System z Implementation Study: Oracle Utility Server

Prepared for:
Gerald Araneo
STG Analyst Relations

By:
Sine Nomine Associates
43596 Blacksmith Square
Ashburn, VA 20147

IBM-2007-03651-E-03
February, 2007

Table of Contents

List of Tables

List of Figures

1 Introduction and Preface

"It is like the story of the three little pigs. You can build in brick, or you can build out of straw or sticks. If your build out of the straw and sticks of commodity systems, you are committing yourself to an ongoing litany of grooming, acquisition and replacement. Building with brick costs more, but you have a more permanent asset. The IT technology brick-equivalent is IBM's System z9." Anne McFarland – The Clipper Group

1.1 Executive Summary

As part of a larger series of papers on the position of mainframe technology in the current IT market, this paper considers the role application consolidation plays in the decision process for organizations that may not be traditional mainframe environments, but are now considering System z as part of larger strategies within an increasingly compliance-oriented marketplace.

1.2 Intended Audience

This white paper illustrates the strengths of the System z platform for pursuing a server consolidation strategy. It is intended primarily for data center technical managers, architects and Oracle database administrators in small and mid-sized environments who are concerned with realizing the benefits of server consolidation. The audience to whom the paper is primarily addressed is expected to benefit most to either the IBM System z9 Business Class (z9 BC), which is targeted specifically to the needs of customers who may not have previously considered their workloads appropriate to the mainframe; and to users of IBM System z9 Enterprise Class (z9 EC) who are wishing to exploit their existing back-end processes with the servers they are consolidating.

1.3 Preface to Discussion

If the 1990's could be characterized by the rise of discrete servers, departmental servers and client server that may be best summarized by the catch phrases “I'll do it myself,” then the 2000's have given birth to the slogan “Do more with less.” The technology that we now associate with this cry is server consolidation.

2 Scenario Discussion

2.1 When to Consolidate

Consolidation is not for every IT infrastructure. Candidate components for consolidation are best identified on a case by case basis. Typically, it is found that the best candidates for consolidation exhibit one of these two characteristics:

1. Low average CPU utilization: This can be a very subjective term, but for the purposes of consolidation on System z it should be remembered that System z can be driven at a sustained utilization of 85-90% quite comfortably, whereas competitive virtualization technologies start to experience problems at the 30-35% mark. This means it is possible to consolidate the workload of more servers with System z than with other platforms.

2. Belonging to 2 of the 3 tiers in a traditional client/server model (typically the backend and middleware).

As a corollary, there are valid circumstances where server consolidation is not advisable:

1. CPU-intensive workloads

2. The required software will not run on the target platform

3. There is no expectation for the system to grow with the business

According to IBM in "Consolidating Oracle Databases to Linux on System z" [2], there are over 900 customers in production using Linux on System z. Linux is running on over 3000 IFLs across this customer base. Many of these customers report that apart from the savings from consolidation they are also enjoying the benefits of being able to offer "on demand" services or "capacity when needed."

2.2 Study Structure and Scope

In this white paper the traditional "discrete" approach to deploying Oracle servers is contrasted with that of the virtual server deployment using the scenario described in one of the seminal papers on server consolidation "Six Case Studies for Priority Investments in the Corporate Server Infrastructure" [1]. When first published, the zSeries (as it was then known) was identified as a platform for migrating Windows services such as file servers and print servers. In the subsequent years the System z has attracted a vast range of ISV support (978 applications for Linux on System z as of January 2007) that make consolidating Oracle instances a viable consideration.

IBM and Oracle have announced support for the Oracle E-Business Suite Database-Tier product on System z's Integrated Facility for Linux (IFL) in support of Oracle's line of enterprise business applications, such as PeopleSoft and Siebel. This also reinforces the belief that Linux on System z is a mature platform capable of running significant applications suites.

Currently, the available Oracle products for Linux on the System z are:

- Oracle9i Release 2 Enterprise Edition

- Oracle Database 10g Release 1 EE
- Oracle Database 10g Release 2 EE
- Oracle Application Server AS 10g
    - AS 10g 10.1.2 Base
    - AS 10g 10.1.3 (J2EE)
    - AS 10g 10.1.4 (Identity Manager)
- Oracle Clustered File System V2 (OCFS2)

The ability to consolidate the IT infrastructure was dependent on several technical advances.

1. Virtualization:
    - Hardware partitioning
    - Hypervisor software
2. Overcoming communication barriers:
    - Constrained WAN bandwidth
    - Proliferation of TCP/IP
    - Application chattiness multiplies the effect of latency

From a non-technical perspective, server consolidation was made viable as the financial side of the business came to understand that deployment of technology is not simply a function of acquiring some hardware and some software. When the total cost of operation (TCO) or total cost per user (TCU) was calculated, these initial acquisition costs were quickly dwarfed by the people and environment costs that it takes to operate and grow the infrastructure. Server consolidation then became an important cost and operations management strategy.

The major advantages of server consolidation have been enumerated in many places and in many different ways. Essentially, the benefits include:

1. Lower TCO/TCU via reduced environmental, people, and software license charges
2. Simplified management
3. Improved service levels

While server consolidation can be accomplished using many different platforms, the System z has several key advantages over the competition:

1. Fine granularity in capacity levels enables a business to choose an entry point into the processor family without grossly over or under provisioning. The business can grow and the System z will grow with the business.

2. A 40 year pedigree of virtualization and a toolset that manages a virtualized environment simply and efficiently.

3. A rock-solid hardware platform that has the reliability, availability and serviceability essential to operating the essential processes of every business.

4. A specialized engine: the Integrated Facility for Linux (IFL) created for server consolidation that lowers hardware costs and does not increase existing software licensing charges.

5. Synergy with existing mainframe-based systems such as z/OS using technologies such as hipersockets which provides for memory transfer speeds between systems.

6. A tool set that facilitates rapid cloning of servers.

7. Exploitation of disaster recovery facilities such as Geographically Dispersed Parallel Sysplex (GDPS) which has the ability to switch to and from production systems located in secondary data centers.

8. Capacity on demand means that unexpected or seasonal spikes in processing requirements can be met without downtime or additional provisioning.

9. Security: the System z has obtained EAL5 certification, which demonstrates that System z9 architecture is designed to prevent the flow of information among logical partitions on a single system.

For Oracle, the advantages of a Linux on System z configuration are primarily two-fold:

1. Reduction in the number of licenses. Oracle is licensed on a per IFL basis rather than per virtual server.

2. Ease of cloning new servers.

3 Distributed Server Deployment

3.1 Deployment Architecture and Design Goals

The configuration under study is a Solaris-based Web Portal implemented using 13 discrete servers performing the following tasks:

1. Web serving
2. Real time data feeds
3. JRUN serving
4. Java applet serving
5. C++ applet serving
6. An Oracle database
7. A File and XML server

For contingency reasons this configuration is replicated across three sites. In addition there are additional servers used for development, QA, test, and pre-production.

3.2 Design Overview

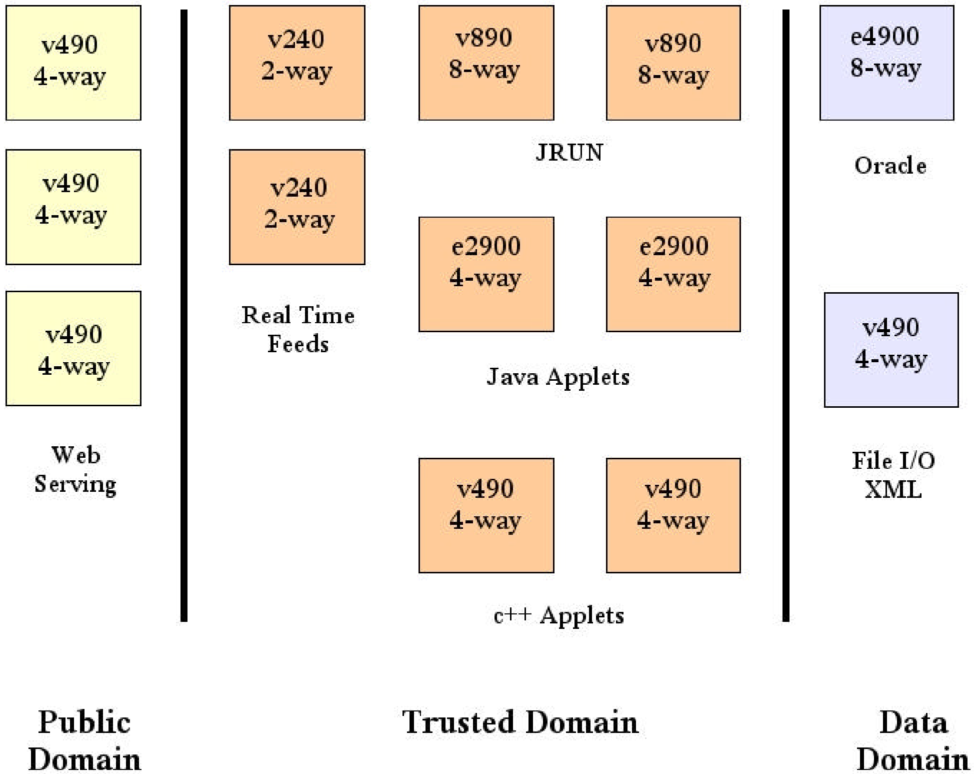

Drawing 1: Distributed Configuration Schematic

The discrete solution segregates the boxes into three domains each with differing security requirements.

3.2.1 CP Utilization

The central processor (CP) utilization during prime shift by application type is as follows:

Component	***Average CP Prime***
Web serving	15%
Real-time feeds	23%
JRUN	4%
Java servlets	27%
C++ servlets	17%

Component	*Average CP Prime*
Oracle	7%
File I/O XML	6.5%

Table 1: Component Prime CP Utilization

3.3 Component List

Using the same combination of servers as described in the original "Capricorn" paper [1] but upgrading the hardware to contemporary levels the list of components is as follows (prices from shop.sun.com):

Component	*Function*	*Quantity*	*Unit Cost 000's*	*Total Cost 000's*
V490	Web Servers	3	59	177
V240	Real Time Feeds	2	8	16
V890	JRUN Application Servers	2	71	142
E2900	Java Applet Server	2	94	188
V490	C++ Applet Server	2	59	118
E4900	Oracle DB	1	240	240
V490	File I/O & XML Server	1	59	59
Oracle 10g	Database	8	40	320
Total				**1260**

Table 2: Component Description and Cost

Given the advances in processor speed since 2004 this configuration is of much greater capacity than the original.

3.4 Analysis

The discrete solution results in 13 physical management points, each with their own cabling, powering and environmental requirements. The above costs are conservative, as the actual cost to the customer would need to include maintenance costs for the software, which is often calculated by CPU as well.

4 Virtualized Server Deployment in System z9

4.1 Deployment Architecture and Design Goals

For the purposes of this comparison, the analysis assumes that server utilization is at or about the levels measured in 2004. The configuration described in the previous section is going to be lifted from its discrete server base to a virtual server complex running in a single z9 EC which is controlled by z/VM 5.2. The use of the z/VM SHARE command is used to establish the relative importance of each of the servers to meet service level requirements.

Section 2 of this analysis explained that good candidates for consolidation are those with low average processor utilization. As you can see by table 1 “low average utilization” does not necessarily mean all components running at under 10%, but that when all servers are examined the across-the-board utilization falls into a range of 10-20%. This selection of components also accounts for spikes in utilization by any one server. A characteristic of many distributed applications is that for every heavy processing component there are several supporting elements that do relatively little processing at all.

The triple site redundancy described previously may now be implemented either by acquiring the services of a Disaster Recovery facility (and thus lowering acquisition costs) or by acquiring additional System z processors at the remote sites and using the high availability techniques described in [4]. This type of provisioning greatly simplifies DR planning and implementation.

Using server cloning, a copy of the environment can be made that would implement QA and developer configurations operating in the same z/VM system or in a different LPAR (or a remote site if the company is planning to do its own disaster recovery).

4.2 Design Overview

The following schematic illustrates the same components as the discrete server but they now exist as virtual machines running under z/VM on a System z9 EC.

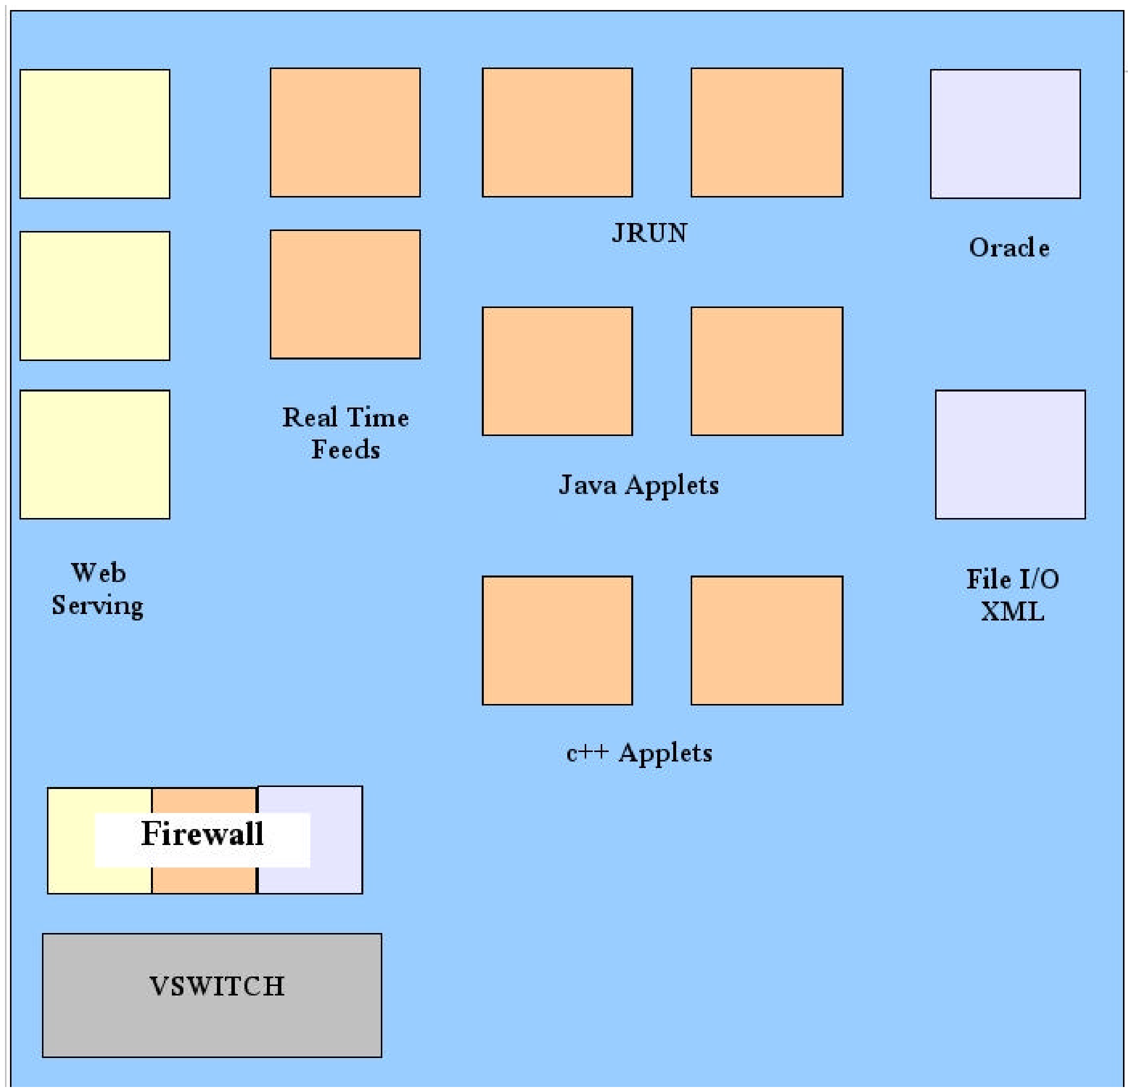

4.3 Components

For this configuration, the analysis assumes that this configuration will use an existing firewall system which, in conjunction with the virtual LAN support, would provide the necessary isolation between the domains at no additional cost. Alternatively, it would be possible to use a Linux on System z based firewall running within the z/VM complex and the virtual LAN facility of the virtual switch to implement public, trusted, and data domains within the System z.

Component	*Function*	*Quantity*	*Unit Cost*	*Total Cost 000's*
z9 BC	System z9 server with 2 IFLs - 64GB memory - 4 FICON cards - 4 OSA cards	1	680	680
z/VM	Hypervisor + Support Products	1	69	69
SUSE SLES 10	Linux O/S (licensed for 2 IFLs)	1	22	22
Oracle 10g	Oracle DB	2	40	80
Total				**851**

Table 3: Virtual Server Components and Cost

4.4 Analysis

The System z based version of the portal system results in one physical point that has power, cooling, and other environmental requirements, in contrast to the 13 of the discrete solution. The System z configuration also affords the ability to duplicate a test and development environment (or even multiple environments) on this same server.

5 Study Results and Analysis

5.1 Architecture and Design Goals

The intent of a virtual server architecture is threefold:

1. Provide a platform which can be grown and managed without regular acquisitions of additional hardware, people resources or other infrastructure.

2. Provide the flexibility to implement variations of the configuration and variations in service levels to meet future needs as well as providing a test and development environment.

3. Simplify disaster recovery requirements.

5.2 Operations Simplification and Business Resilience

The great strengths of consolidating on the System z include its ability to simplify provisioning, operations, and disaster recovery.

5.2.1 Provisioning

In “Consolidating Oracle Databases to Linux on System z” [2], a provisioning solution via server cloning was used by the Government of Quebec that demonstrated the following characteristics:

- Creating/Installation a new server
    - Linux : 30 minutes
    - SUN, AIX, Windows : between 1 week and 3 months (if RFP needed)
- Cloning/Installation an Oracle/DB instance
    - z/VM-Linux : 30-45 minutes
    - SUN : 10-14 hours

Initially growing a server farm in a discrete server environment simply means acquiring a new box, plugging it into an available power source and connecting it to a switch for network connectivity. Often there are delays as suppliers need to ship equipment from warehouse to the glasshouse. However, there comes a point where there are no more available power outlets, the floor space has run out, or the air conditioning system cannot cope with any additional heat sources. Suddenly the $5000 server turns into the $500,000 server as the environmental limitations are overcome by acquisition of additional infrastructure.

Conversely adding a virtual server to a System z can be as simple as a cloning operation. Of course if the System z based virtual server farm continues to grow then there will come a point where additional capacity is required. However, this usually means simply enabling another CPU that's already on a system board. At worst it may mean another frame. In either case the amount of capacity added will provide for significant growth before those new resources are exhausted.

The marginal power, cooling and operational requirements of the environment grow at a much lower rate than in a discrete server configuration.

5.2.2 Operations

From a physical perspective using virtual servers means that operations staff do not need to walk the data center floor to check its connections etc. In large server farms this type of checking can mean traversing relatively large distances in order to check many hundreds of servers. The reduction of the physical requirements of maintenance significantly improves the efficiency of a facility's operations.

Servers running on System z have access to the same tool set as used on the discrete servers. For example, the Tivoli suite for managing all aspects of the servers' operations is fully available on System z. There are also several tools native to the System z environment that greatly simplify the operation and monitoring of virtual servers.

DGTIC in the Government of Quebec case study [2] has stated that they are able to operate their production systems with a 100 servers per system administrator ratio. They also have an additional two z/VM system programmers whose roles, now that the project is maturing, will be reabsorbed in their existing mainframe support group.

5.2.3 Resiliency and Disaster Recovery

Consider the case where a fan on a server is rated to have a MTBF of 50 years. Now consider a server farm consisting of two hundred servers. On average 4 servers per year will suffer a thermal failure. If we also take into account the other physical components of the discrete server farm with similar MTBF ratings, there is a high likelihood that during any given week there will be a physical failure of some sort and that a business process will be interrupted until the server is recovered or the workload moved.

This analysis can be demonstrated with a simple mental calculation: If there are 10 components per server with MTBFs of 50 years (for example, fan, power supply, disk drive, hard disk, CD-ROM drive, network card, management interface, memory card, processor and motherboard), then there will be, on average, 40 outages per year. That means that there is a greater than 75% likelihood of an outage during any given week.

Using virtual servers eliminates most of the exposure to a physical failure. The System z is subject to thermal and other physical or environmental exposures. However, the strength of System z is the amount of redundancy that will overcome or circumvent the vast majority of these exposures. From multiple power supplies, redundant processors, an error correcting circuitry the System z has used its 40 years of mainframe heritage to provide perhaps the most robust server on the market. The total, single-system MTBF for a System z is greater than 55 years [5].

The use of a virtual server farm also simplifies the disaster recovery requirement specifications for a given installation. Even issues such as the administration of CPU-based license keys are greatly simplified when you are dealing with virtual CPUs.

"Consolidating Oracle Databases to Linux on System z" [2] quotes a 6-7 hour recovery time for z/OS, z/VM, Linux on System z, Oracle databases, and the network. Simply having fewer moving pieces and being able to use a common set of tools enables this speed of recovery.

5.3 Security Auditing and Compliance Management

A System z based consolidation has important security benefits:

- Vast reduction in physical cables that means a lower potential of network "sniffing."
- Reduction in physical objects that could be compromised or stolen.
- Access to hardware encryption facilities that are shared across all guests.
- Use of proven security products such as RACF to secure the virtual and physical resources.

5.4 Cost of Acquisition

Often, if a comparison is made purely on the cost of acquiring a single production configuration then the discrete server is going to be cheaper. However, this is a false economy and in many cases since the introduction of the z9 BC the discrete server is **not** the less-expensive option. The discrete server option provides a (limited) solution where the System z option provides a complete platform where deploying development, QA, and pre-production systems is a simple matter of cloning, not provisioning. The System z reduces the total cost of operation across the breadth of computing parameters, from physical components, to floorspace and environmentals, through staff efficiency and recovery time. When the number of users that are able to use the System z's resources (i.e. developers etc.) is considered, then the total cost per user plunges.

The incremental cost to add a server to System z is substantially less than a discrete server configuration. Apart from the costs of the physical machine, the cost of the software licenses is also lower. System z also provides a larger number of servers that can be added before additional infrastructure is required when compared with a discrete server solution. Each discrete server requires floor space, power, and air conditioning. A company can add dozens – or hundreds – of virtual servers before another IFL would need to be added (often by way of enabling a processor already present in the complex). It takes several additional IFLs before another frame is required. This means that potentially thousands of individual instances can be added to a System z before any additional floor space, power or air conditioning is required. For a discrete server solution, thousands of extra servers means exponential growth in the cost of time, space, power, staffing and environmental controls, not to mention the security implications and attendant costs.

Another factor not to be overlooked is that cloning new servers for use in creating new production portals or for development and testing does not result in increased license charges. All virtual machines operate under the same serial number and so are not subject to higher license costs as servers are added.

"Consolidating Oracle Databases to Linux on System z" [2] describes a break even period for their consolidation project of 2 years. However, this explicitly excludes the savings accrued from the reduced environmental requirements. The break-even point would be even sooner when such savings are included as part of the analysis.

5.5 Notes and Discussion

The case study presented is generally conservative. It describes an existing infrastructure and virtualizes it on the basis that the existing servers are under utilized and that the System z in conjunction with z/VM is able to share the resources efficiently. Additionally, the infrastructure proposed enables the cloning of systems for the purposes of QA and system development without provisioning additional hardware. Of course there comes a point where the addition of a new virtual server will require the addition of another IFL. However, this addition does not occur with every server that's added and typically only involves enabling a new engine that already exists on the processor complex.

“Resiliency with Linux and zSeries” [3] takes a far more aggressive approach to consolidation. As of December 2006, the study reports that there are 165 Oracle instances on their z9 EC, with an estimate of around 190 by year's end 2006. 12 of these instances are in production as part of the Government of Quebec's portal. All this is achieved using a 5 IFL z9 EC running z/VM 5.2.

Appendix A References

1. “Capricorn Six Case Studies for Priority Investments in the Corporate Server Infrastructure” David G Heap, IBM 2003.

2. “Consolidating Oracle Databases to Linux on System z” Graylon Baselton, Denny Dutcavich, Karen Ann Plourde, and David Kreuter, IBM Webcast 2006:

   https://whiteglove.on.raindance.com/confmgr/view_stored_doc.jsp?docId=98066172957656428623387327285

3. “Resiliency with Linux and zSeries”, GM13-0618-01, IBM 2004.

4. “High Availability Architectures For Linux on IBM System z” Stephen Wehr,Harriet Morril, Scott Lovelando, Scott Epter, IBM April 2006.

5. “zSeries Mainframe Resurgence Beyond Question. Software/Costs See Major Advances: IBM Tools Break USV Grip” Ian Bramley, Software Strategies January 2004

