Lawson M3 7.1 Large User Scaling on System i

IBM® System i

Paul Swenson
paulswen@us.ibm.com
System i ERP, Lawson Team

Version Date: November 15 2007

Statement of Approval

Lawson M3 has reviewed, verified, and approved results for their applications which are shown in this report.

- Lars Strandner, Sr. Analyst, Lawson Product Development, lars.strandner@se.lawson.com

Introduction

This paper describes testing that was done to test flexibility, integration, and scalability of the new Lawson M3 7.1 running on a System i server.

Lawson M3 is a supplier of collaboration software which focuses on the manufacturing, maintenance, and distribution industries and serves many customers around the world. The Lawson M3 ERP solution tested here is a Java™-based application which runs on System i servers.

To demonstrate the abilities of Lawson M3 7.1 the Lawson M3 Order Entry Benchmark kit was used to simulate key elements of a typical customer transaction. The Benchmark scenario was modified to use a more typical think time. Thus the goal is to show a realistic number of users running each configuration tested.

The test scenario was run using the i5/OS® operation system, version V5R4, and using the new IBM Technology for Java Virtual Machine, or IT4J, available in V5R4. All Lawson M3 applications were installed on the same System i server. The test scenario was as follows: simulated users entered orders through Lawson M3 Workplace into Lawson M3 Business Engine, which includes full back-end processing to invoice the orders. The test team then ran the scenario using a System i model 570 system. Several different test scenarios were run to show the performance and scalability of Lawson M3 7.1 running on a System i server.

All tests achieved excellent results, including subsecond response times and at least 2,500 order lines processed per minute. Overall, Lawson M3 7.1 achieved excellent scalability when running on a System i server.

This report is also a demonstration of the capabilities of a System i server:

- The ability to run multiple complex workloads.
- The integration of the i5/OS operating system, DB2® UDB for System i, Java, and WebSphere® Application Server demonstrates the ability of a System i server to support new application models.
- Flexibility of a System i server to handle workloads on even small environments and grow as your business grows.

Benchmark Methodology

The IBM test team used the Lawson M3 Order Entry Benchmark kit for all results in this report. The benchmark kit uses a load generation tool to simulate a number of virtual users entering orders at a reasonable rate. The Order Entry transaction, OIS100, is what is used for this. The Order Entry transaction was chosen for its relative complexity and connection to a real life scenario. Further, it's easily repeatable, can be run infinitely and allows for a large variance in data.

The results from the order entry tests can be calculated into an entity called Universal Performance Unit or UPX. UPX is the sizing indicator used by Lawson M3 for customer sizings. The UPX is a theoretical transaction consisting of a average CPU time required for a typical customer load. A theoretical number of UPX'es per hour can calculated via a user number, activity and the production timeframe.

Thus even though a customer may have many other transactions types than just the order entry transaction that is running here, the results here likely still apply, since all other transaction can be derived via CPU time from this one transaction using UPX. It's is the overall performance of the Lawson M3 Business Engine to handle high volumes of transactions that is being stressed, the transaction type in not key in this, since all business logic share the same application foundation and architectural design.

A realistic number of virtual users, or vusers, were used to drive the scenario. The number of vuser ran on each configuration represents the maximum number of users that configuration can support at 70% CPU. This is was done to more closely represent a customer environment

Note: The number of vusers used for each configuration tested was based on sizings done by Lawson. These sizings were done assuming an average configuration. A customer sizing may require more or less CPU per user and thus the number of users a system can support may vary from what is shown here.

The benchmark scenario works as follows:

After logging in, the user performed an order entry operation consisting of seven steps:

1. Create a new order
2. Enter order line 1
3. Enter order line 2
4. Enter order line 3
5. Enter order line 4
6. Enter order line 5
7. Close the order

Throughput is calculated by counting the number of fully invoiced order lines once every minute throughout the benchmark test. Once the run was complete, the number of invoiced order lines per hour over a particular measurement interval was calculated from this data. The measurement interval was defined to be a 90 minute period beginning 15 minutes after the run started. The reported throughput metric for this benchmark is *number of invoiced order lines per hour*. An invoiced order line is one that has completed all of the interactive and batch processing required for that order and the order lines to have a status of 77 or completed. To demonstrate that the results were repeatable, each result consisted of two runs with identical parameters.

70% CPU Scaling Results

The chart below shows the scaling results that were achieved on three different processor configurations with a realistic number of users on each configuration. As the chart shows near linear scaling was achieved. In addition CPU was ~70%'s for all tests and sub-second transaction response times were maintained as well.

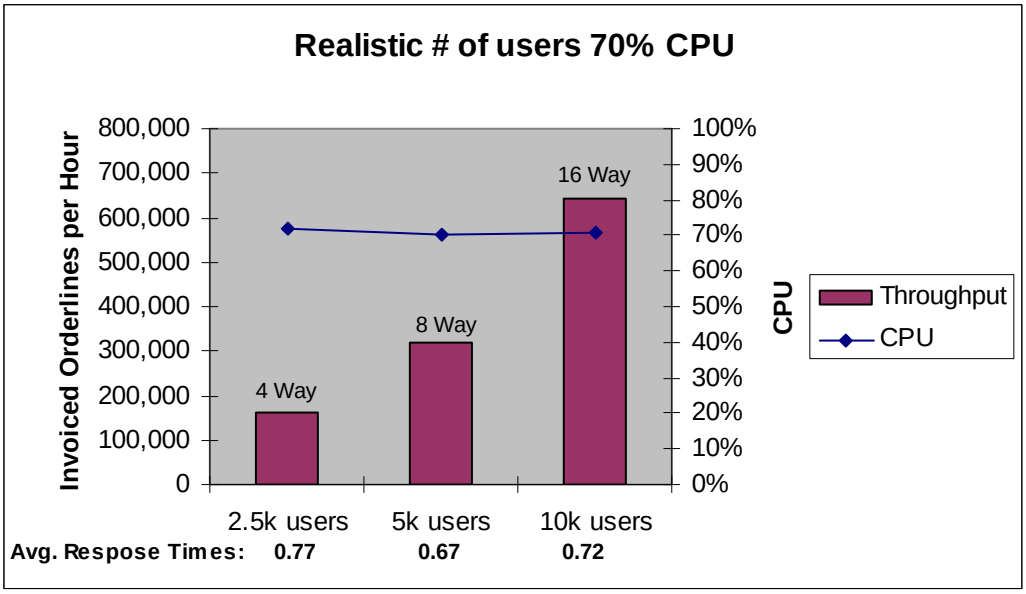

90% CPU Scaling Results

The chart below shows the scaling results that were achieved on two different processor configurations with a increase in the number of users so that CPU was driven to +90% on each configuration tested. As the chart shows near linear scaling was still achieved and sub-second transaction response times were maintained as well.

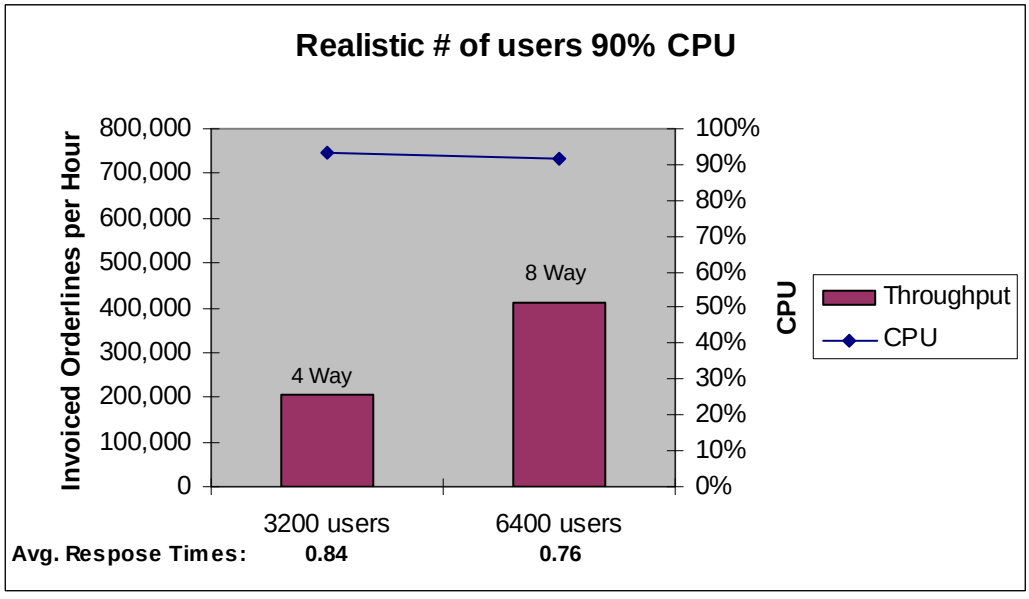

Summary

As the results in this paper have shown, Lawson M3 7.1 sees excellent scalability when running on a System i server using a realistic number of users which were driving the CPU to ~70%. Also CPU of 90% can be achieved as well when additional users are added. Thus both Lawson M3 and the System i server can grow along with your business. Even though only the Order Entry transaction was tested that does not mean the results here don't apply to other transaction types. Since all transactions can be derived from this single transaction, the results shown can be applied to almost all Lawson M3 transaction types.

Note: 16 CPUs at 90% CPU was not attempted due to lack of time to complete the necessary measurements in time for this report.

All of this also demonstrates the unique capabilities of a System i server to provide superior scaling, the ability to run multiple complex workloads, and the integration strength of the i5/OS operating system, DB2® UDB for i5/OS, Java, and WebSphere Application Server.

Appendix A.

System Configuration:

To gauge the performance of Lawson M3 software on a System i server the IBM test team configured a System i model 570 server with the following resources:

System i Power5+ Model 570-7749
Processors: 16
CPW rating: 58,500
Main storage: 256 GB
Disk: 180 arms (type 4327), 15k rpm, 70GB capacity
Disk configuration: Single ASP with device parity protection
Disk IOP: Disks spread across twelve 574F/5583 IOPs (390 MB write cache)
Network: 100Mbps Ethernet, full duplex

For the scaling tests each of the measurement sets used a different number of processors. The measurements sets used an identical configuration for disk, memory, network, operating system, and applications.

Processors	4	8	16
CPW Rating	16,700	31,100	58,500

Appendix B.

Detailed Results

4-way 70% CPU
The 4-way invoiced 160,386 order lines per hour in one run and 160,388 order lines per hour in the second run. In both runs, 2,500 virtual users were used.

Response time. Table 1 shows the average response time for each Web browser transaction on the 4-way. Response times do not include the average think time for each transaction and they measure the entire length of the run, not just the 90 minute measurement interval.

Step	Number per Order	Average Response Time (seconds)	
		Run 1	Run 2
Create Order Head	1	0.137	0.152
Insert Order Line	5	0.097	0.131
Close Order	1	0.083	0.025

Table 1. Average response times for the 4-way at 70% CPU

8-way 70% CPU
The 8-way invoiced 320,761 order lines per hour in one run and 320,696 order lines per hour in the second run. In both runs, 5,000 virtual users were used.

Response time. Table 2 shows the average response time for each Web browser transaction on the 8-way. Response times do not include the average think time for each transaction and they measure the entire length of the run, not just the 90 minute measurement interval.

Step	Number per Order	Average Response Time (seconds)	
		Run 1	Run 2
Create Order Head	1	0.139	0.148
Insert Order Line	5	0.099	0.090
Close Order	1	0.071	0.041

Table 2. Average response times for the 8-way at 70% CPU

16-way at 70% CPU
The 16-way invoiced 640,640 order lines per hour in one run and 641,214 order lines per hour in the second run. In both runs, 10,000 virtual users were used.

Response time. Table 3 shows the average response time for each Web browser transaction on the 16-way. Response times do not include the average think time for each transaction and they measure the entire length of the run, not just the 90 minute measurement interval.

Step	Number per Order	Average Response Time (seconds)	
		Run 1	Run 2
Create Order Head	1	0.143	0.065
Insert Order Line	5	0.102	0.103
Close Order	1	0.167	0.048

Table 3. Average response times for the 16-way at 70% CPU

4-way at 90% CPU
The 4-way invoiced 205,058 order lines per hour in one run and 205,081 order lines per hour in the second run. In both runs, 3,200 virtual users were used.

Response time. Table 4 shows the average response time for each Web browser transaction on the 4-way. Response times do not include the average think time for each transaction and they measure the entire length of the run, not just the 90 minute measurement interval.

Step	Number per Order	Average Response Time (seconds)	
		Run 1	Run 2
Create Order Head	1	0.215	0.155
Insert Order Line	5	0.12	0.127
Close Order	1	0.053	0.026

Table 4. Average response times for the 4-way at 90% CPU

8-way at 90% CPU
The 8-way invoiced 410,566 order lines per hour in one run and 410,665 order lines per hour in the second run. In both runs, 6,400 virtual users were used.

Response time. Table 4 shows the average response time for each Web browser transaction on the 8-way. Response times do not include the average think time for each transaction and they measure the entire length of the run, not just the 90 minute measurement interval.

Step	Number per Order	Average Response Time (seconds)	
		Run 1	Run 2
Create Order Head	1	0.15	0.126
Insert Order Line	5	0.094	0.099
Close Order	1	0.119	0.162

Table 5. Average response times for the 8-way at 90% CPU

Trademarks and Disclaimers

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at http://www.ibm.com/legal/copytrade.shtml.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered
trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.