Lawson M3 7.1 16 Way 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 uses almost no key think times. Thus the goal is to maximize throughput through the Business Engine and only a small number of virtual users are needed to drive a system.

The test scenario was run using the i5/OS® operation system, version V5R4, and 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 4,000 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 standard Lawson M3 Order Entry Benchmark kit for all results in this report. The benchmark kit uses a load generation tool to simulate a small number of virtual users entering orders at a fast 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.

Only a small number of virtual users, or vusers, were used to drive the scenario. The vusers were configured to wait an average of one second between each step. Thus there is almost no key think time between steps of a transaction. This is not at all representative of a customer environment. The goal of the benchmark is to achieve the maximum throughput from the Lawson M3 Business Engine on a given system configuration. The number of users that a given system could support can be calculated based on the CPW rating and by applying some general assumptions on user activity, production timeframe, and peak expectations. For example, using this type of calculation with general assumptions on the 16-way tested here results in 10,000 concurrently active users. Thus number of users that were used is not reported in this report, since that can be calculated and the expectation is, using the above example, that 10,000 users would see throughput in the range that is reported in this report.

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.

Scaling Results

The chart below shows the scaling results that were achieved. Four different processor configurations were tested. As the chart shows near linear scaling was achieved. In addition CPU was in the 90's for all tests and sub-second transaction response times were maintained as well.

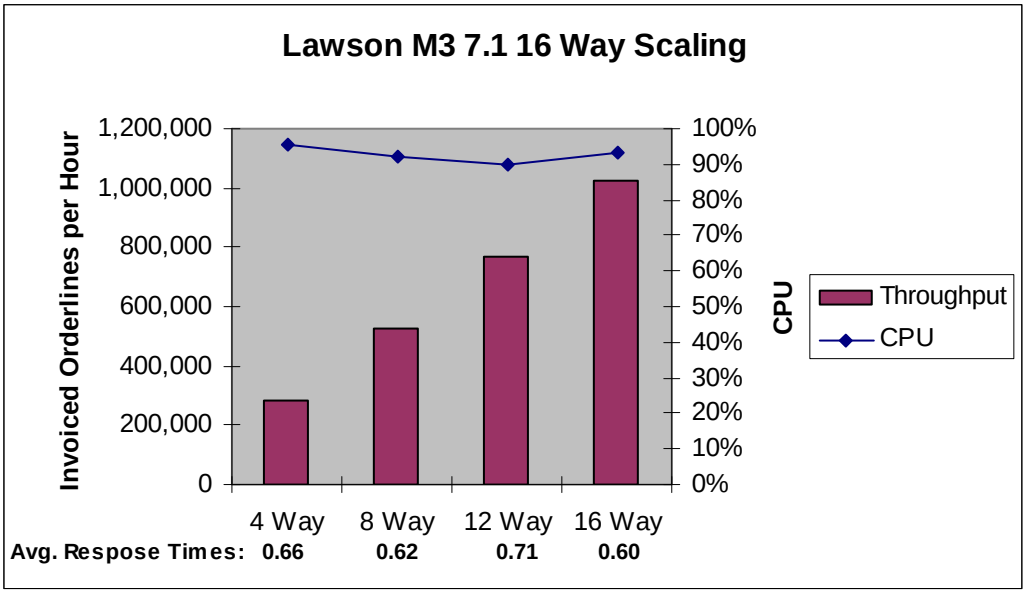

Lawson M3 7.1 Scaling vs. CPW

Commercial Processing Workload or CPW values are relative system performance metrics and reflect the relative system capacity. Each System i server has a CPW rating to show its relative system capacity. The chart below shows how the Lawson M3 7.1 results compare to the CPW ratings for each system configuration that was tested. As you can see Lawson M3 7.1 is scaling in-line with or slightly better than CPW.

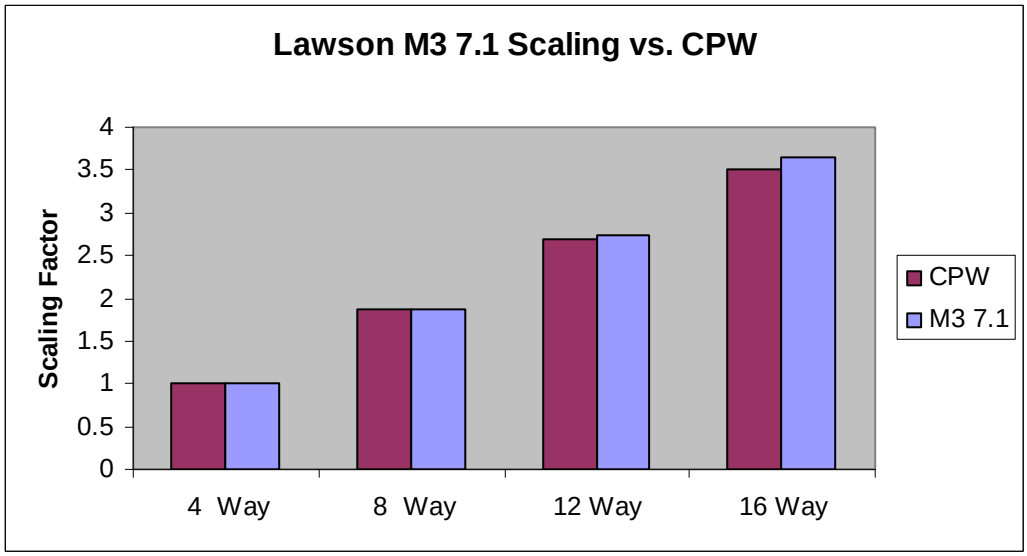

Summary

As the results in this paper have shown, Lawson M3 7.1 sees excellent scalability when running on a System i server and the scaling results are also in-line with CPW. Thus both Lawson M3 and the System i server can grow along with your business. The Benchmark scenario used here has almost no key user think times with the goal being to achieve the maximum throughput through the Lawson M3 Business Engine for each configuration that was test. Also 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.

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	12	16
CPW Rating	16,700	31,100	~45,000	58,500

Appendix B.

Detailed Results

4-way
The 4-way invoiced 282,134 order lines per hour in one run and 278,581 order lines per hour in the second run.

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 one-second 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	0.157	0.162
Insert Order Line	5	0.083	0.084
Close Order	1	0.077	0.081

Table 1. Average response times for the 4-way

8-way
The 8-way invoiced 524,853 order lines per hour in one run and 521,084 order lines per hour in the second run.

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 one-second 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	0.151	0.146
Insert Order Line	5	0.080	0.080
Close Order	1	0.077	0.073

Table 2. Average response times for the 8-way

12-way
The 12-way invoiced 762,171 order lines per hour in one run and 768,701 order lines per hour in the second run.

Response time. Table 3 shows the average response time for each Web browser transaction on the 12-way. Response times do not include the average one-second 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	0.158	0.165
Insert Order Line	5	0.093	0.092
Close Order	1	0.087	0.087

Table 3. Average response times for the 12-way

16-way
The 16-way invoiced 1,015,063 order lines per hour in one run and 1,036,560 order lines per hour in the second run.

Response time. Table 4 shows the average response time for each Web browser transaction on the 16-way. Response times do not include the average one-second 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	0.140	0.140
Insert Order Line	5	0.078	0.077
Close Order	1	0.068	0.067

Table 4. Average response times for the 16-way

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.