WebSphere® for System z® Prescriptive Use Cases

Continuous Improvements to System z Performance and Price

June 9, 2011

Addendum: Linux for System z Comparison

October 28, 2011

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. ALTHOUGH EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS DOCUMENT, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.

IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS DOCUMENTOR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS DOCUMENT IS INTENDED TO, OR SHALL HAVE THE EFFECT OF CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS; OR ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.

The prices used in this comparison are current as of the date of this paper and are subject to change by IBM without notice.

z/OS Software Used in This PUC V3 Effort

z/OS® V1.11 PUT level 1005

CTG V7.0

CICS® Transaction Server for z/OS® V3.2

DB2® V10 for z/OS, PUT1012

WPS for z/OS V7.0.0.2 64-bit

WAS for z/OS V7.0.0.11 64-bit

JDBC V3.62.57

System z Hardware Used in This PUC V3 Effort

z196 2817 in LPAR mode with:

- 2 general purpose and 1 zIIP processor (for single LPAR tests)
- 1 general purpose and 1 zIIP processor (additional LPAR for Type 4 driver test)

In this paper a comparison is made to a prior test using the following environment. The primary difference in the two environments was the use of zEnterprise hardware and the use of DB2 V10. Also, slightly newer levels of WPS, WAS and z/OS were used in the new environment compared to the prior test.

z/OS Software Used in PUC V2 (July 2010) Test Effort

z/OS® V1.10

CTG V7.0

CICS® Transaction Server for z/OS® V3.2

DB2® V9.10 for z/OS, PUT0903

WPS for z/OS V7.0.0.1 64-bit

WAS for z/OS V7.0.0.9 64-bit

System z Hardware Used in PUC V2 (July 2010) Test Effort

z10 2097 in LPAR mode with:

2 general purpose processors

Table of Contents

Executive Overview

In September 2009 and July 2010 IBM released the results of studies that proved the price/performance competitiveness for deploying WebSphere® middleware natively on the System z® platform (z10™ hardware running the z/OS® operating system). These studies validated the premise that running realistic WebSphere workloads on z/OS, collocated with existing online and batch workloads and the data associated with these workloads, is price/performance competitive with running these same workloads on alternative (distributed) platforms (for example, AIX running on System p®) communicating with the same back-end z/OS resources (for example, CICS and DB2).

Customer response to the original white papers was extremely positive, fueling continued interest and increased customer investment in the z/OS platform as a platform of choice for deploying similar WebSphere workloads. In those original white papers, we provided our customers with answers to the following questions:

- Why and when should I deploy WebSphere middleware on System z?
- Is the price/performance of System z competitive with alternative distributed deployment options? [1]

What have we done since publishing those first 2 white papers?

Is IBM innovating - making the kind of technical investments that customers expect?

This current study confirms the cost benefits of deploying new WebSphere workloads on System z for customers who have already invested in System z from both a hardware and a software perspective. Additionally, this prescriptive use cases effort reassures customers that IBM is continuing its technical investment in the System z platform.

In fact, we are publishing this paper now because new testing shows significant improvements in the performance and price of the newer releases of System z hardware and software:

1. Performance improvements as measured in terms of throughput rates[2]
2. Improvements in the cost of 10K transactions

As you will see, the results of this latest prescriptive use cases effort, summarized in this document, show that a WebSphere workload collocated in the same LPAR with a Type 2 driver offers the best option from both a performance and a price perspective.

[1] Further analysis of the Linux for System z option and new analysis of the new zEnterprise solution with zBX deployment option are under consideration for a future paper.

[2] The previous publications introduced the concept of business services per hour to describe the unit of measure from the test. In this paper the cost of running 10,000 transactions is introduced to provide an additional view of the results.

Introduction to the WebSphere for System z Prescriptive Use Cases

Before the WebSphere for System z prescriptive use cases effort, IBM provided limited collateral that exhibited solid proof points for customers who wanted to understand the cost competitiveness of deploying new WebSphere workloads on System z. Most of the available information was anecdotal data related to the qualities of service (QoS) provided by WebSphere Application Server or WebSphere Process Server on z/OS on System z. The QoS of z/OS (such as security, resilience and scalability) are generally well understood by customers.

Hypothesis

The expectation for the prescriptive use cases effort is that the performance and price improvements of the latest class of System z hardware and selected key software provide solid grounds for continued customer investment in System z. The choice of system configuration can impact the performance and price of the System z solution; therefore, the use cases discussed here also shed light on implementation choices that can help to optimize the performance and price of the workload on System z. Stated more succinctly, the purpose of this effort is to demonstrate the following hypothesis:

A WebSphere workload collocated with traditional z/OS online and batch workloads in the same LPAR with a local connection offers the best value for z/OS customers.

High-Level Description of the Validation Effort Deliverable

The key deliverable from the validation effort is documented performance and price competitiveness based on quantifiable costs, thereby establishing the proof points for continued customer investment and deployment in z/OS on System z. How did we do this? New innovations in System z hardware and software, together with the implementation of recommended system configurations and new pricing enhancements, offer significant improvements in performance and cost.

To Which Customers Does This Apply?

The prescriptive use cases focus on existing System z and z/OS customers who have already deployed CICS, IMS™, and/or DB2 workloads on z/OS; that is, customers who have an investment in online and batch processing on z/OS and are considering WAS and z/OS as a viable option for new Java-based transactional workloads. These customers:

- Might have extra capacity to support new Java workloads.
- Understand the value of z/OS and its inherent qualities of service (such as security, resilience, governance, scalability, manageability, and energy efficiency).
- Might elect to add capacity to further exploit the advantages of their System z environment.

A Customer Business Scenario for Investment in New Hardware and Software: Key Steps

First, a customer business challenge is identified. One such challenge, Enterprise Business Transformation (EBT), is necessary for many IBM customers. An example of EBT is when a bank acquires additional banks and must consolidate and then deploy its existing business processes to the newly acquired banks.

Second, a customer accepts Service Oriented Architecture (SOA) as the infrastructure to deploy the EBT solution. IBM has been very successful in helping customers to understand that IBM has a solution that implements SOA: the WebSphere portfolio.

In the third step, given that IBM understands the business problems and can provide a good architecture, infrastructure, and product set, the next question is raised: W*here should the new workload be deployed*? In the first 2 prescriptive use case white papers, we proved that deploying the new workload on System z is price competitive with deploying to alternate distributed platforms. Since those papers were published, we have been concentrating our attention on the price performance of the z/OS option[3].

The prescriptive use cases can be applied when performance and price assessments are needed, such as:

- Deployment of new WebSphere applications and business processes.
- Significant anticipated growth of an existing WebSphere application; that is, organic growth or a merger or acquisition.
- Optimization of an existing server or set of existing servers and WebSphere application infrastructure.
- Significant changes to business requirements for improved business agility and qualities of service in the IT infrastructure (maybe as a result of a security audit or disaster recovery planning).

Note: It is important to understand the value of keeping the hardware and software current. IBM has worked to improve the price/performance of the environment through changes in technology, software configurations, and pricing policies.

Applying the Rule of Four

Four basic rules, applied here in the form of questions, can help to determine whether to deploy on System z:

1. Is the hardware current (z10, z196), that is - is the customer positioned to take advantage of the latest hardware capabilities and the best software pricing?
2. Does the business application require frequent access to z/OS data (DB2 z/OS or IMS-DB) or transactions (CICS, IMS, WebSphere MQ)?.
3. Are the related workloads highly dynamic, unpredictable, or of high business value?
4. Can the workloads take advantage of specialty processors (IFLs, zIIPs, or zAAPs)?

Rule 1 says that you already have an investment in System z. It just makes sense to leverage your existing investment. This analysis assumes zEnterprise System 196 but there is more to the solution than hardware. Being on the latest System z hardware not only has the performance benefit of the new zEnterprise, but for 2011 there is a new software pricing option

as well. Integrated Workload Pricing (IWP) is designed to remove barriers to collocating WebSphere applications and existing transaction managers of CICS or IMS. More information about this pricing technique can be found in the IBM U.S. Announcement letter 211-011 dated January 18, 2011. IWP can provide:

- Significant improvement in price/performance when running new WAS and z/OS Java collocated workloads with CICS or IMS.
- Improved alignment between software utilization and software charges.
- Increased flexibility in application deployment and software licensing.

Rule 2 says that collocating the applications (WPS or WAS) close to the data (DB2 z/OS) and transactions (such as CICS, IMS, or WebSphere MQ) is a smart choice. z/OS excels here and is the platform of choice when data is on z/OS. Having everything under the control of z/OS has many benefits that can provide significant value to customer applications and business. This analysis assumes DB2 for z/OS as the data store and shows the value of specific data access configuration options.

Rule 3 says that System z excels in providing high hardware utilization rates and running business-critical applications. A z/OS on System z deployment addresses this rule. This rule also leads to further platform differentiation and operational benefits; that is, WPS can inherit the z/OS clustering technology (Sysplex) from WAS for z/OS . WPS supports, for example, the WAS z/OS controller/servant architecture and also exploits Workload Manager (WLM) in z/OS.

This analysis also assumes workload growth in the environment and not a single snapshot-in-time for price performance. z/OS excels at supporting hundreds of concurrent users and is therefore included as part of the dynamic nature of the workload.

Rule 4 relates to specialty processors, which represent significant price/performance benefits for deploying on z/OS on System z. If the application workload exploits these specialty processors, System z often delivers a lower implementation cost, as proven by the testing described in this document. This analysis exploits zAAP for WebSphere JAVA processing and zIIP for DB2 connectivity.

The focus of the prescriptive use cases is on quantifying Rule 2 and Rule 4 (collocation and the leveraging of specialty processors) so that when a new workload is placed on the same system as the existing workloads, it provides a distinct and cost-justifiable benefit. Also providing a cost benefit is the consolidation of new workloads with existing distributed workloads that currently run on under-utilized distributed processors – thereby lowering the costs of managing multiple servers.

There are numerous value points for deploying on z/OS on System z, only one of which is collocation – a key consideration because it can be quantified in terms of throughput, CPU consumption, and cost.

Description of the Test and Workload

The comparison included:

- **A single, realistic scenario**: a mixed workload based on WebSphere Process Server (WPS) Version 7 and enterprise business process modernization.
- **Proof points** (Two units of measure are provided):
    - *Business services per hour*. (This measurement is described in "Proof Points" in this paper.)
    - *Cost of 10K transactions.*
- **Topology:**
    - System z10 vs zEnterprise (z196)
    - Separate logical partition (LPAR[4]) with Type 4 driver versus the same LPAR with Type 2 driver

Single, Realistic Scenario – Mixed Workload

A WPS workload was chosen for this comparison for two reasons:

- A set of real-world z/OS applications (new and existing), various configurations, and workload generation tools were available for use within IBM.
- *Business Process Management* (BPM) using WebSphere Process Server (WPS) is a strategic and growing investment area for many companies.

A mixed, collocated workload was chosen, which represents a combination of real workloads that are modeled after some of IBM's largest client test environments and workloads for WPS on z/OS.

Compared to a typical benchmark, which involves a single transaction or a single application that is run to high stress levels or transaction rates, this effort involved a mixed application representing a real-world scenario, in which different types of Java™ and non-Java transactions were blended.

Note: While this analysis is focused on business process management, the conclusions set forth in this document have applicability to other WebSphere Application Server workloads beyond BPM. The BPM workload on WPS is a set of transactions not unlike any other transactions that would be run in a WebSphere environment. The data presented is specific to the use case tested but is felt to be representative of any WebSphere workload.

Proof Points

Business Services per hour and the ***cost of 10K transactions*** were selected as proof points for this test. Both are discussed in the following paragraphs.

Note: In this context, a Business Service is defined as a collection of related, structured activities that together produce a service or product that meets the needs of a client. These activities are critical to virtually any organization because they can be used to generate revenue and they often represent a significant percentage of a company's operating costs. Business Services represent meaningful business-critical transactions to virtually all IBM customers.

[4] Customers are generally familiar with the term *LPAR* in the System z context as the ability of System z hardware and firmware to provide a virtual computer environment, separate from other virtual computer environments running on the same physical hardware.

Examples of the Business Services That Were Used in This Effort

A few of the WebSphere Process Server transactions and the services that are called include:

- **Enter a New Order** - A NewOrder request message is received and the request is put on a HumanTask work list for processing. When the work list item is selected for processing, it is either accepted, after which the NewOrder service component architecture (SCA) service is called and a NewOrder is entered into the WebSphere Process Server system, or it is rejected and the process bypasses the NewOrder SCA service and ends.
- **Give a Price Quote (or Price Change) –** An item number and quantity are entered into a list. The application generates a price quote for the information that was entered by calling the PriceQuote transaction. From the price quote reply, the application selects the first item in the list and changes the price for that item by calling the PriceChange transaction. The test case executes a read-and-update transaction from a process.
- **Process Payments (Credit/Debit) -** Funds transfer information is entered on the Web application screen. The application executes a credit transaction and a debit transaction from two separate accounts. This test case runs read and update transactions from a process.

How the Business Service Measurement Differs From Typical Benchmarks

The test was centered on running steady-state workloads in System z just as a typical customer would experience in a peak hour. By design, we kept the workloads on the systems (z10 and z196) exactly the same; the amount of memory that was available was the same, and the actual business services that were run were relatively consistent across the measured runs. This strategy provides a foundation that is sound and easy to understand in making a comparison.

We use the phrase *business services per hour* to show the relative workload of the different environments. This unit of measure is intended to describe the real world nature of the workload – it is a metric that a business may want to monitor in the day-to-day running of a company. The analysis for this paper was run for use cases of 87000, 176000, and 262000 business services per hour; we show the price/performance with 262000 business services per hour, as shown in Figure 1, 3 Year Total Cost of Acquisition (TCA).

This prescriptive use cases effort is *not* a performance benchmark; that is, it does not involve running the machine at the highest utilization to arrive at a maximum performance measurement. Industry standard benchmarks are typically difficult to correlate to real-world workloads because they represent a unit of work that is standardized. Measurements against such benchmarks are hard (or impossible) to relate to the activities of a large enterprise with mission-critical data. The prescriptive use cases effort focuses on real-world workloads that reflect the type of activities that are familiar to customers.

What Differentiates the 10K Transaction Measurement

In this paper we introduce a new unit of measure for comparison purposes. Based on the business services per hour measurements and the associated 3 year total cost of acquisition (TCA), we derived and examined the cost of running 10,000 transactions. The data point we chose to compare was 26200 business services per hour, as shown in Figure 1.

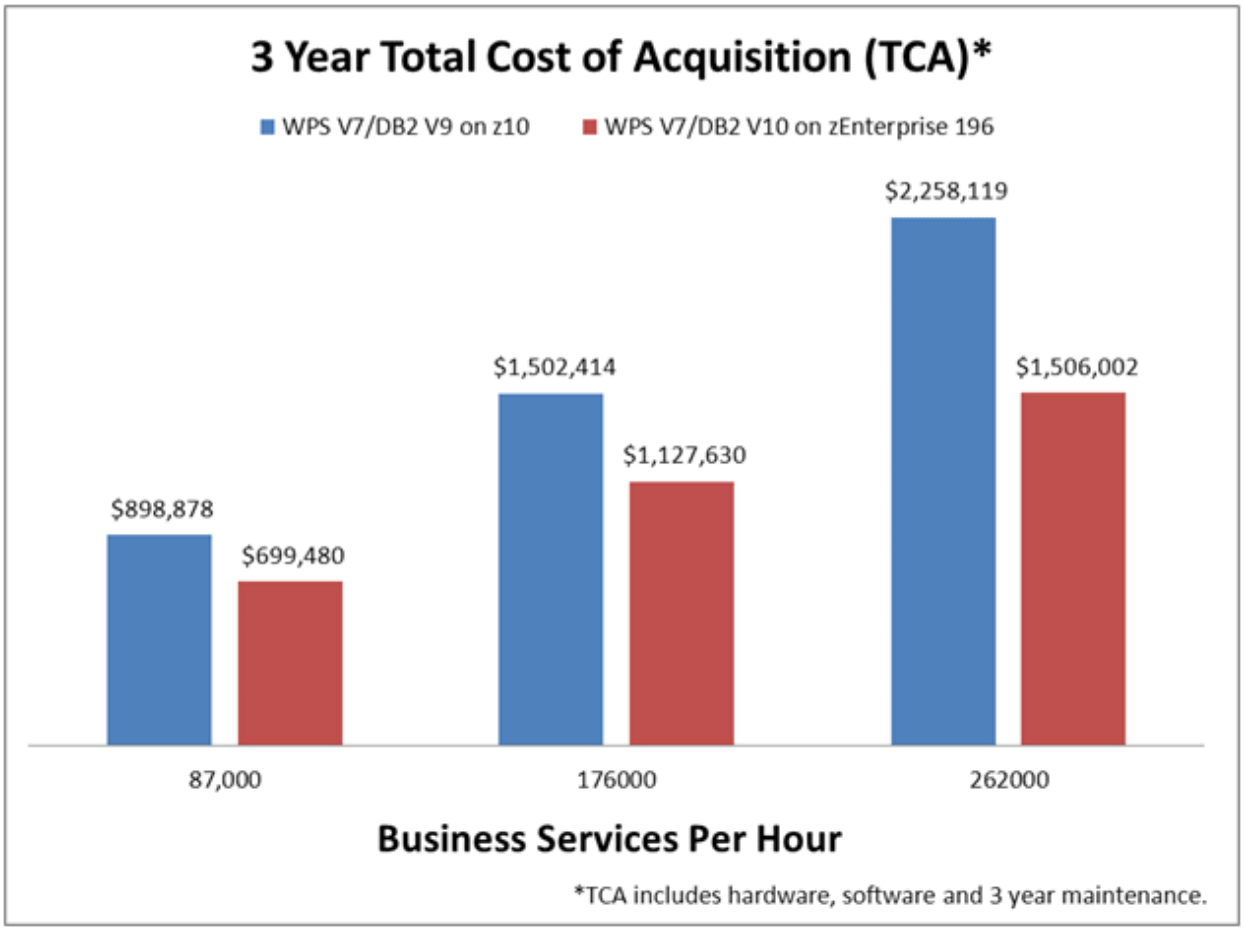

Figure 1 – 3 Year Total Cost of Acquisition (TCA)

Like the business service measurement, the 10K transaction measurement is not meant to be a performance benchmark. The internal throughput rate (ITR) is measured as the external throughput rate divided by the CPU percentage used.

Why use this measurement? Most ITR measurements look at the capacity of the hardware to see the maximum number of transactions that can be driven out of the maximum CPU capacity. Customers, however, are focused on specific workloads. While the “maximum” type of numbers might be interesting and useful in other ways, we wanted to look at a new workload that would use a percentage of the available CPU capacity, taking into consideration how many transactions a customer might need to drive. With this understanding, we decided to focus on driving incremental levels of transactions at less than full CPU capacity.

Topology: What This Effort Compared

This latest effort compared the cost of 10K transactions using:

- z10 with WPS version 7 and DB2 version 9
    - In a separate LPAR with Type 4 Java Database Connectivity (JDBC) driver

  - In the same LPAR with Type 2 JDBC driver
- z196 with WPS version 7 and DB2 version 10
  - in a separate LPAR with Type 4 JDBC driver
  - In the same LPAR with Type 2 JDBC driver

Note: *Type 2* and *Type 4* describe the software configuration of the communication method used between software on System z. Type 2 (local) communication is cross-memory communication between two pieces of software on the same LPAR. Type 4 (remote) denotes communication using standard TCP/IP network communications – even when communicating within the same LPAR.

Separate LPAR Configuration with Type 4 JDBC Driver

The configuration in Figure 2 shows WPS running in a new LPAR with CICS and DB2 running in a separate LPAR. SOA-SCA services are used to access CICS transactions. The Type 4 driver is used for DB2 access.

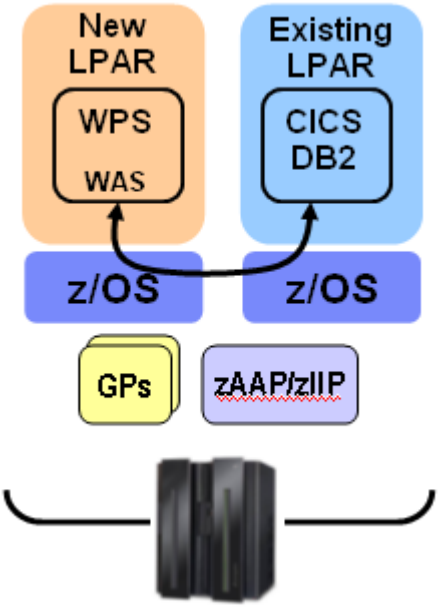

Figure 2 – Separate LPAR (Type 4)

The typical multiplatform topology for high availability includes the traditional applications, CICS and DB2, which house the key business data for the traditional batch and online workloads and the new WebSphere workload. The new workload communicates through the connection points to the existing online and batch environment, and uses the CICS and DB2 data on z/OS.

Note: Figures 2 and 3 are general representations of the system; they do not represent all of the specific features that were included in the pricing calculations. Actual configurations could include zIIPs, zAAPs, IFLs, etc. These figures represent a starting option for 80 – 90% of the intended audience for this paper.

Same LPAR Configuration with Type 2 JDBC Driver

In Figure 3, the configuration shows WPS running in the same LPAR with CICS and DB2 running in the same LPAR. SOA-SCA services are used to access CICS transactions. The Type 2 adapter is used for DB2 access.

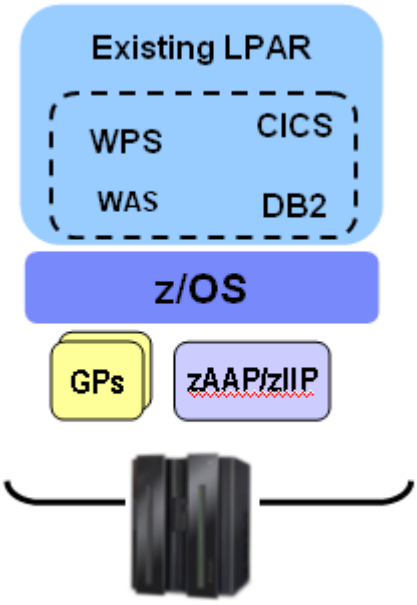

Figure 3 – Same LPAR (Type 2)

The following items were configured, tested, priced, and compared in the selected topologies:

- Business Services rates and the cost for adding a WPS workload to the traditional z/OS (collocated) workload
- 10K transactions

Execution of the Workloads

Validating the prescriptive use cases included measuring and pricing the deployment of the collocated z/OS on System z topology:

- **Collocated topology:** Test runs to determine Business Services per hour at multiple utilization rates with the new workload running in a single z/OS image
- **Pricing** the topologies and assessing the cost comparison.

Results and Impact of the Validation

The following analysis is divided in two different ways:

1. Performance improvements
2. Cost improvements

Performance was measured in terms of internal throughput rate (ITR), that is, the number of transactions/CPU busy second. It is worth noting that the environment used for this data collection was treated as a real-world scenario of adding workload to an existing CICS and DB2 environment and was not simply an exercise of seeing how fast the benchmark workload could be executed. The internal throughput rate is, in fact, based on 262, 000 business services per hour. The point is z/OS - given adequate hardware - can take on this additional WPS workload and show performance improvements utilizing the latest z hardware and software.

- For the z10 system configured with a separate LPAR and Type 4 driver, running WPS and DB2 version 9, the ITR was measured at 78.77.
- For the z10 system configured with WPS and DB2 version 9 running in the same LPAR with a Type 2 driver, the ITR was measured at 114.14, a **44.9 % improvement** over the z10 system with a Type 4 driver.
- For the z196 system configured with a separate LPAR and Type 4 driver, running WPS and DB2 version 10, the ITR was measured at 144.49, a **26.6% improvement** over the z/10 system with a Type 2 driver and running DB2 version 9.
- For the z196 system configured with WPS and DB2 version 10 running in the same LPAR with a Type 2 driver, the ITR was measured at 203.38, a **40.7% improvement** over the z196 system with a Type 4 driver and running DB2 version 10.

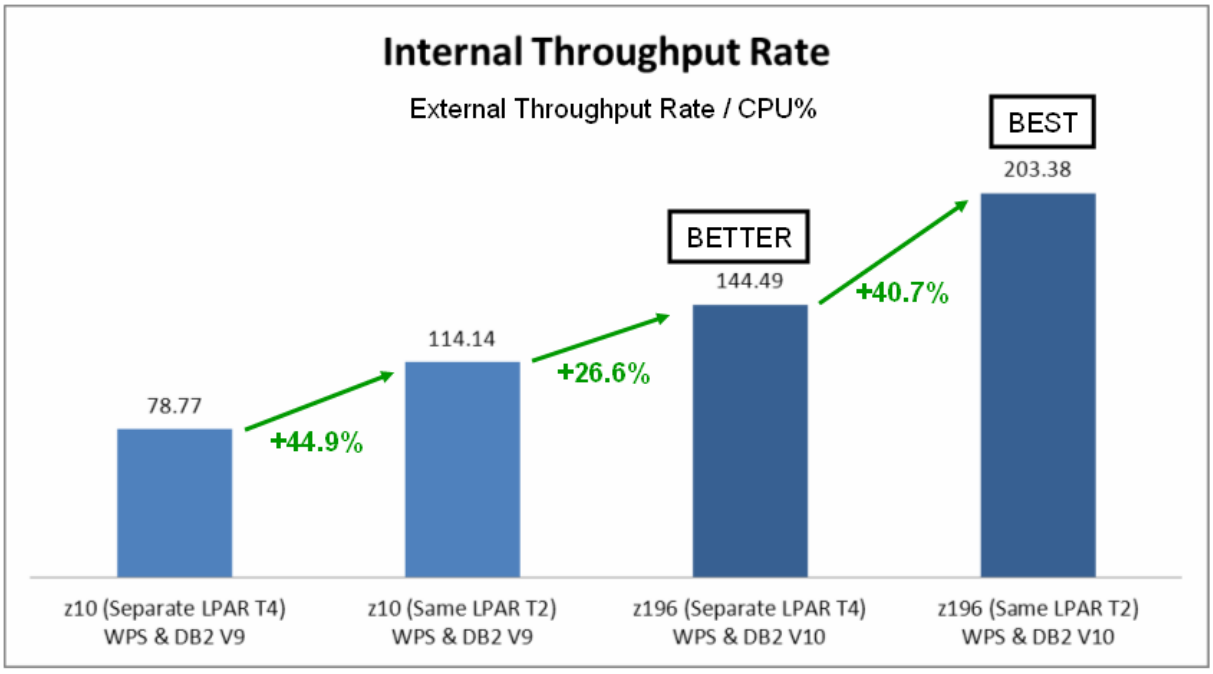

Figure 4 – Performance Improvements

Where did these improvements come from? Well over 80% of the improvement is derived from running on a z196 system rather than the z10 system and from the new levels of

WebSphere software. Running collocated workload in the same LPAR offers additional performance improvements. Collocated workload reduces the number of software instructions needed for transaction. Fewer instructions leads to more efficient execution of the workload.

From a software perspective, an improvement of 26% is attributed to running DB2 version 10 rather than version 9. The zEnterprise hardware has a faster processor, new processing capabilities for managing memory (improved levels of caching) and support for preparing instructions for execution prior to the actual processing (known as out-of-order pipeline design) that all add up to improved Java performance.

The significant improvements we have achieved in the z196 system demonstrate that we are not just cutting costs, although we have achieved significant gains in cost improvements, but that we continue to invest in the technical improvements that make the z196 an excellent platform.

Cost improvements are measured in terms of the cost of 10,000 (10K) transactions. In general, when CPU capacity goes up, the cost goes up correspondingly. We show you with the cost improvements we have made that a proportional incremental cost increase is not a given.

- The cost of 10K transactions while running a z10 system configured with a separate LPAR and Type 4 driver, running WPS and DB2 version 9, was calculated to be $3.16.
- For the z10 system configured with WPS and DB2 version 9 running in the same LPAR with a Type 2 driver, the cost was calculated to be $2.73 per 10K transactions, a **13.6% decrease in cost** over the z10 system with a Type 4 driver.
- The cost of 10K transactions while running a z196 system configured with a separate LPAR and Type 4 driver, running WPS and DB2 version 10, was calculated to be $2.33, a **17.2% improvement** over the z/10 system with a Type 2 driver and running DB2 version 9.
- For the z196 system configured with WPS and DB2 version 10 running in the same LPAR with a Type 2 driver, the cost was calculated to be $2.18 per 10K transactions, a **6.5% improvement** over the z196 system with a Type 4 driver and running DB2 version 10.

Cost has many components. In this study we focus on hardware and software but it is worth pointing out that labor costs, power costs, systems management, disaster recovery plans and so on are not considered in this study. System z has strengths in these cost areas as well but is outside the scope of this paper to analyze all the costs of the environment.

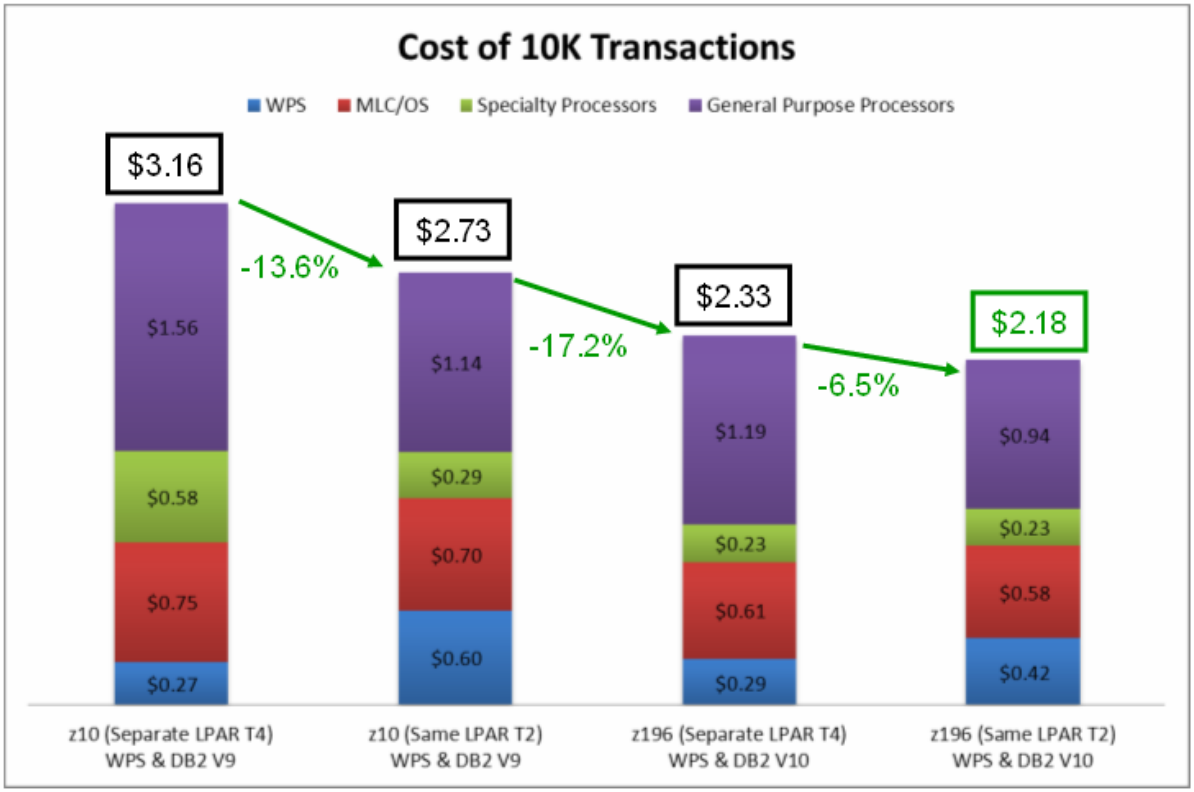

Figure 5 – Cost Improvements

Take time to study the last two columns of the graph in Figure 5. Notice that separate LPARs actually consume more general-purpose processing than a single LPAR configuration. Many customers try to minimize their software charges by creating separate, small LPARs. However, additional LPARs can have increased costs compared to a single LPAR, including additional CPU cycles to operate multiple LPARs.

For the WPS cost, the graph in Figure 5 shows a cost reduction from the z10 to z196 environments for the single LPAR use case. The greater efficiency of WPS, even with the scaling factor of 2x, has actually narrowed the gap in licensing costs.

Single LPAR price performance continues to be a focus area for IBM via innovations in hardware, software, and pricing. One example of a pricing innovation is Getting Started Sub-capacity Pricing (GSSP). (Explaining GSSP is beyond the scope of this paper; see Announcement letter 209-162, dated April 28, 2009 for more information about GSSP.)

It is worth noting that the savings gained from more efficient processing of transactions in a Type 2 configuration (for example, lower CPU costs and MLC) is contributing to the best overall price performance of $2.18 even though the WPS license is slightly higher than in the Type 4 configuration. As you can see, it is more important to look at the overall configuration that drives the cost of 10K transactions rather than the individual components of the configuration.

Let us further demonstrate cost improvements when making use of all the System z technology capabilities and software pricing techniques available by looking at the cost of 10K transactions in the same LPAR on z/OS (shown in Figure 6 as $2.18 for 10K transactions).

- Running WPS version 7 and DB2 on a z10 system using General Purpose Engines only, we calculated the **cost of 10K transactions to be $8.82.**
- Adding specialty engines to the z10 system **decreased that cost by 39% to $5.36 per 10K transactions.**
- Adding Getting Started Sub-capacity Pricing (**GSSP**[5]) to the z10 system **further decreased that cost by another 41% to $3.15 per 10K transactions.**
- Running DB2 version 10 and Integrated Workload Pricing (IWP[6]) on a z196 system noticeably **decreased that cost by a further 30% to $2.18 per 10K transactions!**

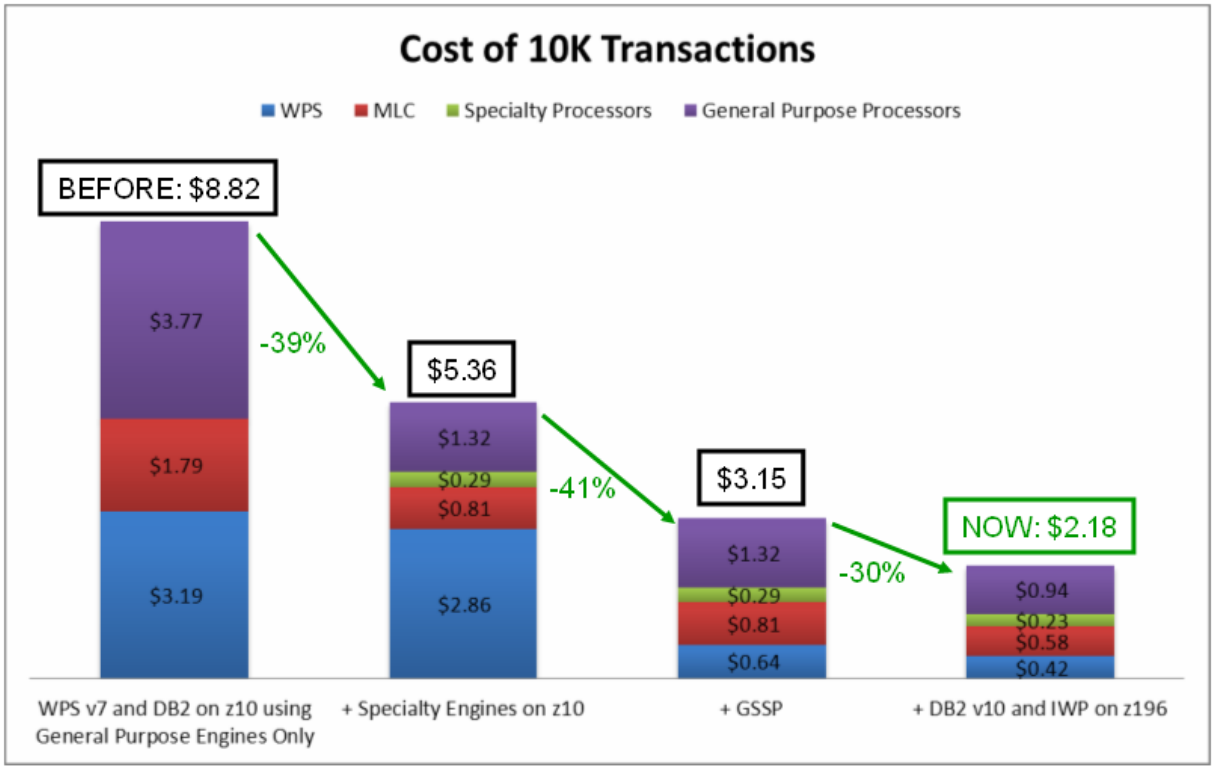

Figure 6 – Continuous Cost Improvements: Same LPAR on z/OS

As you can see, running in the same LPAR has not always been the most cost-effective use case. We have spent the last 3 to 4 years working to lower this cost, however. Factors that have contributed to cost include:

- Today's complex business transactions tend to be more CPU-intensive than traditional transactions.
- WebSphere running on z/OS is licensed on a per-LPAR basis, where the size of the LPAR tends to be large and increase year over year.
- Adding new workload on z/OS increases the size of the LPAR, which increases the license cost of existing software that runs in the same LPAR.

[5] For additional information about GSSP see announcement letter 209-162, dated April 28, 2009.

[6] For additional information about IWP see announcement letter 211-011, dated January 18, 2011.

These factors contributed to a cost of $8.82 for 10K transactions, as shown in Figure 5. With the introduction of System z Application Assisted Processors (zAAPs) and System z Integrated Information Processors (zIIPs) (meaning hardware that was less expensive and no software charge), IBM started to tackle these costs, leading to a new cost that was lower by 39%.

In 2009, IBM introduced Getting Started Sub-capacity Pricing (GSSP), which permits WebSphere workload to be priced based on the size of the **workload** instead of the size of the LPAR. New in 2011, IBM introduced Integrated Workload Pricing (IWP), which eliminates the increase in cost of the peer collocated transaction servers: CICS and IMS when running the WebSphere family of products on z/OS. This latest pricing innovation, available for zEnterprise 196 customers, can improve the overall software price performance and enhance the value of collocation on z/OS.

The combination of these pricing innovations and technical advancements in the hardware and software resulted in an **improvement in cost of over 4X**.

The Operational Value of Collocation of New Workloads on z/OS

Deploying WebSphere workloads on z/OS in a single LPAR with a type 2 driver offers the advantages of higher performance and simplified workload management and control.

With the type 2 driver, local cross memory communications between applications is faster. The overhead of using data communication encryption (such as SSL) is reduced to zero when everything is collocated in a single LPAR because there is no need for security encryption. When the WebSphere workload is collocated, the new thread of execution that is normally required when remote calls are received by the data system on z/OS is not needed; the thread of execution is carried through, reducing the switching overhead. The combination of these three things speaks to one thing: higher performance.

Better workload management and control is achieved through a variety of solutions, such as a single, end-to-end Workload Manager (WLM) classification. Local calls enable WLM to manage the entire request flow under a single classification, which increases the manageability and control over remote calls. Flowing security over a remote connection often requires that you define a remote alias, which then requires that you create ID and password definitions across multiple servers, which can make updates and coordination more difficult to control. Not needing to define identity aliases outside a central repository means more simplification, better control. Finally, through the z/OS Resource Recovery Services, z/OS provides for highly efficient synchronization point coordination, particularly during transaction recovery, unlike remote calls, which often require distributed, two-phase commit protocols, such as Java Database Connectivity (JDBC) XA.

Additional information about the technical advantages and the resulting business value of WAS on z/OS is available in the *Why WebSphere Application Server for z/OS?* Presentation (Document ID WP101532; available on http://www.ibm.com/support/techdocs).

Summary

Customers demand that the dynamic and complex transactions that they are running in their businesses be executed and managed in the most secure and reliable infrastructure. No other combination of operating system and hardware can match the z/OS operating system running on a z/Enterprise System z196.

A confluence of the latest middleware, operating system, and hardware architecture of the System z platform come together to meet the most stringent business and technical requirements and are proven to provide the best value to the customer in terms of transaction efficiency and cost. Are your customers taking advantage of the benefits of the System z hardware innovations and pricing options discussed in this paper?

For customers who are already running workloads on System z, we have demonstrated that a WebSphere workload collocated in the same LPAR with a Type 2 driver offers the best option from both a performance and a price perspective. Just look at the numbers:

- **78% performance improvement** over our previously published (July 2010) figure (based on System 10, same LPAR, Type 2 driver with WPS and DB2 V9).
- **20% cost improvement** over our previously published (July 2010) figure.

Customers are encouraged to consult their local IBM Sales representative to learn how these new technical and pricing innovations can help their business.

References

"The Value of Co-Location: Locating WebSphere Application Server in the same z/OS instance as data" and the update to this paper, published June 2011.	Find this document on the web at http://www.ibm.com/support/techdocs Search for document number WP101476 in the category of White Papers.
"Top-10 Reasons the mainframe is the Most Cost-Efficient Platform Available Today"	Go to http://www.mainframezone.com/it-management/top-10-reasons-the-mainframe-is-the-most-cost-efficient-platform-avail/P5 By Jonathan Adams for MainframeZone.com, May 17, 2010

Addendum: Linux on System z Comparison

Significant improvements in the performance and price of the newer releases of System z® hardware and software led to the publication of the WebSphere® for System z Prescriptive Use Cases white paper in June of this year. Testing and analysis of the Linux for System z option had not been completed at that time. However, the story we had was so compelling that we decided to publish the z/OS® data that was available, with the understanding that we would re-publish the white paper when the Linux on System z data became available and had been analyzed. That time is now.

The most recent configurations that we focused on:

- WPS V7.0 CICS® V3.2 and DB2® V10 collocated on z/OS
- WPS V7.0 on Linux on System z accessing CICS® V3.2 and DB2® V10 on z/OS.

For the z/OS collocated configuration, we used the same configuration that we had used for the June testing.

For the Linux on System z in a separate LPAR configuration:

- The hardware consisted of a z196 2817 in LPAR mode with:
    - 2 general purpose and 1 zIIP processor (for the z/OS DB2 system)
    - 2 IFL processors (for the z/VM® and Linux system).
- Software included:
    - z/VM Version 5 Release 4.0, service level 0901 (64-bit)
    - SUSE Linux Enterprise Server 10 (s390x) SP3
    - WAS and WPS on Linux were at the same levels that were used for the June testing:
        - WPS for z/OS V7.0.0.2 64-bit
        - WAS for z/OS V7.0.0.11 64-bit.

More recent versions of some of these products (including BPM Version 7.5 and CICS TS for z/OS Version 4.2) became available after the June paper was published; however, we maintained the same product levels for the purpose of an “apples-to-apples” comparison.

Linux on System z in a Separate LPAR

The configuration in Figure A1 on page 20 shows WPS V7.0 running in a new Linux on System z LPAR. CICS Transaction Server for z/OS V3.2 and DB2 for z/OS V10 are running in a separate z/OS LPAR. CICS Transaction Gateway V7.0 is used to access remote CICS transactions. A JDBC Type 4 driver is used to access DB2.

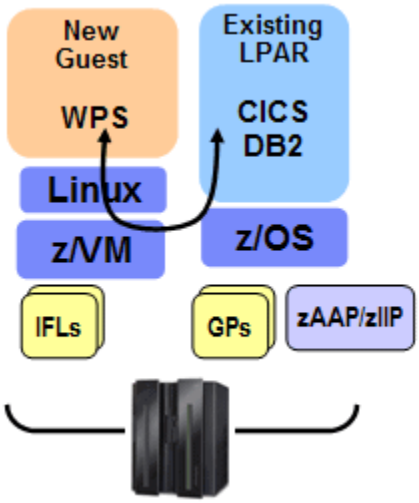

Figure A1 – Linux on System z in a separate LPAR

Collocated Configuration

The collocated configuration, seen in Figure A2, shows WPS V7.0 running in the same LPAR with CICS TS and DB2. Again, CICS CTG is used to access CICS transactions. In this configuration, a JDBC Local Type 2 adapter is used for DB2 access.

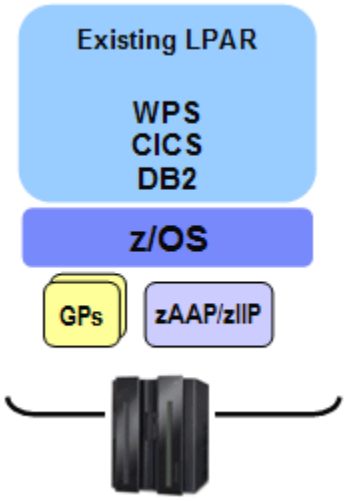

Figure A2 – Same LPAR Collocated

Performance Comparison

As in the previous testing earlier in 2011, performance was measured in terms of internal throughput rate (ITR), the number of transactions per CPU busy second.

- Separate LPARS: For the WPS V7.0 on a Linux on System z LPAR accessing DB2 V10 on a z/OS LPAR, the ITR is **157.02**.
- Collocated configuration: For WPS V7.0 collocated with CICS TS 3.2 and DB2 V10 on a single z/OS LPAR, the ITR is **203.38*, a performance boost of 30%!***

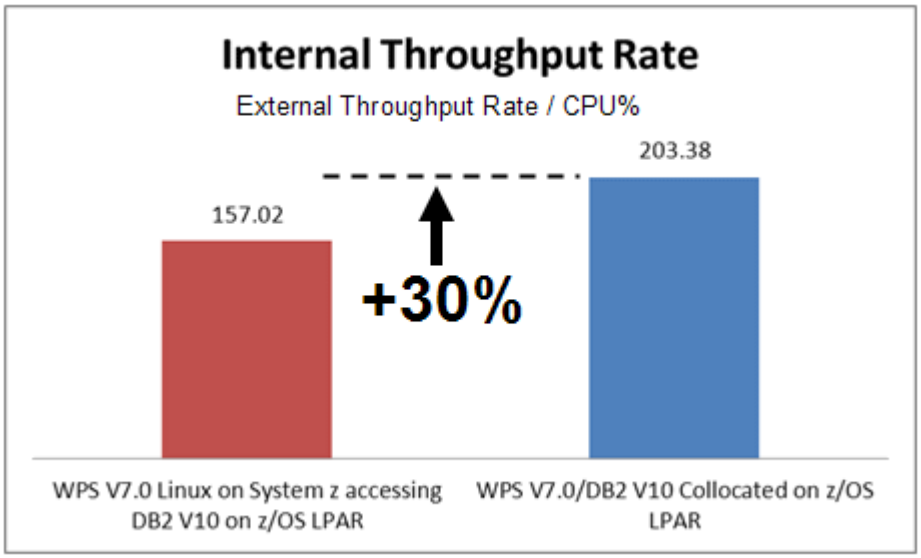

Figure A3 – Internal Throughput Rate Comparison

A second performance measurement was taken to compare the end user response time for each configuration.

- Separate LPARS: For the WPS V7.0 on a Linux on System z LPAR accessing DB2 V10 on a z/OS LPAR, the end user response time is **0.0281 seconds**.
- Collocated configuration: For WPS V7.0 collocated with CICS TS 3.2 and DB2 V10 on a single z/OS LPAR, the end user response time is 0.0220 seconds and as shown in Figure A4 on page 22, this is ***an improvement of 22%!***

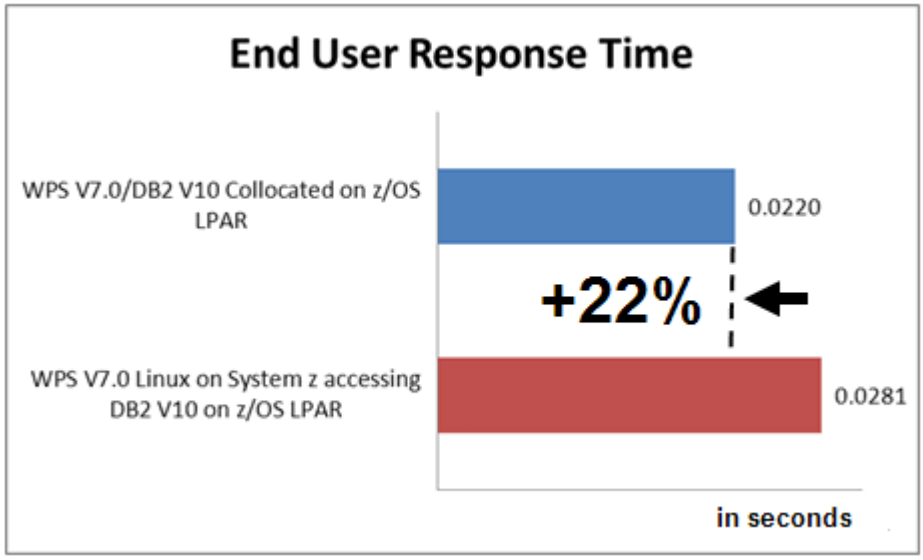

Figure A4 – End User Response Time

In both comparisons, running a workload in the collocation configuration proved to be the better option from a performance perspective. Why? We investigated further to find out two things:

- The source of the performance differences.
- The relative CPU usage with the increased workload.

The extra processing comes mainly from three sources:

- Operating system infrastructure: maintaining the environment under z/VM and the Linux operating system increases the number of software instructions that are needed for transactions.
- Network: sending transactions between two separate logical partitions using a Type-4 remote connection instead of a Type-2 local connection within one logical partition.
- WPS, DB2, and CTG: executing extra instructions to send and receive external transaction requests within and between the WPS, DB2, and CTG products.

A breakdown of the CPU usage is shown in Figure A5 on page 25. You can see that the CPU usage is significantly higher with every measurement of business services per hour, with the usage gap per engine increasing as the volume of business services per hour increases.

What is clear in each measurement of business services per hour is that the CPU usage per engine is higher for the Linux on System z configuration than for the z/OS collocated configuration.

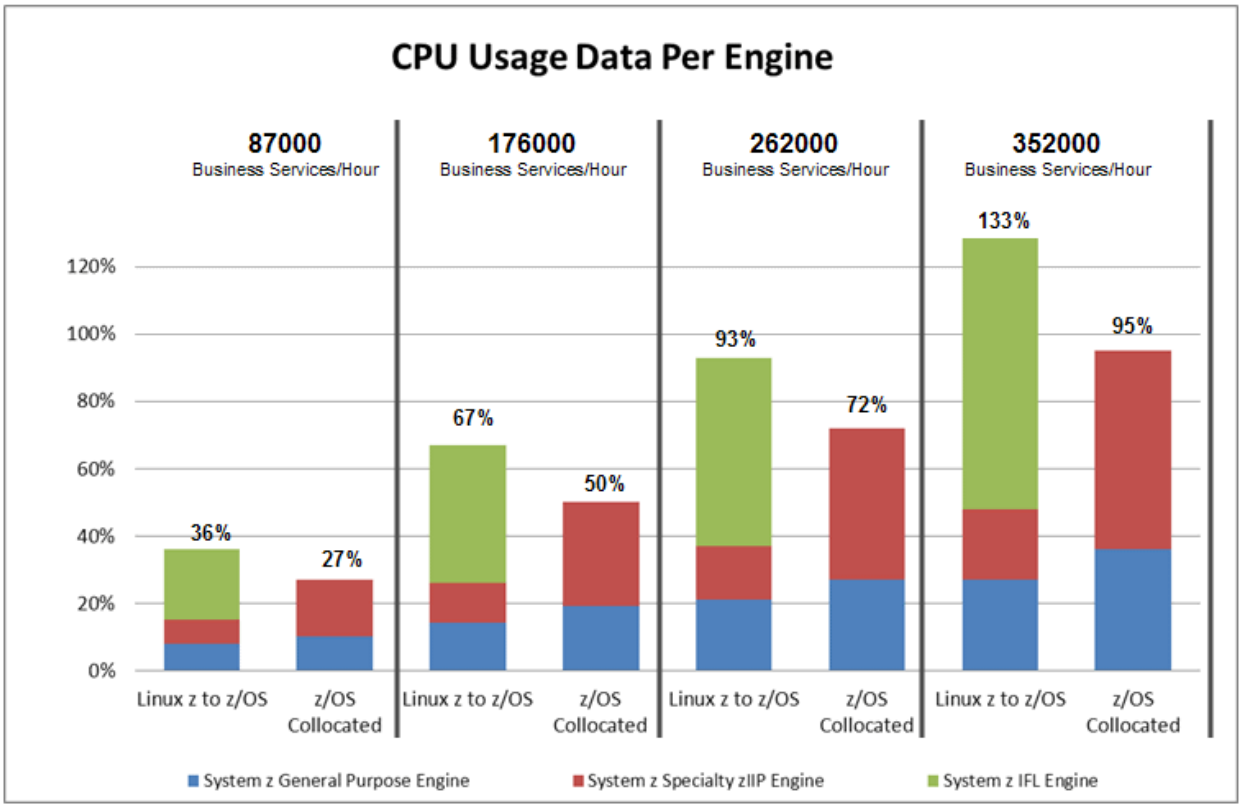

Figure A5 – CPU Usage Data Per Engine

The CPU usage is further analyzed by product in Figure A6 on page 24 to help understand what is driving the usage numbers. The data in this figure is based on 262, 000 business services per hour. WebSphere Process Server is the biggest CPU consumer, at 52.80% for the zLinux configuration.

Although WPS* remains the largest CPU consumer for the z/OS collocated configuration, the use of the Type-2 local connection and the collocation itself provide a benefit for this configuration of virtually no network time on z/OS. For this configuration, the CPU usage for WPS is 45.00%; **WPS uses 17% less CPU resource in the z/OS collocated configuration than in the Linux on System z separate LPAR configuration.**

* WebSphere Application Server is included with WPS for these measurements.

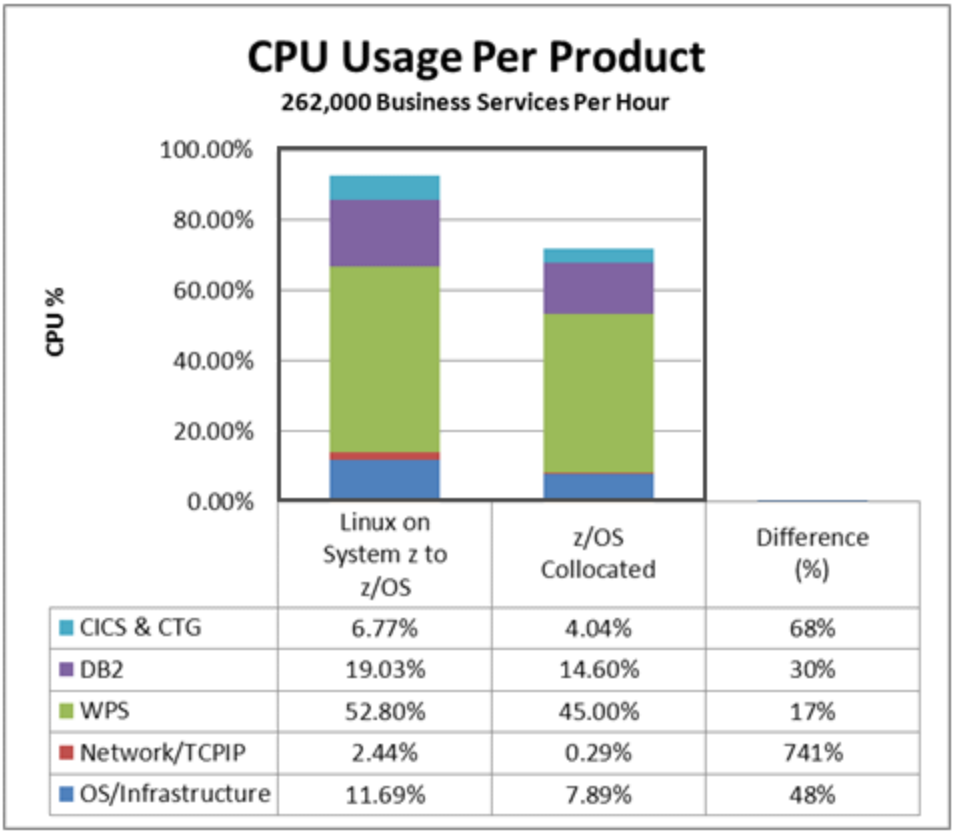

	Linux on System z to z/OS	z/OS Collocated	Difference (%)
CICS & CTG	6.77%	4.04%	68%
DB2	19.03%	14.60%	30%
WPS	52.80%	45.00%	17%
Network/TCPIP	2.44%	0.29%	741%
OS/Infrastructure	11.69%	7.89%	48%

Figure A6 – CPU Usage Per Product

How do CPU usage and workloads contribute to cost? We expected that not only would we see higher performance with the z/OS collocated configuration, we would also see a more efficient use of resources, which is a contributing factor to the total cost of a system.

Is the increased end user response time and more efficient CPU usage enough to offset the higher unit costs of individual components? The overall story is very encouraging for z/OS and might surprise some customers who traditionally have thought that Linux on System z is the better option for cost reasons. In our cost comparison, we looked at the average cost of 10K transactions and the total cost of acquisition as measured in business services per hour, to show how individual components contribute to the overall cost.

Cost Comparison

The June 2011 figures showed cost measurements in terms of the cost of 10,000 (10K) transactions. As shown in Figure A7 on page 25, we have taken the average cost of 10K transactions, where the cost of 10K transactions is measured for 100, 200, 300 and 400 users and then averaged. As seen in the June 2011 measurements, looking at the overall configuration that drives the cost of 10K transactions is more important than the individual components of the configuration.

With a difference of approximately 1% in the average cost of 10K transactions, you can easily see that the cost is fairly neutral between the compared configurations. Clearly, this is not a differentiating factor when you are looking for the best system solution for your workloads.

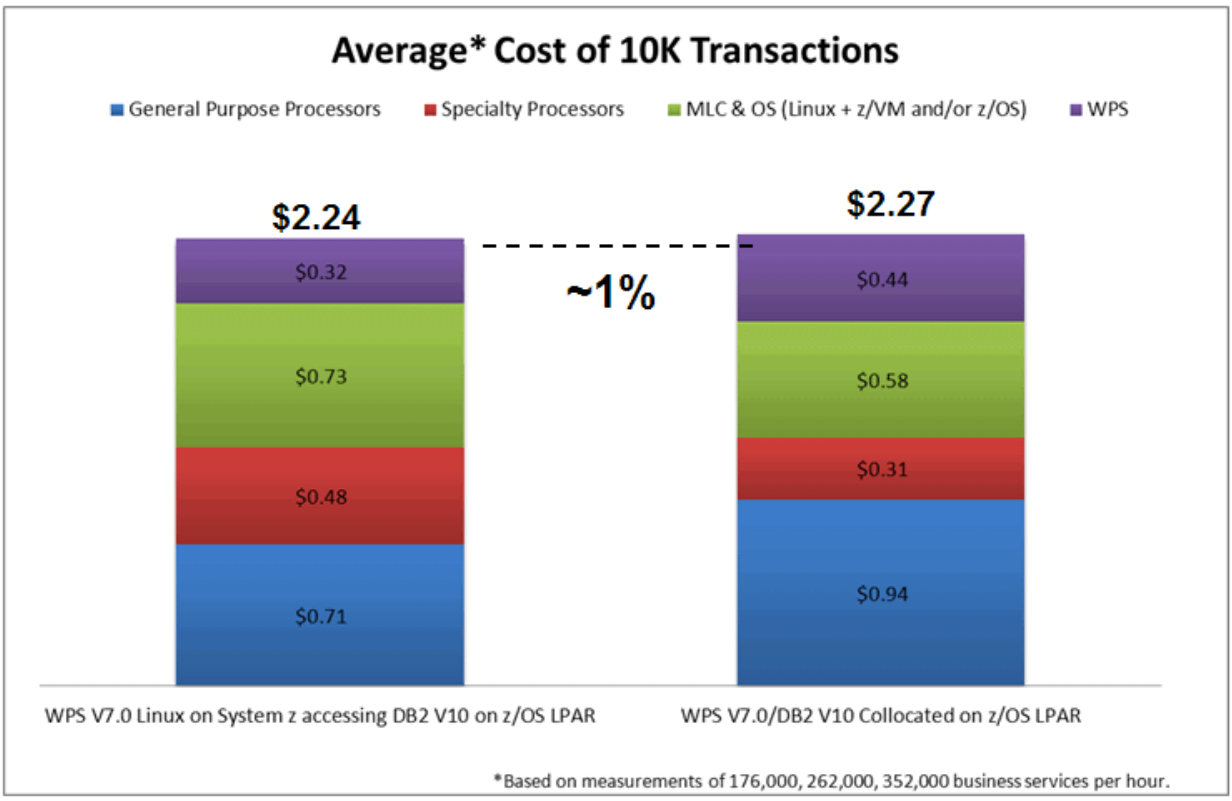

Figure A7 – Average Cost of 10,000 Transactions Comparison

Another way to look at cost is to compare the total cost of acquisition (TCA); in Figure A8 on page 26, we look at the 3-year TCA based on four quantities of business services per hour. What is immediately clear is that there is no obvious trend. Again, it is generally (but not always) the case that the collocated workload is less costly than the separate LPARs configuration. In most cases, the difference is relatively small.

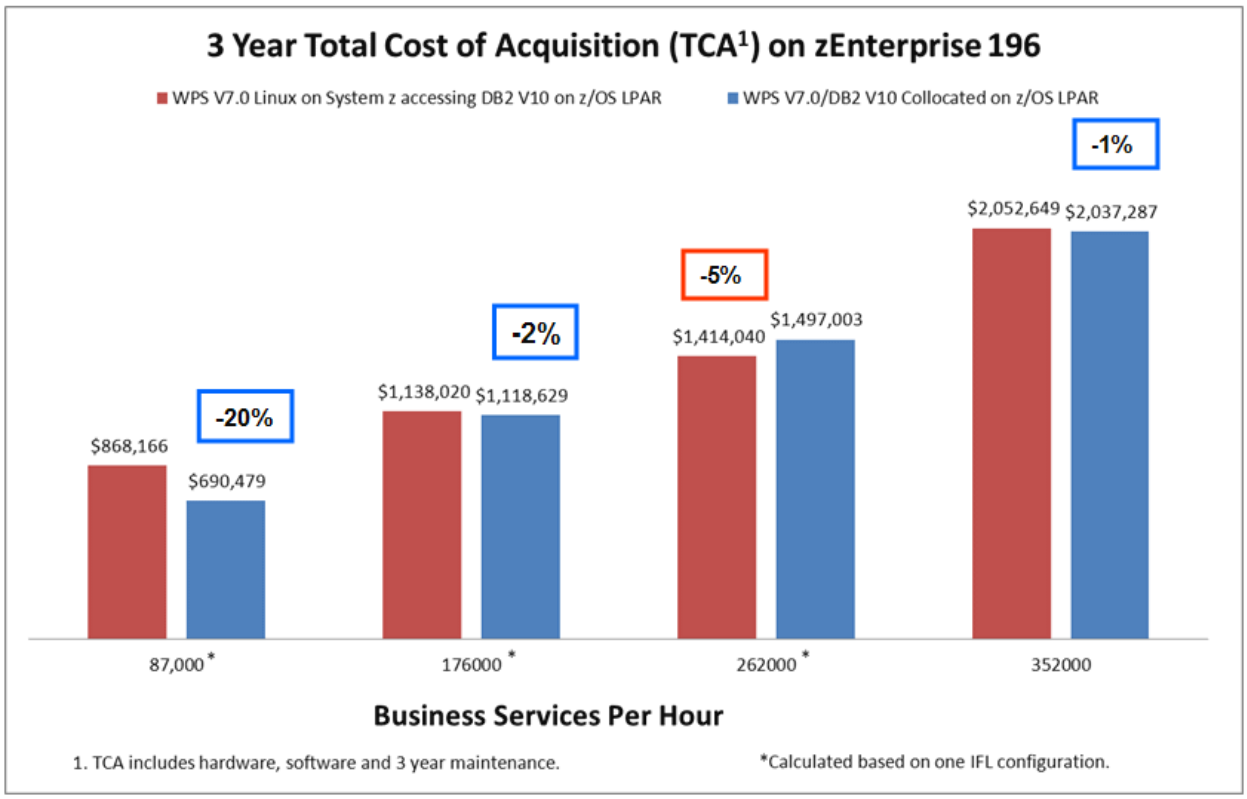

Figure A8 – 3-Year Total Cost of Acquisition Comparison

Note: The configuration in Figure A1 shows two IFLs. However, in Figure A8 the cost of the two IFLs is normalized to one IFL for 87K, 176K, and 262K business services per hour.

Conclusion to the Addendum

When you compare the price/performance of workloads on separate LPAR configurations versus collocated workloads, you see that there is a much greater difference between the configurations in terms of performance, with a much smaller difference in cost. **With a cost difference of approximately 1%, the 30% performance improvement weighs more heavily** in the decision of where to place different workloads in customer organizations.

Based on a clearly differentiated performance value, it is architecturally better to look to a collocated workload configuration rather than placing workloads in separate LPARs.

References

Optimizing process management with IBM System z, ZSW03193-USEN.

Appendix

About the WPS Application Used in This Effort

The WPS production workload is a cross-industry application that is centered on the activity of processing orders. The WPS production workload application is enabled for WAS z/OS and is representative of actual existing IMS, CICS, and DB2 workloads. This application was used to create a composite application using SOA principles with WPS for z/OS V7.

The WPS Order Processing System serves as the principal customer-like application and workload for validating new releases of WAS z/OS. It is also used cooperatively by many IBM internal development teams to drive z/OS workloads and benchmarks on System z hardware.

The WPS workload is a typical example of the mission-critical online processes for an IBM customer. What is important about the workload is:

- The variation of the Business Services
- The fact that Business Services are run concurrently from multiple users
- The ability to drive the workload efficiently at multiple Business Services per hour rates.

Activities of the WPS workload

The WPS application includes typical activities that most companies conduct, such as:

- Entering a new order
- Processing payments
- Checking the status of the order
- Delivering products or services
- Checking the stock level
- Changing the price of an item
- Giving a price quote
- Performing a customer inquiry.

Differentiation Value for Deploying Java Workloads on z/OS

WPS/WAS z/OS differentiation and customer value occurs *below* the open specification line. How the differentiating value (the features, functions, attributes, and QoS) is realized depends on the platform where it is deployed. Some notable examples of WPS/WAS active exploitation of System z and z/OS are:

- **z/OS Workload Manager (WLM) and WebSphere Application Server for z/OS Controller/Servant architecture** provide intelligent dynamic capacity expansion, flow control, and routing. This combination provides ***availability and scalability*** like no other WebSphere platform and allows WAS on z/OS to scale with the multiprocessor capability of System z.

- **z/OS Resource Recovery Services (RRS) and two-phase commit processing** offer superior ***reliability***. Distributed implementations can implement XA[7], but it is less efficient than RRS in a parallel sysplex environment. RRS is ready to exploit from the start.
- **EAL4+ Certification** through logical partitions (LPAR), cryptography for Clear Key and Secure Key, and z/OS Security Server, including RACF, provide ***multi-level security.***
- **WAS z/OS V7,** announced and available since 3Q 2008, introduced features (such as Fast Response Cache Accelerator [FRCA]) that offer significantly improved ***performance***. For workloads that can benefit from FRCA, improvements of 40 percent to 70 percent have been measured. In addition, enhancements such as thread hang recovery increase ***availability***. These features are offered *only* in WebSphere for z/OS.

[7] The XA specification defines how applications use a transaction manager to coordinate a distributed transaction across multiple resource managers. Distributed transactions enable independent resources to participate in a single unit of work that assures data integrity, even if one resource cannot commit or crashes while the transaction is in progress. A transaction will commit only if all resources can commit successfully, otherwise it is rolled back.

IBM Corporation
Produced in the United States of America
October 2011