Infor M3 on IBM POWER7+ and using Solid State Drives

IBM Systems & Technology Group

Robert Driesch

cooter@us.ibm.com

This document can be found on the web, www.ibm.com/support/techdocs

Version Date: January 31, 2014

Table of Contents

Statement of Approval

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

- Lars Strandner, Sr. Analyst, Infor M3 Product Development, lars.strandner@infor.com

Introduction

This paper describes testing that was done with Infor M3 10.1 on the new IBM® POWER7+™ 740 and the IBM i 7.1 operating system. This report highlights the benefits of the latest IBM POWER® hardware and IBM i™ operating system as well as the benefits that usage of solid state drives or SSDs can provide. First it shows how runtime performance on the new POWER7+ has improved over a similarly configured POWER7® system. Second it shows the benefits SSDs can provide in an I/O constrained environment.

Overall the results show a 25% runtime performance improvement with the latest POWER7+ 740 model over a comparable POWER7 model. Also the results show that SSDs can allow for better application scaling in a disk constrained environment. Thus a significant improvement can be seen when moving to the latest IBM POWER hardware and IBM i operating system.

Benchmark Methodology

The IBM test team used the Infor 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 Infor M3 for customer sizings. The UPX is a theoretical transaction consisting of an average CPU time required for a typical customer load. A theoretical number of UPX'es per hour can be 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 Infor M3 Business Engine to handle high volumes of transactions that is being stressed, the transaction type is 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, however, the goal of the benchmark is to achieve the maximum throughput from the Infor M3 Business Engine on a given system configuration.

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 all vusers had 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.

IBM POWER7+ Results

System Configuration:

POWER7 750:

Model:	IBM POWER7 750	**Edition Feature**:	8336
Number of Processors:	4	**Chip Speed**:	3.55 GHz
CPW Rating:	26,200		
Main Storage:	64 GB		
Disk:	48 arms (type 433D), 15k rpm, 280GB capacity		
Disk Configuration:	Single ASP with device parity protection		
Disk IOP:	Disks spread across one 572F IOPs (390 MB write cache)		
Network:	100Mbps Ethernet, full duplex		

POWER7+ 740:

Model:	IBM POWER7+ 740	**Edition Feature**:	EPCR
Number of Processors:	4	**Chip Speed**:	4.2 GHz
CPW Rating:	32,250		
Main Storage:	32 GB		
Disk:	8 arms (type 19A1), 15k rpm, 283GB capacity		
Disk Configuration:	Single ASP with device parity protection		
Disk IOA:	Disks spread across one 57CB IOA (1.8 GB write cache)		
Network:	1Gbps Ethernet, full duplex		

Results:

The chart below shows the results for the new IBM POWER7+ 740 versus a similarly configured IBM POWER7 750. As the chart shows a 25% improvement is seen when moving up to the new IBM POWER7+ 740.

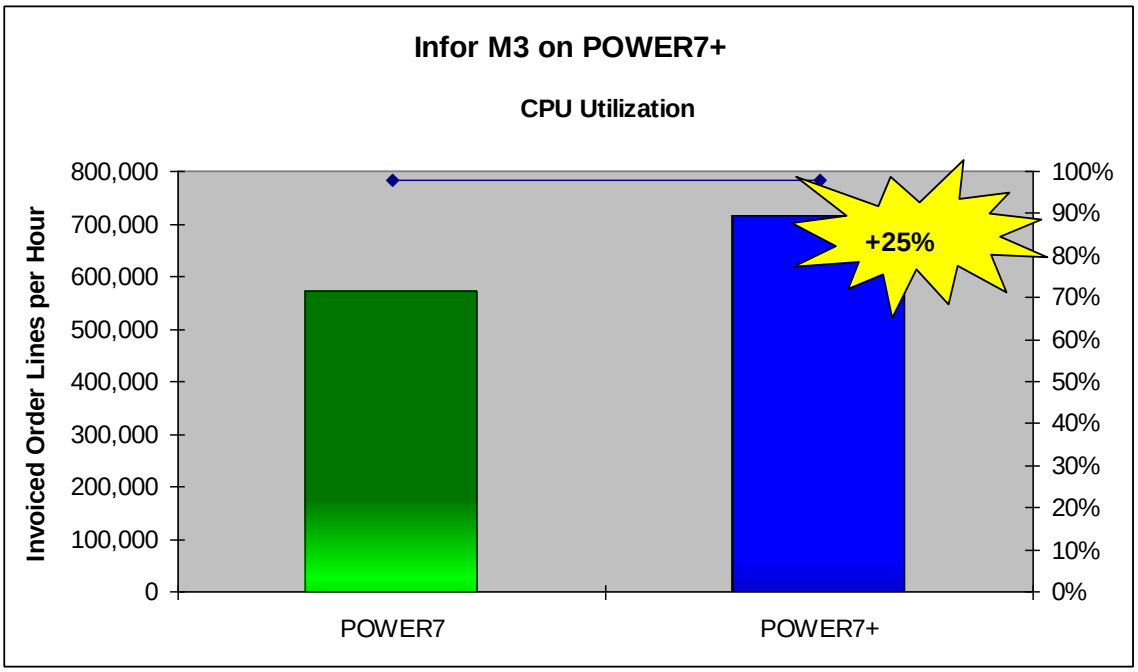

Figure 1 Invoiced Order Lines per Hour at max CPU Utilization

SSD Results

System Configuration:

Scenario 1 (CEC + Spinning):

Model:	IBM POWER7+ 720	**Edition Feature**:	EPCM
Number of Processors:	4	**Chip Speed**:	3.6 GHz
CPW Rating:	28,150		
Main Storage:	32 GB		
Disk:	8 spinning (type 19A1), 15k rpm, 283GB capacity		
Disk Configuration:	Single ASP with device parity protection		
Disk IOP:	Disks spread across one 57CB IOA (175 MB write cache)		
Network:	1Gbps Ethernet, full duplex		

Scenario 2 (CEC + SSD):

Model:	IBM POWER7+ 740	**Edition Feature**:	EPCR
Number of Processors:	4	**Chip Speed**:	4.2 GHz
CPW Rating:	32,250		
Main Storage:	32 GB		
Disk:	8 SSDs (type 58B8), 387GB capacity		
Disk Configuration:	Single ASP with device parity protection		
Disk IOP:	Disks spread across one 57CB IOA (175 MB write cache)		
Network:	1Gbps Ethernet, full duplex		

Scenario 3 (Expansion + Spinning):

Model:	IBM POWER7+ 740	**Edition Feature**:	EPCR
Number of Processors:	4	**Chip Speed**:	4.2 GHz
CPW Rating:	32,250		
Main Storage:	32 GB		
Network:	1Gbps Ethernet, full duplex		
Disk Configuration:	Single ASP with device parity protection		
Disk:	8 spinning (type 19A1), 15k rpm, 283GB capacity		
Disk IOP:	Disks spread across one 57B5 IOA (1.8 GB write cache)		

Results:

The chart below shows the results of the three scenarios tested. *Scenario 1 (CEC + Spinning)* had 8 spinning disks under an IOA with a small cache; this was a disk constrained environment. *Scenario 2 (CEC + SSD)* made use of 8 disks with the same IOA as *Scenario 1*, but instead using SSDs. *Scenario 3 (Expansion + Spinning)* had 8 spinning disks under an IOA with a large cache.

As the chart below shows, *Scenario 1* was unable to fully utilize a four CPU partition. In repeating that same test making use of SSDs, *Scenario 2*, we can now successfully scale up to eight CPUs. *Scenario 3* shows that the same results can be achieved with spinning disks with the addition of an expansion drawer and an IOA with a large write-cache. Hence SSDs can allow a system to be fully utilized making use of less hardware than what is required with spinning disks.

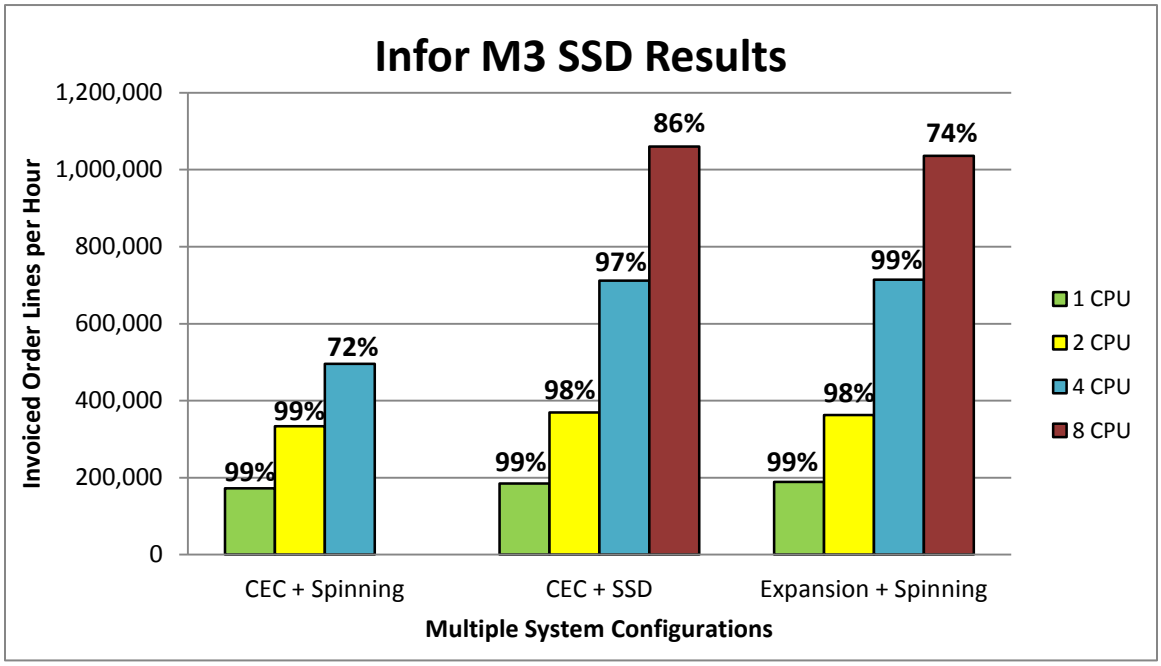

Figure 2 Invoiced Order Lines per Hour across multiple disk configurations

Note: CPU utilization was measured at 95% or higher for all runs listed in Figure 2.

Summary

As the results in this paper have shown, Infor M3 10.1 sees significant performance benefit when moving to the latest IBM POWER hardware and IBM i operating system.

Overall the results show a 25% runtime performance improvement with the latest POWER7+ 740 model over a comparable POWER7 model. Also the results show that SSDs can allow for better application scaling in a disk constrained environment. Thus a significant improvement can be seen when moving to the latest IBM POWER hardware and IBM i operating system.

Appendix A

Detailed Results: IBM POWER7+

POWER7 750:

The POWER7 750 invoiced 564,472 order lines per hour in one run and 565,352 order lines per hour in the second run. In both runs, 237 virtual users were used.

Table 1 shows the average response time for each Web Browser transaction on the POWER7 750. 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	0.122	0.126
Insert Order Line	5	0.069	0.072
Close Order	1	0.066	0.069

Table 1 Average response times for the POWER7 750

POWER+7 740:

The POWER7+ 740 invoiced 720,535 order lines per hour in one run and 708,019 order lines per hour in the second run. In both runs, 300 virtual users were used.

Table 2 shows the average response time for each Web Browser transaction on the POWER7+ 740. 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	0.101	0.097
Insert Order Line	5	0.066	0.066
Close Order	1	0.057	0.053

Table 2 Average response times for the POWER7+ 740

Detailed Results: Scenario 1 (CEC + Spinning)

Scenario 1 with 1 CPU:

Scenario 1 with 1 CPU invoiced 172,256 order lines per hour in one run and 172,007 order lines per hour in the second run. In both runs, 79 virtual users were used.

Table 3 shows the average response time for each Web Browser transaction for *Scenario 1* with 1 CPU. 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	0.282	0.286
Insert Order Line	5	0.163	0.168
Close Order	1	0.150	0.155

Table 3 Average response times for Scenario 1 with 1 CPU

Scenario 1 with 2 CPUs:

Scenario 1 with 2 CPUs invoiced 334,100 order lines per hour in one run and 332,974 order lines per hour in the second run. In both runs, 145 virtual users were used.

Table 4 shows the average response time for each Web Browser transaction for *Scenario 1* with 2 CPUs. 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	0.163	0.169
Insert Order Line	5	0.098	0.105
Close Order	1	0.087	0.092

Table 4 Average response times for Scenario 1 with 2 CPUs

Scenario 1 with 4 CPUs:

Scenario 1 with 4 CPUs invoiced 495,916 order lines per hour in run one. In this run, 200 virtual users were used.

Note: With 4 CPUs this scenario started showing signs of a disk constrained environment making it difficult to get 2 consistent consecutive runs. In this case only, just one data point was given to show the maximum order lines per hour that could be achieved in this scenario.

Table 5 shows the average response time for each Web Browser transaction for *Scenario 1* with 4 CPUs. 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
Create Order Head	1	0.023
Insert Order Line	5	0.038
Close Order	1	0.048

Table 5 Average response times for Scenario 1 with 4 CPUs

Detailed Results: Scenario 2 (CEC + SSD)

Scenario 2 with 1 CPU:

Scenario 2 with 1 CPU invoiced 184,932 order lines per hour in one run and 185,010 order lines per hour in the second run. In both runs, 85 virtual users were used.

Table 6 shows the average response time for each Web Browser transaction for *Scenario 2* with 1 CPU. 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	0.284	0.273
Insert Order Line	5	0.174	0.133
Close Order	1	0.173	0.162

Table 6 Average response times for Scenario 2 with 1 CPU

Scenario 2 with 2 CPUs:

Scenario 2 with 2 CPUs invoiced 367,644 order lines per hour in one run and 370,541 order lines per hour in the second run. In both runs, 157 virtual users were used.

Table 7 shows the average response time for each Web Browser transaction for Scenario 2 with 2 CPUs. 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	0.140	0.128
Insert Order Line	5	0.098	0.086
Close Order	1	0.090	0.077

Table 7 Average response times for Scenario 2 with 2 CPUs

Scenario 2 with 4 CPUs:

Scenario 2 with 4 CPUs invoiced 709,475 order lines per hour in one run and 714,639 order lines per hour in the second run. In both runs, 285 virtual users were used.

Table 8 shows the average response time for each Web Browser transaction for *Scenario 2* with 4 CPUs. 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	0.067	0.061
Insert Order Line	5	0.045	0.041
Close Order	1	0.033	0.030

Table 8 Average response times for Scenario 2 with 4 CPUs

Scenario 2 with 8 CPUs:

Scenario 2 with 8 CPUs invoiced 1,066,248 order lines per hour in one run and 1,067,767 order lines per hour in the second run. In both runs, 440 virtual users were used.

Table 9 shows the average response time for each Web Browser transaction for *Scenario 2* with 8 CPUs. 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	0.063	0.069
Insert Order Line	5	0.046	0.049
Close Order	1	0.045	0.037

Table 9 Average response times for Scenario 2 with 8 CPUs

Detailed Results: Scenario 3 (Expansion + Spinning)

Scenario 3 with 1 CPU:

Scenario 3 with 1 CPU invoiced 186,749 order lines per hour in one run and 190,709 order lines per hour in the second run. In both runs, 90 virtual users were used.

Table 10 shows the average response time for each Web Browser transaction for *Scenario 3* with 1 CPU. 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	0.364	0.353
Insert Order Line	5	0.217	0.198
Close Order	1	0.229	0.212

Table 10 Average response times for Scenario 3 with 1 CPU

Scenario 3 with 2 CPUs:

Scenario 3 with 2 CPUs invoiced 363,859 order lines per hour in one run and 360,997 order lines per hour in the second run. In both runs, 157 virtual users were used.

Table 11 shows the average response time for each Web Browser transaction for Scenario 3 with 2 CPUs. 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	0.175	0.179
Insert Order Line	5	0.107	0.110
Close Order	1	0.108	0.109

Table 11 Average response times for Scenario 3 with 2 CPUs

Scenario 3 with 4 CPUs:

Scenario 3 with 4 CPUs invoiced 720,535 order lines per hour in one run and 708,019 order lines per hour in the second run. In both runs, 300 virtual users were used.

Table 12 shows the average response time for each Web Browser transaction for Scenario 3 with 4 CPUs. 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	0.101	0.097
Insert Order Line	5	0.066	0.066
Close Order	1	0.057	0.053

Table 12 Average response times for Scenario 3 with 4 CPUs

Scenario 3 with 8 CPUs:

Scenario 3 with 8 CPUs invoiced 1,038,240 order lines per hour in one run and 1,033,785 order lines per hour in the second run. In both runs, 440 virtual users were used.

Table 13 shows the average response time for each Web Browser transaction for Scenario 3 with 8 CPUs. 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	0.105	0.108
Insert Order Line	5	0.068	0.073
Close Order	1	0.061	0.190

Table 13 Average response times for Scenario 3 with 8 CPUs

Trademarks and Disclaimers

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AS/400®	POWER6®
i5/OS®	POWER7®
IBM®	POWER7+™
IBM i ™	Redbooks®
iSeries®	Redbooks (logo) ®
OS/400®	System i™
POWER®	WebSphere®
POWERHA®	

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 and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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.