

Implementation Design Study WebSphere Application Server

Prepared for:
IBM

Sine Nomine Associates
43596 Blacksmith Square
Ashburn, VA 20147

March 14, 2007

Copyright © 2007 Sine Nomine Associates. All rights reserved.

Document History

<i>Date</i>	<i>Revision/ Reissue</i>	<i>Nature of Change</i>
March 14, 2007	Version 1.0	Initial Public Release

Table of Contents

1	INTRODUCTION AND PREFACE	1
1.1	Executive Summary	1
1.2	Intended Audience	1
1.3	Preface to Discussion	1
2	SCENARIO DISCUSSION	2
2.1	Deployment and Uses of WebSphere	2
2.2	Business and Technical Objectives of Deployment	2
2.3	Study Structure and Analysis Scope	3
3	DISTRIBUTED SERVER DEPLOYMENT.....	4
3.1	Deployment Architecture and Design Goals.....	4
3.2	Design Overview.....	4
3.3	Component List.....	6
3.4	Analysis	7
4	VIRTUALIZED SERVER DEPLOYMENT IN SYSTEM Z9.....	8
4.1	Deployment Architecture and Design Goals.....	8
4.2	Design Overview.....	8
4.3	Component List.....	9
4.4	Analysis	10
5	STUDY RESULTS.....	12
5.1	Architecture and Design Goals	12
5.2	Operations Simplification and Business Resilience	12
5.3	Security Auditing and Compliance Management	12
5.4	Cost of Acquisition	12
5.5	Cost of Operations and Maintenance	13

5.6 Notes and Conclusions13

List of Tables

Table 1: Component List for Discrete Server Solution.....	6
Table 2: Component List for Virtualized Solution on System z	10

List of Figures

Figure 1: Block Diagram of Network/Systems Module	4
Figure 2: Module Internal Structure	5
Figure 3: Proposed Virtualized Architecture on System z.....	8

1 Introduction and Preface

1.1 Executive Summary

This report is part of a series that examines common designs in the discrete server environment and presents the equivalent design using virtual machines for System z9 and Linux for System z. The paper presents the original design for a WebSphere environment, the System z9 alternative design and a discussion of the pros and cons of the alternate design from a technical and business viewpoint.

1.2 Intended Audience

This paper is targeted to decision makers and technical staff responsible for evaluating platform hosting decisions. Some technical familiarity with infrastructure design principles is assumed.

1.3 Preface to Discussion

The discrete system design in this paper is drawn from an actual deployment for a Canadian bank performed in 2003. The original design goals are retained, however pricing and systems availability are adjusted to reflect the availability and pricing model changes as currently available. Readers should note that any pricing information cited here is list pricing and does not reflect any discounts or other incentives that may be available from IBM or other distributors for specific customers or situations.

2 Scenario Discussion

2.1 Deployment and Uses of WebSphere

WebSphere Application Server and other WebSphere component products are employed to speed development and deployment of customer applications. The WebSphere software suite accomplishes this by providing a rich library of pre-defined routines that handle common application programming tasks. This allows customers to concentrate on the business and presentation logic of the application without needing to actively consider maintaining a consistent user interface or integration with a larger application infrastructure. WebSphere provides the platform on which the application logic runs, without the application logic needing to worry about how it is interacting with the wider computing environment. The customer project considered in this report explored the use of different technologies on the SPARC and Intel platforms at a time when some WebSphere components were not yet available. As a result, the customer chose to deploy using the platforms available at that time.

2.2 Business and Technical Objectives of Deployment

The overall goal of the project was to integrate the presentation of customer information into a centralized portal application, allowing bank customers to perform a wide range of housekeeping and personal-finance related tasks from a single location. The bank, like most institutions, had a number of internal systems that had evolved at various times to provide support for the institutional and regulatory services required to support various financial products, but all were developed at different times by different groups, and thus presented different images for different products.

The business goals of the project were to provide:

- A unified presentation of services including a common “look and feel” for the products and services of the bank
- Self-service access to common customer data requests and updates, such as current balance information or mailing address changes
- Minimized development, testing and integration costs for new products and services.

The technical challenges of the deployment were legion: a large number of pre-existing applications (including some commercial applications without access to source code), a combination of SNA and TCP/IP connectivity, a stringent, regulated environment with complex encryption and authentication requirements, and a number of reporting requirements to internal and external agencies. These requirements were also layered on a mandate for more extensive testing and auditing of the applications before public deployment, including a regression test environment that could be used to test incremental upgrades in a realistic copy of the production environment (in order to exercise interconnections between applications and possible interactions in resource utilization between multiple applications sharing a application server).

2.3 Study Structure and Analysis Scope

This study examines the project in light of the advances in IBM System z software and hardware technology and presents an alternative approach that addresses the same requirements at lower cost and complexity than the currently deployed discrete system solution. The alternate design is in evaluation at the customer site as a possible technology upgrade.

3 Distributed Server Deployment

3.1 Deployment Architecture and Design Goals

The original WebSphere architecture was engineered using eight discrete Intel-based servers running Windows 2000 Server and two, SPARC-based, Sun 4500 systems running Solaris. The customer was familiar with both technologies from internal development and support for the other discrete applications within the environment and felt confident with the supportability of the Sun/Intel-based approach.

3.2 Design Overview

The design was originally intended to be replicated in the bank's secondary data center, and a parallel test configuration was provisioned in a lab site in Alberta near the development teams (due to implementation cost overruns, the additional production deployment was never completed). Use of "standard" enterprise hardware and networking infrastructure were significant issues. Previous projects had often left hardware selection to individual application groups, resulting in difficult-to-manage, one-off configurations. Use of Cisco networking hardware and IBM or HP Intel systems, along with the Sun SPARC hardware, was required to attain uniform manageability and easy availability at business continuity facilities. The configuration was deployed in the primary data center and connected to the bank's core network infrastructure as a modular element (network and systems integrated as a modular unit). Figure 1 below illustrates a block diagram of the design and the interconnected systems with the bank's environment. The solution component representing this project is the circled component.

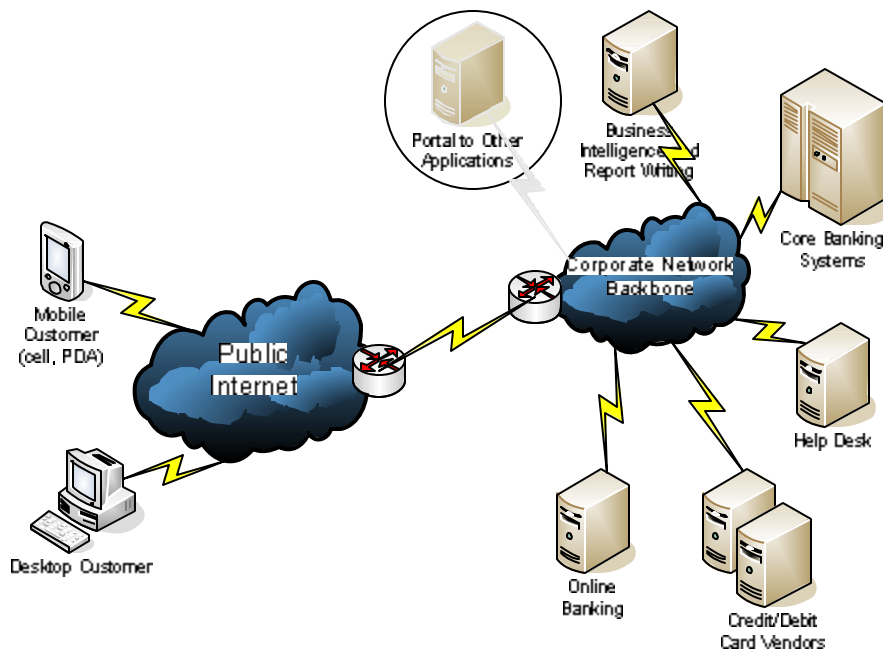



Figure 1: Block Diagram of Network/Systems Module

The internals of the module can be illustrated as shown below in Figure 2:



3.3 Component List

Hardware

Description	Number	Cost (List)	Subtotal
<i>Server Hardware</i>			
HP DL360 rackmount server (prod)	20	\$ 2,978.67	\$ 59,573.40
HP DL360 rackmount server (test/QA)	2	\$ 2,978.67	\$ 5,957.34
HP DL360 rackmount server (dev)	2	\$ 2,978.67	\$ 5,957.34
Sun E4500 8 way, 16G, 2 FCP HBA (prod)	4	\$ 82,965.72	\$ 331,862.88
Sun E4500 8 way, 16G, 2 FCP HBA (test/QA/dev)	1	\$ 82,965.72	\$ 82,965.72
Total Server Hardware			\$ 486,316.68
<i>Network Hardware</i>			
Cisco 4507 (chassis)	3	\$ 7,680.00	\$ 23,040.00
12 port GigE blade	4	\$ 3,487.00	\$ 13,948.00
24 port 10/100 blade (test/dev/QA)	6	\$ 2,654.00	\$ 15,924.00
48 port 10/100 blade	2	\$ 2,955.00	\$ 5,910.00
Cisco Content Switch (prod)	2	\$ 9,787.00	\$ 19,574.00
Cisco Content Switch (test/dev/QA)	1	\$ 9,787.00	\$ 9,787.00
Total Network Hardware			\$ 88,183.00
<i>Support Hardware</i>			
7 ft 19" racks	3	\$ 3,200.00	\$ 9,600.00
Network Cables	109	\$ 7.00	\$ 763.00
40A PDU connections	3	\$ 2,675.00	\$ 8,025.00
Total Support Hardware			\$ 18,388.00
Software			
<i>Operating Systems</i>			
Windows 2000 Server	24	\$ 876.00	\$ 21,024.00
Solaris 2.8	5 included		\$ -
<i>Application Middleware</i>			
Websphere Application Server	29	\$ 15,782.25	\$ 457,685.25
<i>Monitoring/Management Tools</i>			
Netbackup Client	29	\$ 87.00	\$ 2,523.00
Total Software			\$ 481,232.25
Acquisition Total			\$ 1,074,119.93

Table 1: Component List for Discrete Server Solution

In addition to the itemized costs listed above, there are other, exogenous costs to the configuration such as failover configuration costs; people/administration costs; non-functional costs such as authentication flows, audit complexity, multi-tiered logging, business resilience solutions, security, compliance, availability, development and support costs. The costs itemized here are merely those that can be directly specified for this analysis.

3.4 Analysis

While the overall deployment was successful, and the project met the technical and business goals, there were several issues that surfaced as the environment operated over a longer term:

- Capacity demand in the environment was highly cyclical, however the peaks were more than 50 times the nominal operational load, requiring investment in substantial infrastructure to support a high load needed for only a few weeks during the year. This cyclical demand required investment in additional WWW servers and application servers to accommodate the demand at peak, requiring the purchase of two additional Sun systems and 12 additional WWW servers with accompanying hardware maintenance and software licensing costs despite utilization of less than 15% for the majority of the year.
- This additional capacity requirement impacted space and physical connectivity requirements for the module to accommodate peak load. The additional 13 systems and related networking hardware **doubled** the floor and rack space requirements, and imposed two additional 40A power circuits, requiring a \$32,000 power distribution system upgrade.
- The licensing cost model for WebSphere is by physical system and number of processors and in the discrete environment, must reflect peak load capacity. The initial deployment plus the additional capacity necessary to service the peak capacity demand increased the software licensing costs substantially, and continues to increase each time additional capacity is added to the total deployment, even during times of reduced load.

These issues have lead to the bank considering an alternative method to manage the deployment of this application and several additional applications.

4 Virtualized Server Deployment with System z9

4.1 Deployment Architecture and Design Goals

The shift of the design to a virtualized deployment within a System z infrastructure was intended to minimize the number and type of physical servers involved in the operation of the portal, and to optimize the pricing and delivery structure of the WebSphere product lines for this type of application where there was a strong tie to mainframe-hosted data as well as distributed data.

4.2 Design Overview

As shown in Figure 3, the application deployment logic is similar to the discrete system deployment (leveraging the existing application engineering and management tooling). However, all of the application server and web server infrastructure has been moved into virtual machines hosted by a z/VM system running on the existing mainframe hardware. An additional two processors were activated on the existing mainframe system to support the additional workload under the control of the System z Capacity on Demand microcode, and additional network adapters were deployed to provide capacity for data paths to the networking hardware.

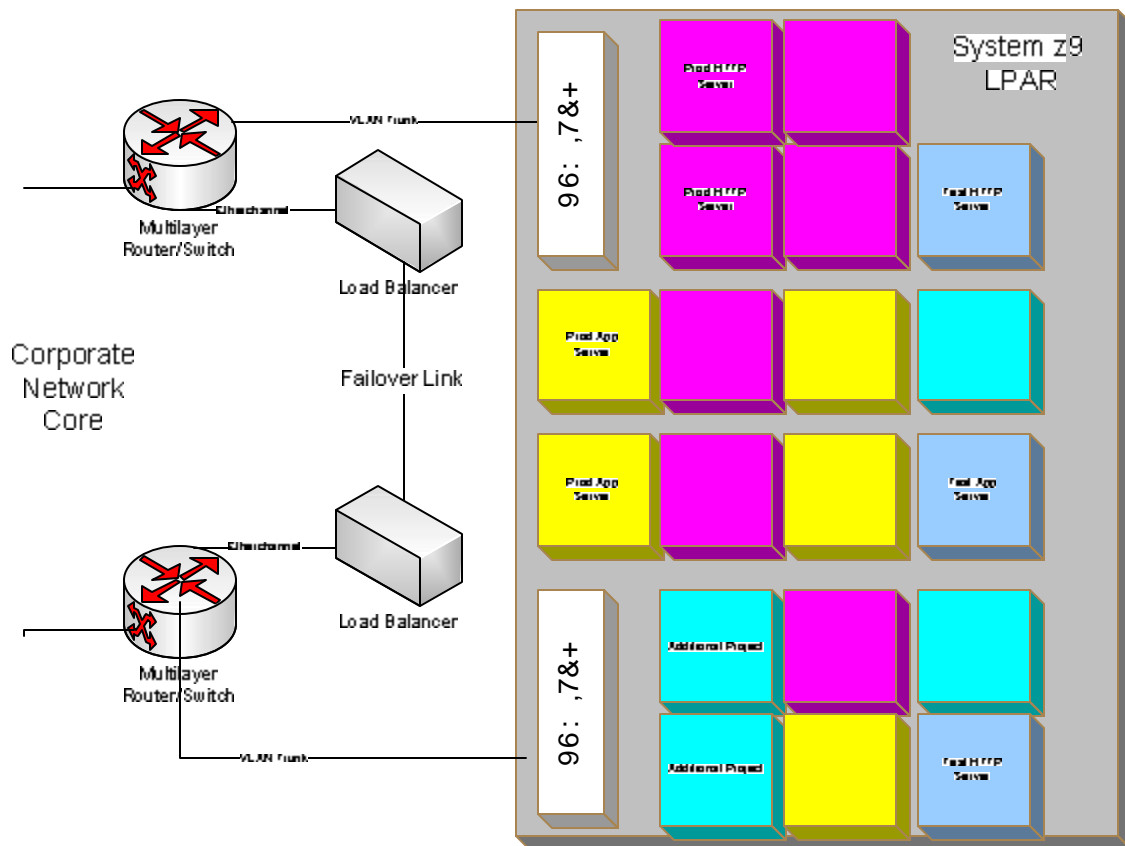


Figure 3: Proposed Virtualized Architecture on System z

Where possible, the external networking hardware and load balancers were retained. This application of existing hardware provided standardized IP traffic management and traffic shaping capability already present in the existing network hardware in a way that did not impact transaction performance or overall mainframe capacity, and also to allow additional future applications to be layered onto the same basic environment by extensive use of IEEE 802.1q VLAN trunking functions. Two additional 1000baseT network adapters were added to the configuration to provide additional connectivity options for the hosting environment.

The z/VM environment available on the IBM System z9 is also able to connect to the existing disk hardware, allowing virtualization of the SAN infrastructure with complete isolation for each server by utilizing the N-port FCP virtualization functions of z/VM 5.2 to arbitrate access to the physical FCP fabric.

4.3 Component List

As much of the networking equipment was recycled from the original deployment, the additional capital acquisition expenditure necessary for the upgrade concentrated on these areas:

- Acquiring the additional capacity in the existing mainframe
- Enabling the virtualization components
- Optimization of software licensing contracts to reflect the new environment.

Additional applications deployed into this environment would not incur any additional physical system installation costs (up to the physical computational capacity of the installed mainframe processor), however permanently enabling additional capacity from the reserve capacity available in the frame would incur additional expense to reflect the additional hardware capacity and increased per-processor licensing cost of the WebSphere middleware..

Hardware

Description	Number	Cost (List)	Subtotal	
<i>Server Hardware</i>				
Enable IFL on System z9	1	\$ 125,000.00	\$ 125,000.00	
Activate CoD on z9	1	\$ 36,785.00	\$ 36,785.00	
Total Server Hardware				\$ 161,785.00
<i>Network Hardware</i>				
Cisco Content Switch (prod)	2	\$ 9,787.00	\$ 19,574.00	
Cisco Content Switch (test/dev/QA)	1	\$ 9,787.00	\$ 9,787.00	
Total Network Hardware				\$ 29,361.00
<i>Support Hardware</i>				
7 ft 19" racks	1	\$ 3,200.00	\$ 3,200.00	
Network Cables	4	\$ 7.00	\$ 28.00	
40A PDU connections	0	\$ 2,675.00	\$ -	
Total Support Hardware				\$ 3,228.00
Software				
<i>Operating Systems</i>				
z/VM	1	\$ 22,500.00	\$ 22,500.00	
SuSE Linux for zSeries	1	5125	\$ 5,125.00	
<i>Application Middleware</i>				
WebSphere Application Server	2	\$ 15,782.25	\$ 31,564.50	
<i>Monitoring/Management Tools</i>				
Netbackup Client	30	\$ 87.00	\$ 2,610.00	
Total Software				\$ 61,799.50
Acquisition Total				\$ 256,173.50

Table 2: Component List for Virtualized Solution on System z**4.4 Analysis**

By taking advantage of Capacity on Demand, activating the application in a z/VM-hosted environment allows usage-based resource allocation on a pay-as-you-go basis for the 20 to 25 days per year where peak capacity is required. The design is intended to be elastic enough to support additional application deployment based on actual demand for resources, not the projected peak demands that would require over-provisioning and the accompanying maintenance costs for idle capacity at times of low demand. Use of virtual machines and the ability to impose strict resource allocation and consumption controls on test virtual machines allowed trivial duplication of the production environment for testing and development purposes, and eliminated the requirement for separate test infrastructure at the Alberta development laboratory and the accompanying connectivity to the bank's primary data center site.

Retaining the networking equipment allows unified management of the networking components with other parts of the enterprise network, as well as allowing the System z processor cycles to be concentrated on the business transaction logic rather than utility functions. Leveraging the existing management tooling and infrastructure allowed rapid deployment of the enhanced

functionality without requiring substantial re-engineering of the network design or infrastructure at large, thus reducing the cost of management.

From an auditing perspective, the preservation of the firewall and security functionality allowed minimization of recertification for the application, as the security management and audit trail functions remained the same from the old application environment.

5 Study Results

5.1 Architecture and Design Goals

While both approaches achieve the business and technical goals set out in the scenario, and both provide supportable models for deploying WebSphere-based applications, the virtualized model concentrates the management and licensing costs into a fewer number of systems, and allows resources to be trivially reallocated on an actual demand-based model. The fewer of systems involved consume less power and space, and over time provide higher levels of stability and utilization per unit cost.

5.2 Operations Simplification and Business Resilience

The virtualized environment reduces the operations management costs from more than twenty systems to two console screens. Resource management and performance are managed as part of the overall resource measurement handled by the Capacity on Demand microcode.

From a business resilience standpoint, the virtualized environment participates in the overall mainframe recovery plan, allowing the entire service infrastructure – servers, application and data – to be restored in a single step, and activated at the same time the mainframe backend is restored. This approach provides substantial improvements in time to availability for a critical customer-facing application (in some cases reducing an entire day of hardware installation, software configuration and data recovery to minutes). The configuration also allows leveraging parallel restores of data volumes and real-time replication of the data in the environment using existing copy sets already maintained by the existing mainframe applications and software.

5.3 Security Auditing and Compliance Management

As noted above, the auditing and compliance management process can be identical to the distributed design, leading to no change and no additional workload to adopt the virtualized model. If additional optimization is permitted, auditing and logging information can be redirected to facilities managed by the z/VM hypervisor, placing them outside the reach of an individual guest operating system and consolidating the audit information into a single common logging and performance management collection tool. This modification is simple to implement (console logs configured to be written to the guest system consoles and captured using the Single Console Image Facility code included in z/VM), but is not required by the transition to the virtualized environment, reducing the training and management tooling impact of implementing the System z-based virtual environment.

5.4 Cost of Acquisition

The cost of continued expansion of the discrete server implementation drove the creation of this study, however the cost of acquisition for the virtualized environment is still a significant factor. A design requirement for a replacement was to incorporate as much of the existing equipment as possible to limit the startup cost. The virtualized design involves activation of some of the pre-installed (but administratively disabled) capacity in the existing mainframe, and adding additional network connectivity to support the additional traffic to the portal application. The initial startup

costs (at list price) are estimated to be \$125,000 for an additional CPU and \$18,000 for two 1000baseT adapters.

While this may seem to be a large startup cost, the amount of capacity activated is controlled by the Capacity on Demand microcode in the System z processor, from a minimum of one processor to a maximum of four. The bank is billed regularly for the software and hardware costs to support the minimum capacity dedicated, and pays for additional capacity temporarily activated to support sustained bursts of demand that exceed 24 hours in duration by one day increments. By structuring the acquisition of the necessary capacity based on actual demand, the initial costs are amortized across a number of applications, and the acquisition cost reflects demonstrated need directly – unused capacity is immediately deactivated and the software and hardware bills remain flat for a majority of the time. Additional workload funds additional capacity.

The virtualization design also reduces the licensing cost for the WebSphere middleware by a substantial number. In the discrete environment, each discrete system required a separate WebSphere license and separate software maintenance for each system. The virtualized environment on System z is licensed on a per-active-processor basis, reducing the required number of licenses for production support from 20 to 1 permanent license, and a small graduated daily amount when additional capacity is enabled during peak periods (estimated average usage indicates 2 licenses will be required for normal production use and to accommodate testing). The ability to host both production and test on the same hardware allows sharing of the licenses, eliminating a further 4 WebSphere licenses entirely – from 24 licenses to two (2) plus some demand-related fees. Based on a list price quote from a Canadian business partner, this will modify the bank's costs for WebSphere from CDN\$360,000/year to CDN\$30,000.

The implementation of the virtualization environment does require licensing z/VM and some management tools, however the cost of licensing z/VM, the IBM management tools, and support for z/VM and the tools is approximately \$45,000 per year – trivially paid from the license savings quoted above.

5.5 Cost of Operations and Maintenance

The virtualized environment integrates operations for the environment into the existing management tools for the mainframe, removing licensing costs for monitoring agents and additional network ports to provide a management LAN to the discrete hosts. The cost of hardware and software maintenance of the environment is also reduced due to the dynamic adjustment of capacity discussed in the previous section. Power and environmental costs are fixed within the maximum capacity of the existing mainframe system (the power and environmentals are the same regardless of load), leading to a 50% reduction in environmentals (removal of 100 sq ft of floor space requirement) and reduction of power requirements by 100 amps (5 20A circuits). Using the customer's notional operational cost per square foot computation of CDN \$48 per square foot per month, this reflects a CDN \$4800/month savings in operations cost, independent of power consumables costs (which were not directly reported).

5.6 Notes and Conclusions

While not all workloads are well suited to the virtualized environment on System z, based on the expansion and contraction of the load on this application due to cyclical demand, and the expansion of other WebSphere-based projects within this customer, we believe that the System z environment is an excellent value for this type of workload. The ability to directly gate the

expenditures related to hardware and software licensing on demand – and the ability to concretely record and report the utilization of these resources – ties into the current demand for more efficient operation and utilization of IT expenditures, and to reflect that utilization directly in value to the enterprise. The System z9 z/VM and Linux virtualized server environment provides an elegant way to effectively balance demand and cost incentives for WebSphere applications, combining a cost-friendly licensing model and a simplified management capability to speed both time to delivery and developer convenience at an attractive long-term operational price.