



Research Report

The Mainframe Virtualization Advantage: How to Save Over Million Dollars Using an IBM System z as a Linux Cloud Server

Executive Summary

Information technology (IT) executives should be aware that it is possible to save *over a million dollars* by consolidating Linux servers on a mainframe as opposed to deploying Linux on a group of x86-based multi-core blade servers.

The reason for this advantage is that large scale architecture does not require as much headroom (spare capacity) as smaller systems to execute heavy I/O workloads due to specialized input/output (I/O) characteristics. Accordingly, a mainframe can run 240 virtual machines (VMs) as compared with about 10 VMs on an Intel 8 core system (in this comparison, both systems run the same workload at the same service level).

By our estimate, the price of running a **heavy I/O workload** on 240 virtual machines running on a mainframe (at 70% CPU utilization with a high reliability service level profile) should be approximately \$3,300,000. The cost of running the same environment (on 240 virtualized machines, 24 blades, 192 CPUs) on a group of Nehalem EP-based Intel Xeon servers should be approximately \$4,800,000 (see Chart 1). *Choosing an IBM System z as a Linux/Cloud consolidation server, therefore, has the potential to save IT buyers over a million dollars!*

Chart 1 — 240 Virtual Machines on a Mainframe vs. Intel Multi-core Architecture

1 X System z 196 with 32 IFLs and associated software:	\$ 3,300,000
vs.	
24 X Xeon 8 core (192 CPUs) and associated software:	\$4,800,000

Notes: To be clear, we believe that Intel's x86 multi-core Xeon architecture is a solid microprocessor architecture that can run serial, parallel and data-oriented tasks in a balanced manner (x86 multi-cores are a good design). This architecture is particularly well suited to run infrastructure-oriented applications such as HTTP serving, firewalls, file/print, directory, and Web-facing workloads (we call these "light workloads") — but it is not the best architecture to choose for heavy I/O workloads. We also note that Intel servers can be configured to run more than 10 virtual machines when executing heavy I/O workloads, but if more than 10 are configured, then they will not meet comparable service level requirements.

The Mainframe Virtualization Advantage: How to Save Over a Million Dollars Using an IBM System z as a Linux Cloud Server

In this *Research Report*, *Clabby Analytics* explains why we believe that IBM System z 196 mainframe — when configured as a Linux/cloud consolidation server — can help enterprises save BIG MONEY when running heavy I/O workloads.

Fixing Bad Assumptions

One of our biggest concerns in the IT industry right now is that many IT buyers believe that Intel x86-based multi-core server are a better deal due to their lower total cost of acquisition (TCA) than other types of servers. Accordingly, these buyers opt for x86 multi-cores without doing a proper TCA cost analysis. For example, many IT buyers will frequently compare costs by only looking at the cost-of-acquisition of the hardware — and, after conducting an analysis such as this, Intel architecture wins every time...

What these IT buyers are missing, however, are a few fundamentals. They should not be comparing these systems based-on the basis of CPU processing power or hardware cost (because these systems have radically different headroom and I/O characteristics) — they should be looking at how much work each server can do. And that's when things get really interesting because large scale systems have more headroom to work with — and can thus create and operate significantly more virtual machines, than smaller scale servers.

Mainframe Architecture Has a Major Virtualization Advantage Over Smaller Scale Architectures

A mainframe, running at 70% CPU utilization while processing heavy I/O intensive workloads can run up to **240 virtual machines**. An Intel-based Xeon 8 core server running the same workload at the same CPU utilization rate with the same service level requirements can run approximately **10 virtual machines**. To reach 240 Linux virtual machines you would need 24 Intel x86-based multi-core machines (192 CPUs).

Why is this the case? The answer to this question is simple and logical. The whole design point of large scale mainframe architecture is based on sharing resources (the mainframe is known as a “*shared everything*” architecture). Mainframes share memory; they share a very large internal communications bus; they share central processing units, disk, and so on. In a scale-up shared environment, all of these resources (memory, the communications bus, CPU processing power, etc.) can be made available to a common pool — all within a single chassis (a self-contained mainframe architecture).

Demand for resources in this pool are constantly fluctuating (usage peaks and valleys) just like in a distributed environment — but, interestingly, the peaks-and-valleys tend to balance out a lot better in a mainframe large scale resource pool.

Smaller scale servers (such as Intel Xeon multi-cores) do not have as large a base of resources available in a common pool (CPU, memory, and I/O is bounded within each server). And, as a result, smaller scale servers need to be over-provisioned (more headroom needs to be allocated) to handle usage peaks and valleys within a given small server environment.

Where's the Proof?

In order to prove that this mainframe headroom advantage is real, we relied heavily on finding published in a benchmark study conducted by IBM's Software Group Project Office (this report can be found at:

The Mainframe Virtualization Advantage: How to Save Over a Million Dollars Using an IBM System z as a Linux Cloud Server

<ftp://public.dhe.ibm.com/common/ssi/ecm/en/zsw03125usen/ZSW03125USEN.PDF>).

Now, before you say “this study is biased because it was produced by IBM”, we ask you to consider the following points:

- IBM sells a lot of x86 Xeon multi-core systems. It does not help IBM to disparage x86 server platforms;
- This study was done in 2009, before Nehalem EP (Intel’s first real Xeon multi-core server architecture) was released. So you could argue that the report compares a mainframe to older Xeon architecture. To remedy this, we supply an updated graph (see Figure 1) based on more current Xeon architecture; and,
- We have verified the core principle of this paper — that scale-up mainframe architecture manages headroom/capacity better than x86 servers — with IT executives who use both architectures (and they, too, have observed this phenomenon).

The Spirit of This Benchmark Study


It is important to understand the spirit of this benchmark. IBM engineers were looking for a way to explain why mainframes are able to host more virtual machines than smaller x86 multi-core environments — not to disparage x86 servers. So they constructed a model that:

- used the same workload;
- measured CPU and memory usage;
- considered service level requirements; and that
- measured the number of virtual machines that could be sustained while meeting the same service level agreements on each architecture.

What these engineers found was that, when running the same workload at the same CPU capacity, a mainframe requires less headroom. Using formulae described in the above mentioned report — and by running the same workload across three environments (System z, Power Systems, and x86 multi-core servers) — these engineers were able to actually measure how many virtual machines could be created in each environment.

Figure 1: Headroom Requirements for IBM System z and Power Systems Compared to x86 Multi-core Architecture

**The Mainframe Virtualization Advantage: How to Save Over
a Million Dollars Using an IBM System z as a Linux Cloud Server**



Source: IBM Corporation — January, 2011

Figure 1 shows that an 8 core x86 blade running a heavy input/output intensive application in a high-availability profile can run about 10 virtual machines. Accordingly, to reach 240 virtual machines, an IT buyer would require 24 x Intel Xeon multi-core servers. Figure 1 also shows several calculations for headroom on an IBM System x (Intel Xeon multi-core), an IBM Power System blade, and a System z mainframe. In the case of the Xeon blade server, take the amount of CPU used (the mean — shown as “M” — which in this case is about 32% of the CPU) and multiply it by 2.581 (the amount of headroom needed to run 10 virtual machines running heavy I/O workloads and the calculation comes out to 98% (almost 100%) total capacity. The mainframe, running at 73% CPU utilization needs only 1.323 of the remaining server capacity (headroom) to support 240 virtual machines. x86 servers need much more headroom (the mean times 2.581). So the bottom line is this: each x86-based server needs to allocate almost twice as much headroom as an IBM System z in order to meet its quality of service requirements.

Other Factors That Work in Favor of a Mainframe as Compared to x86 Architecture

Systems design plays a huge role in the mainframe advantage over x86 architecture.

IBM’s Integrated Facility for Linux (IFL) is a specialty processor for executing Linux and z/VM on the mainframe at a lower price point. In addition to the general or specialty processors like the IFL, the mainframe also has System Assist Processors (SAPs) for offloading certain kinds of processing tasks. Instead of having the mainframe CPU heavily involved in handling IO tasks, mainframes offload these processing tasks to specialized hardware devices like these I/O processors. The reason that this is important is that the number of virtualized servers a given system can support is highly dependent on how much CPU power is available — and because mainframes offload so much extra processing like I/O handling, more CPU is made available to run additional virtual machines.

The Mainframe Virtualization Advantage: How to Save Over a Million Dollars Using an IBM System z as a Linux Cloud Server

Mainframe design includes a very high-speed internal bus architecture that speeds process-to-process communications within a mainframe chassis. Contrast this with distributed systems that frequently have to drive their inter-process communications between various servers through internal network adapters and across external network hub, bridges, and routers (which frequently become bottlenecked when workload spikes occur) and the performance of internal, high-speed mainframe I/O should be readily obvious.

In addition to its high speed bus advantage, mainframes are also easier to manage than distributed systems environments. Logically, IT buyers need to ask themselves this question: “Which is easier to manage — a single system with a large internal communications bus, co-located CPUs, and a large cache of internal memory — or a group of distributed computers that rely heavily on external networking devices and that have trouble sharing memory across systems boundaries?” If your answer is “a self-contained, single system environment”, you’re right...

(Note: we did not include a comparison of the security and management costs for mainframes and distributed systems, but if we had done so, we suspect that difference in managing and securing a mainframe would be *tremendously* less than the costs to secure and manage a distributed systems environment).

The Mainframe Virtualization Advantage: How to Save Over a Million Dollars Using an IBM System z as a Linux Cloud Server

Getting Your Head Straight: Focus on Workloads Rather Than Acquisition Costs

Clabby Analytics likes Intel x86 multi-core Xeon architecture. It does a very balanced job of handling data, parallel and serial workloads. And it is indeed attractive from a cost of acquisition perspective. But what IT buyers need to understand is that no processor architecture does all jobs well — and that system design differences can make a huge difference in how much it costs to process various workloads.

Next month we will be publishing a report that takes a closer look at which workloads belong on which types of servers (x86, Power Systems, or mainframes). But, we do have a few findings that we are willing to share with our readers at present.

We have come to believe that:

- Major differences in processing efficiency start to show up when running different workloads on x86, RISC, and mainframe servers. For instance:
 - When running *heavy input/output workloads* (transactions with many reads to memory, the network, and many writes to disk), mainframes are a far better choice than x86 servers. It can be significantly less expensive to run these workloads on highly-scalable, high-speed internal bus mainframe architecture as opposed to either Unix/RISC or x86 servers;
 - When it comes to *heavy workload processing* (the processing of compute-intensive tasks), Unix/RISC machines are the most efficient processor/-server combo; and,
 - x86 servers are most cost effective when used to process numerous, heavily-threaded *light workloads* (multiple, fast, low-cost threads) that require light memory use, and that require comparatively lower quality-of-service (QoS) than offered on mainframes and RISC machines.

Based-upon these initial findings, we believe that enterprises trying to build the most cost effective cloud architecture — and that have varied workloads — may be best served by building a heterogeneous systems architecture where workloads can be assigned to the server types best suited to serve those workloads.

Summary Observations

To us, the three most important messages in this report are:

- 1) *CPU efficiency is the throttle that controls the amount of headroom needed (the more efficient your processor, the less headroom needs to be allocated). IBM's System z CPU is the fastest in the industry (clocked at 5.2 GHz) — and this makes a huge difference in the amount of work that it can process as compared to other commercially available microprocessors.*
- 2) *The "shared everything" design point of large servers also has a major impact on the number of virtual machines that can be created. If you cannot easily share all of a system's resources, you lose the headroom advantage that large systems offer; and,*
- 3) *In addition to the general processors, a System z has 14 SAPs (System Assist Processors) + up to 336 channel I/O processors + 1024 logical channels. These additional processors give the System z a huge advantage in terms of processing I/O (no x86 server design has an I/O subsystem even remotely like a mainframe).*

The Mainframe Virtualization Advantage: How to Save Over a Million Dollars Using an IBM System z as a Linux Cloud Server

All of these features work in tandem to enable mainframes to claim more headroom.

No single processor/systems architecture does all jobs optimally. Mainframes, with their fast CPU and with their multiple, additional communications processors are particularly strong at heavy I/O workloads. By comparison, Intel architectures are particularly strong at multi-threaded “light-and-fast” workloads. IT executives should understand these differences, *because choosing the wrong architecture can waste millions of IT budget dollars.*

What should you do next? Contact IBM and get the company to tell you about its special packaging and pricing for its mainframe Linux server environments. The company is focusing on growing Linux workloads on the System z platform — and has special discount programs in place to promote the deployment of new workloads on Linux on System z. Also, ask IBM specifically about its Enterprise Linux Server (for more details on this server visit: <http://www-03.ibm.com/systems/z/os/linux/els.html>).

Finally, our bottom line is this: if your organization needs to deploy hundreds of Linux servers and needs to run heavy I/O workloads, it should be looking at an IBM System z. If a z is not considered, your organization could be wasting hundreds of thousands and potentially millions of dollars by deploying the wrong workloads on x86-based servers.

Clabby Analytics
<http://www.clabbyanalytics.com>
Telephone: 001 (207) 846-6662

© 2011 Clabby Analytics
All rights reserved
February, 2011

ZSL03117-USEN-00

Clabby Analytics is an independent technology research and analysis organization. Unlike many other research firms, we advocate certain positions — and encourage our readers to find counter opinions — then balance both points-of-view in order to decide on a course of action. Other research and analysis conducted by Clabby Analytics can be found at: www.ClabbyAnalytics.com.