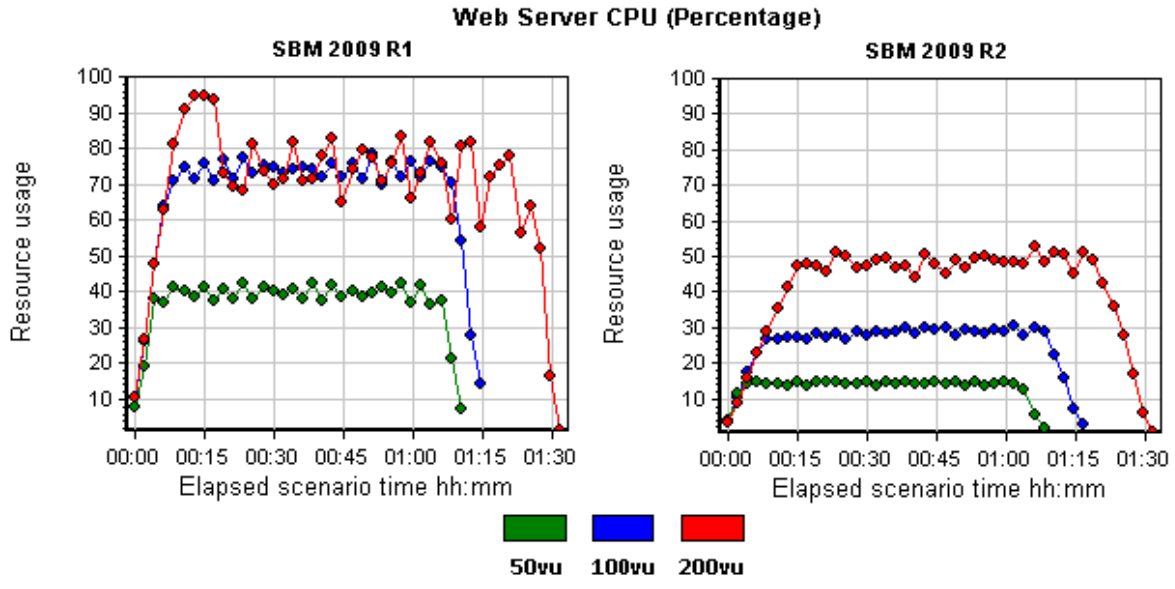

Web Server Memory Usage

Server caching optimization in version 2009 R2 reduced and stabilized the Web server memory footprint, as shown in the following charts.

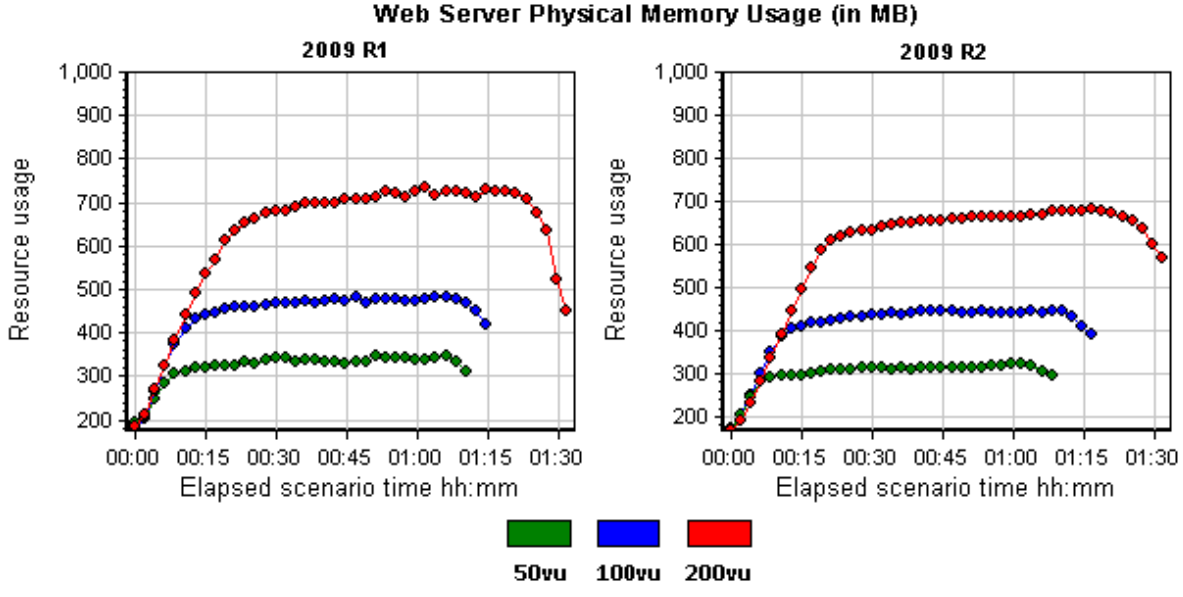

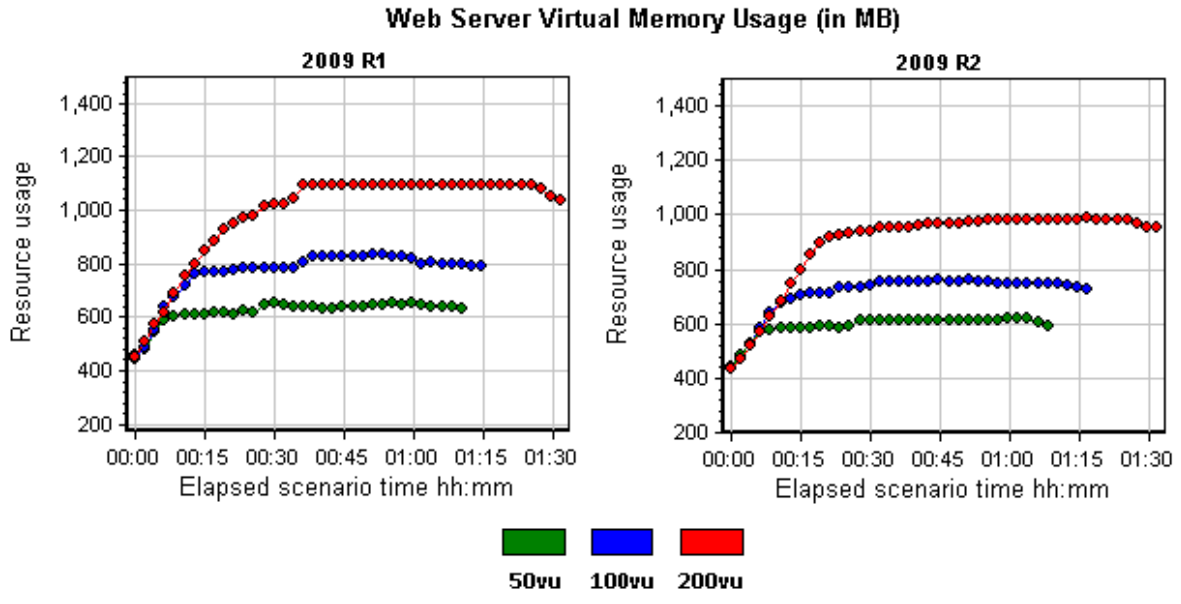

DATABASE SERVER CPU USAGE

Increased activity generated by the Web server caused the database server resource usage to increase between version 2009 R1 and 2009 R2. This change can be attributed to removing resource contention by reducing a Web server CPU bottleneck and by improving transaction response times and application throughput.

Note that the testing was performed using Microsoft SQL Server, but similar results can be expected for systems using Oracle.

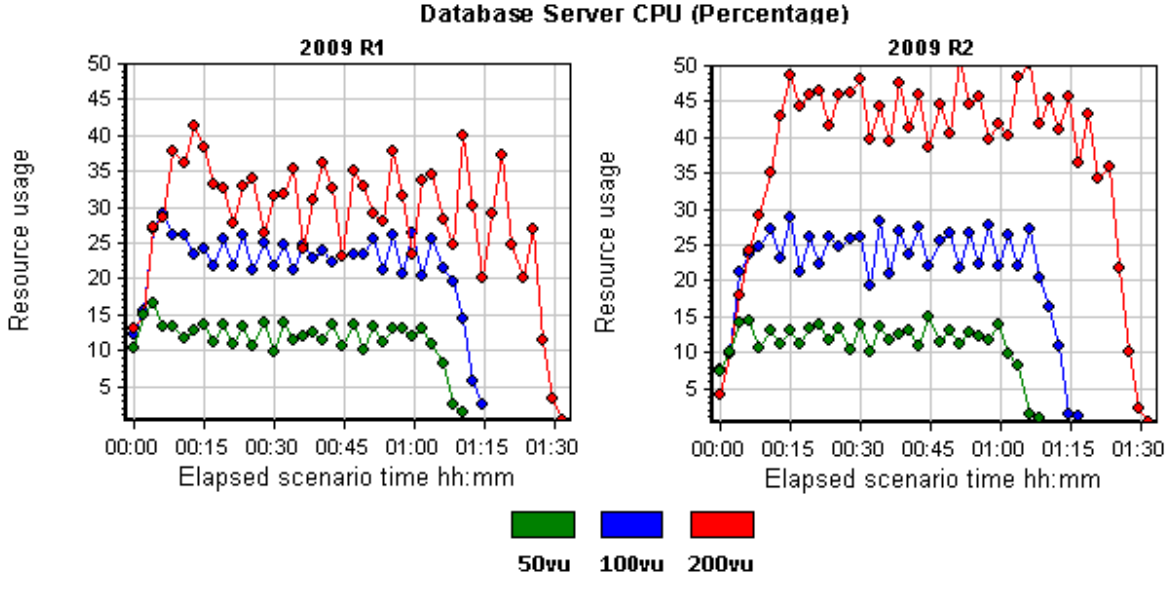

Comparison Test for Orchestrations

Comparison test results indicate improved or stable performance in the following categories:

Average Transaction Response Times

The following chart reflects average transaction response time for load levels of 20, 40, and 80 unique virtual users, demonstrating lower and more stable response times with version 2009 R2.

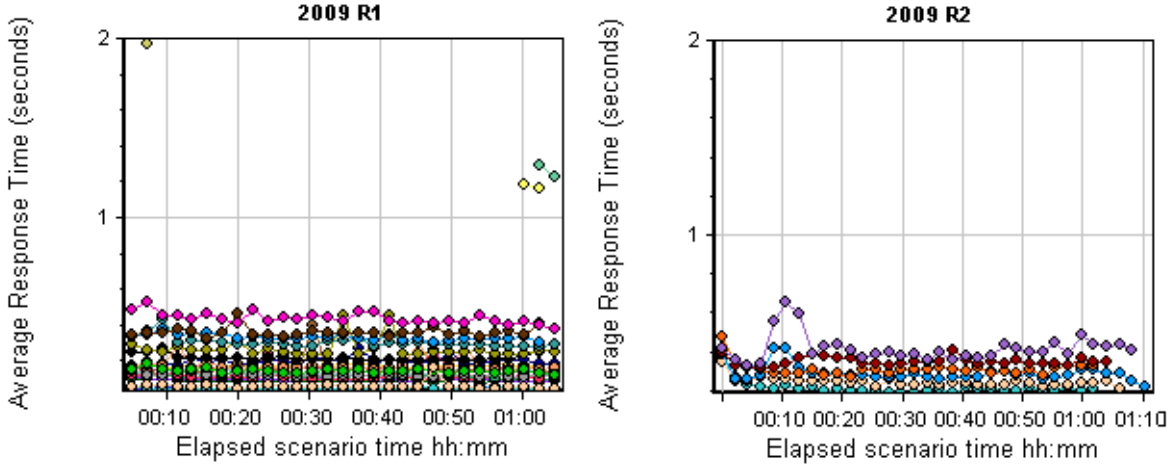

Transaction Throughput

The following chart reflects the number of synchronous and asynchronous transactions processed during the one-hour test by a scaling number of virtual users.

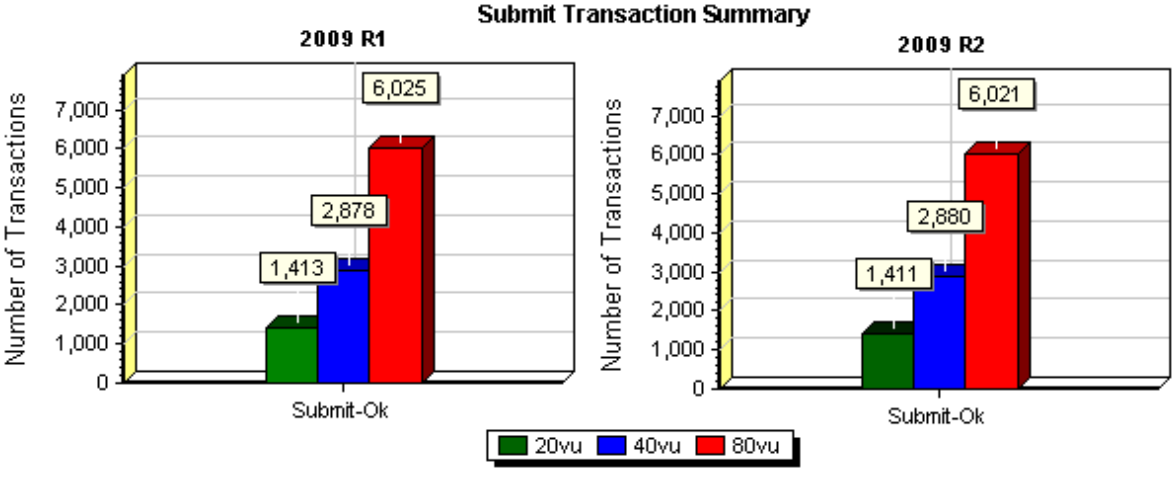

JBoss Server CPU Usage

The JBoss server manages system interaction activities and manages SBM Orchestration Engine execution. JBoss server CPU usage remains stable in version 2009 R2 using the default server configuration. You can further enhance server resource usage by implementing message queues and internal throttling.

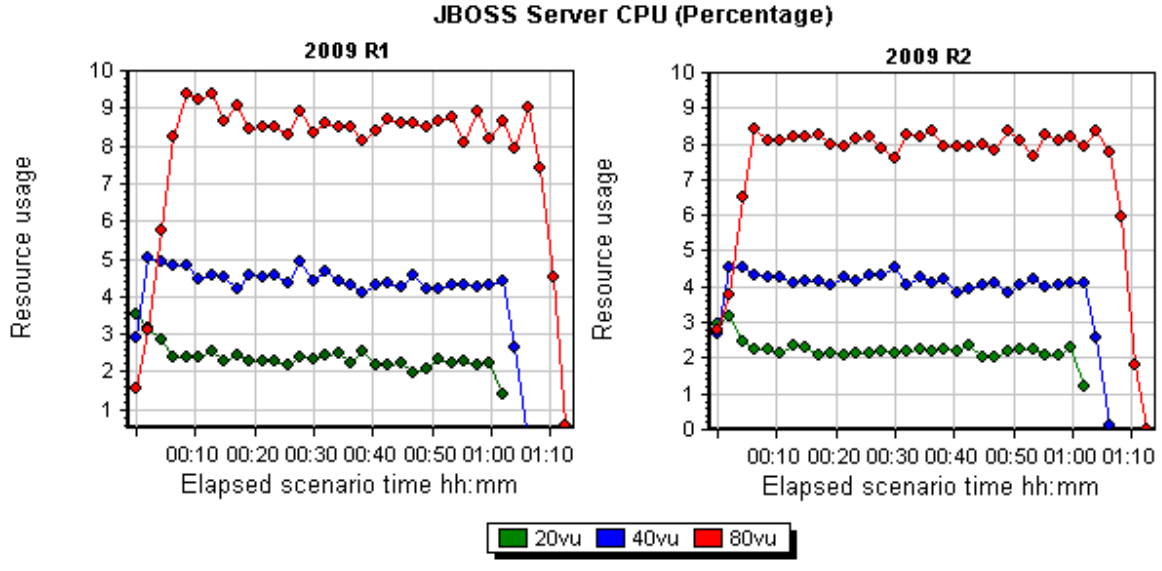

JBoss Server Memory Usage

JBoss server memory usage increased slightly from version 2009 R1 to version 2009 R2, as shown in the graphs below.

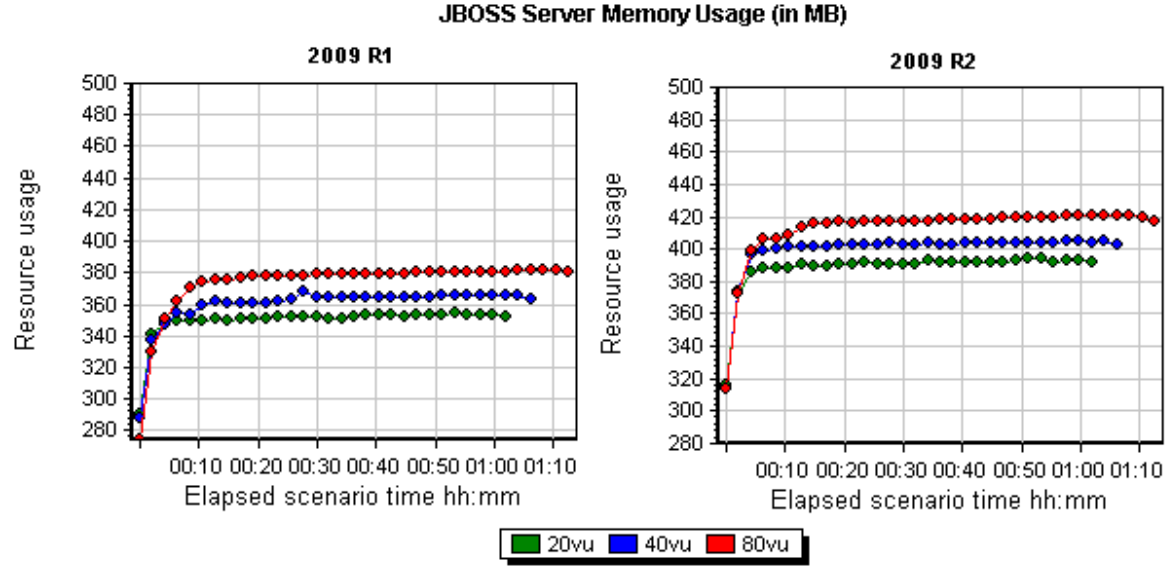

Legacy Keyword Searching vs. Full-Text Searching

Full-text searching capability was added to version 2009 R2, enabling you to leverage the highly scalable full-text search features in Microsoft SQL Server and Oracle. (In Oracle, full-text searching is implemented using Oracle Text.) Test results show that you can greatly optimize keyword searching performance by enabling full-text searching for heavily used primary and auxiliary tables that contain large amounts of data.

Full-text searching applies to the following features:

- Basic Search
- Advanced Search
- Global Search
- Keyword searching options in report search specifications
- Value Find searching on forms for *Single Relational* and *Multi-Relational* fields if the value display format for the relational field table contains a *Text* field

For implementation details, refer to S137570 in the Knowledgebase at serena.com.

Side-by-side comparison tests results are provided for the following categories:

- Microsoft SQL Server Improvements [page 20]
- Oracle Improvements [page 22]

Microsoft SQL Server Improvements

The following charts represent dramatic performance improvements in response times and database server CPU usages when full-text searching is used with Microsoft SQL Server. The charts on the left show performance results for legacy keyword searching when 200 unique virtual users perform Basic searches over the test period; the charts on the right show results for the same test after full-text searching is implemented.

Response Times

Significant improvements can be seen in response times for keyword searches after full-text searching is implemented. The majority of full-text searching response times remain under 1 second.

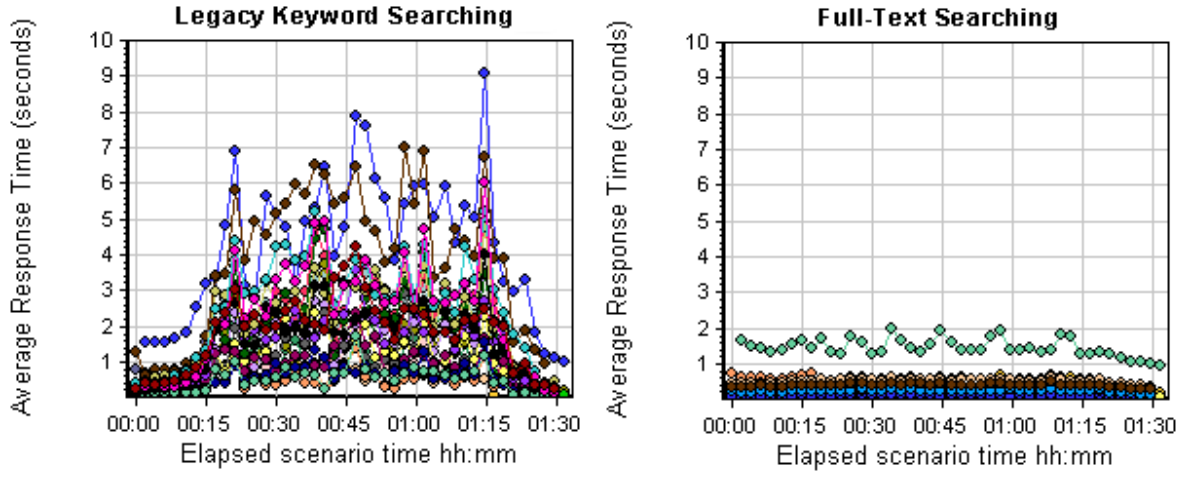

Database Server CPU Usage

Full-text searching greatly reduces the workload on the database server, as shown in the results below.

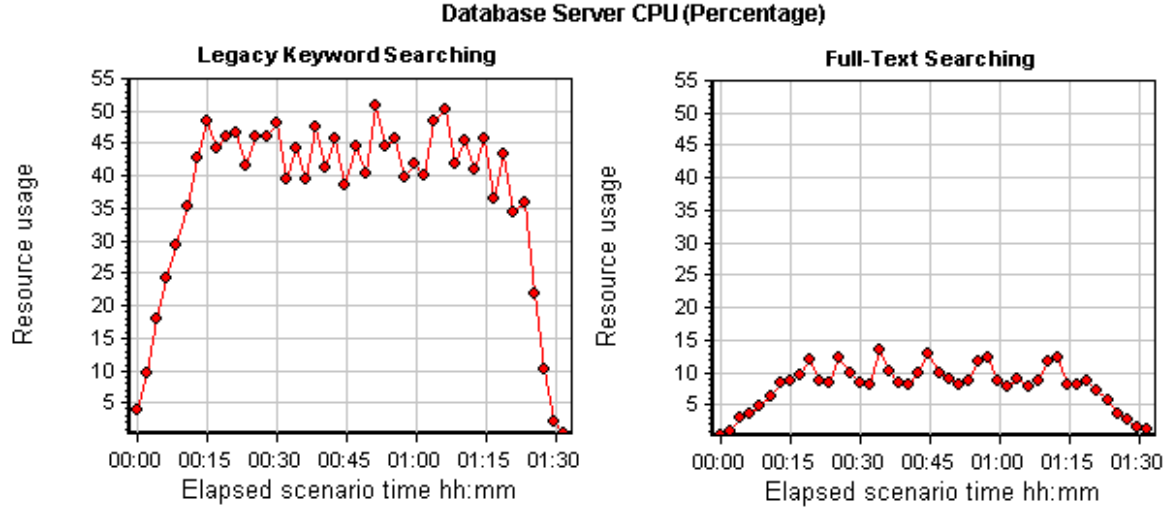

ORACLE IMPROVEMENTS

Oracle Text provides similar capabilities as SQL Server's full-text searching. The following charts reflect the significant performance improvements in response times and database server CPU usages when full-text searching is used with Oracle. The charts on the left show performance results for legacy keyword searching when 200 unique virtual users perform Basic searches over the test period; the charts on the right show results for the same test after full-text searching is implemented.

Response Times

Significant improvements can be seen in response times for keyword searches after full-text searching is implement. All transaction response times remain under 3 seconds, with the majority of response times under 1 second.

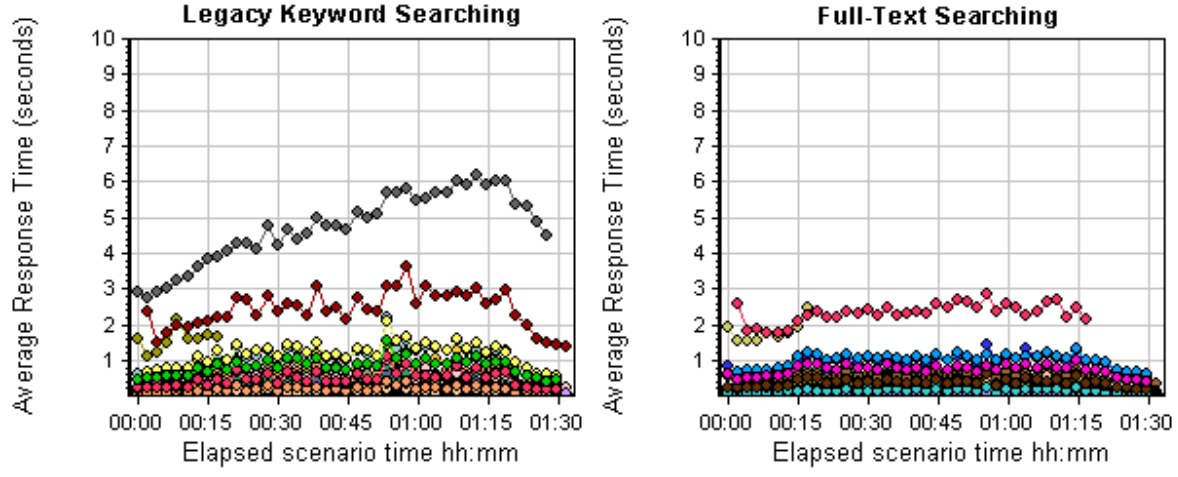

Resource Usage

Full-text searching greatly reduces the workload on the database server, allowing for faster and more efficient search queries.

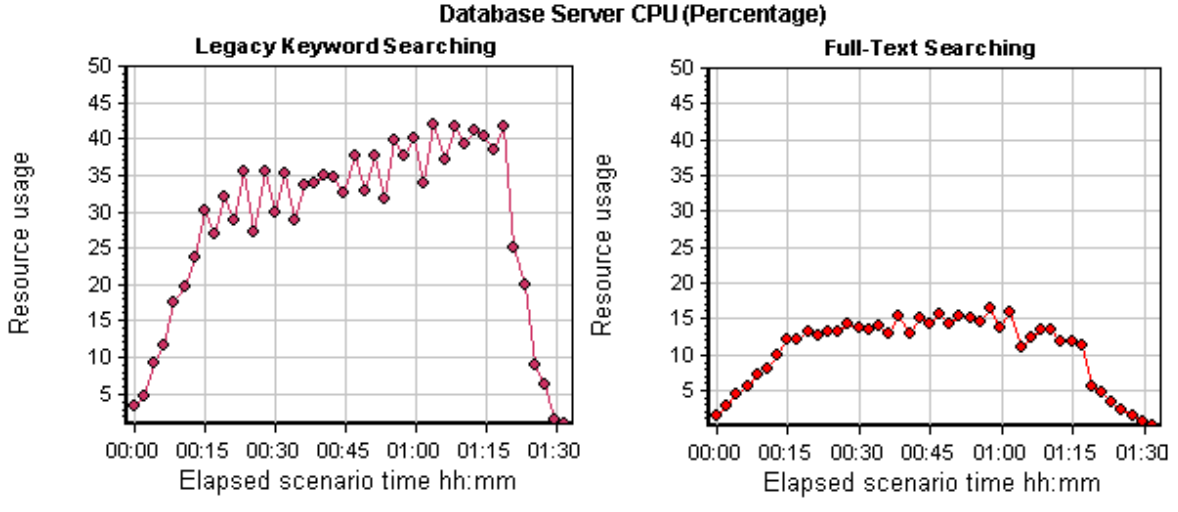

Concurrent Licenses vs. Seat Licenses

Beginning with SBM 2009 R1, seat licenses, which determine the number of users that can be active in an SBM system, are available for user and managed administrator product access types.

You can optimize your system's performance by converting to seat-based licenses because concurrent license checks are eliminated. The charts below show the average response times for systems using concurrent licenses and seat licenses. In each test, one virtual user logs in to the system every 5 seconds for a total of 50 virtual users over the test period. Each virtual user performs typical defect tracking tasks, such as submitting items and performing keyword searches. The charts reflect greatly improved performance with seat licenses, particularly during the first 5 minutes of each user's login period. This comparison test was performed using version 2009 R2.

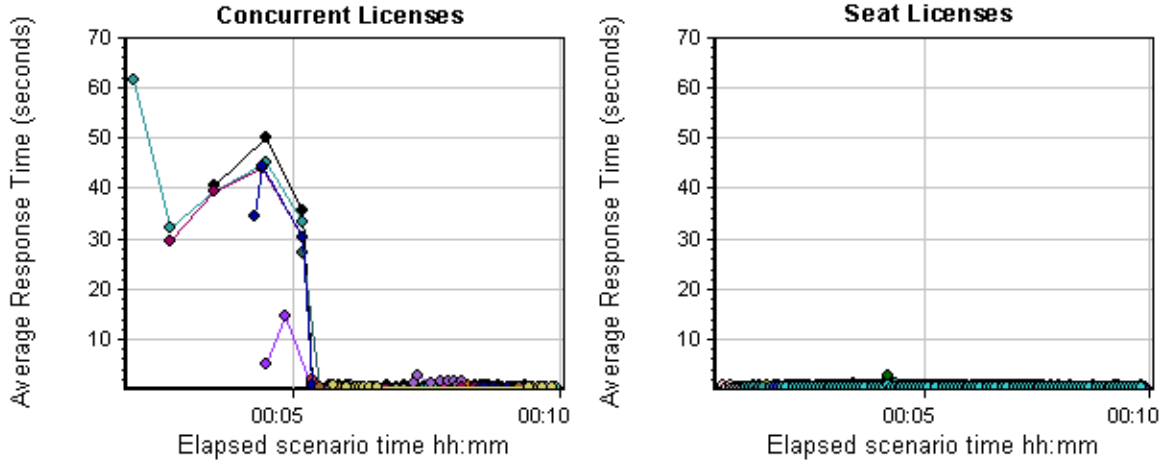

Contact your Serena sales representative for more information about seat-based licenses.

ABOUT

Serena Software, Inc., the Change Governance leader, helps more than 15,000 organizations around the world—including 96 of the Fortune 100 and 90 of the Global 100—turn change into a business advantage. Serena is headquartered in Redwood City, California, and has offices throughout the U.S., Europe, and Asia Pacific.

CONTACT

Website: http://www.serena.com