July 2009

Advantages of a Dynamic Infrastructure: A Closer Look at Private Cloud TCO

Scott A Bain
Innes Read
John J Thomas
Fehmina Merchant
IBM SWG Competitive Project Office

Table of Contents

Executive Summary

Many companies are finding their need for greater business agility being frustrated by an increasingly costly and rigid IT infrastructure. The culprits are many. Maintenance of the current environment accounts for over 70% of the IT budget, leaving less than 30% available for new projects. Annual operational costs (power, cooling, and management) of distributed systems and networking exceed their acquisition cost by 2-3X and continue to climb. Utilization rates of these commodity servers hover around 5% on average, leading to excess capacity going to waste. Time to provision new servers can be as long as six months, hampering lines-of-business efforts to quickly respond to competitive threats or new opportunities. As a result, LOB units are beginning to go outside the datacenter to public cloud providers like Amazon in hopes of lowering their costs and improving their responsiveness. To avoid disintermediation, IT needs to re-invent the datacenter by moving towards a more dynamic infrastructure. One that takes out cost through the use of virtualization and consolidation to improve utilization levels with a commensurate reduction in power consumption. One that embraces a private cloud model that dynamically provisions IT services in minutes/hours rather than months (and at lower cost) via self-service portals.

This paper examines the Total Cost of Ownership (TCO) for a dynamic infrastructure built around private cloud services and compares it to public cloud alternatives as well as conventional one-application-per-distributed server models. The results show that private cloud implementations built around larger virtualized x86 and IBM System z® servers can be up to 75% less expensive than public cloud options over a five year period and around 80% less than a distributed stand-alone server approach.

Take Cost Out Through Virtualization and Consolidation

A recent IBM internal study of its nearly 4000 distributed servers showed annual operational costs attributed to each server to be over $34,000, with almost 90% due to software maintenance and systems administration. It stands to reason that reducing the number of physical servers to fewer, larger, more capable machines can serve to greatly reduce these costs. Indeed, the virtues of virtualization and consolidation to accomplish this have been well-publicized. What has proven to be more elusive, however, is the quantification of these benefits. How many workloads can actually be consolidated onto a given platform while maintaining acceptable service level agreements? Which platform gives you the greatest economy of scale, producing the lowest cost per virtual machine image/workload?

One approach to answering these questions and estimating the TCO of a private cloud environment, as well as compare it to other alternatives, is as follows:

1. Determine the expected consolidation ratio for a given workload type (e.g. Windows, Linux®, etc.)
2. Estimate the annual cost to operate the virtualized servers (over 3 or 5 years)
3. Compare to stand-alone provisioning or public cloud services

To determine the consolidation ratio, the CPO conducted a series of performance benchmarks. A sample banking application was built using IBM WebSphere® Application Server and DB2® running on Linux and run on an older 4-way (single core each) IBM x366 server using 3.66 GHz Intel® processors and 12 GB of memory. Average CPU utilization was 5%, throughput was 4.5 transactions per sec, and average response time was 40 milliseconds. A VM image of that workload was then created and placed on an 8-core IBM x3950 server (four 3.5 GHz dual-core processors) with 64 GB of total physical memory and running a popular x86-based hypervisor. Multiple running instances of this VM image were added incrementally to the server until it could no longer handle any additional throughput. CPU utilization, throughput, and response time metrics were captured throughout. This same test was then applied to a single frame IBM System z10® Enterprise Class (z10 EC®) machine (8 IFL cores @ 4.4 GHz) running z/VM® as a hypervisor. The results are shown in the charts below:

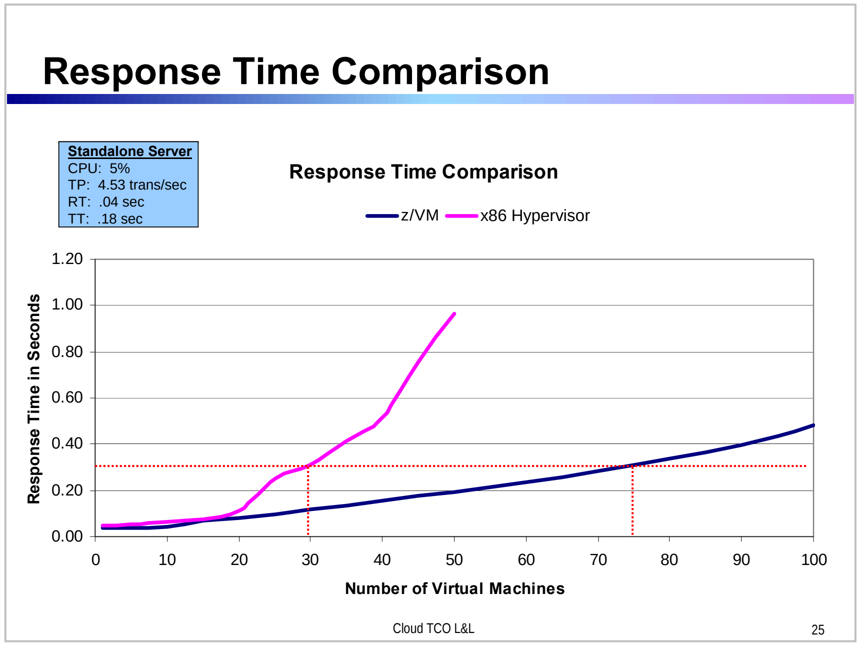
Response Time Comparison
Standalone Server
CPU: 5%
TP: 4.53 trans/sec
RT: .04 sec
TT: .18 sec
Response Time Comparison
z/VM
x86 Hypervisor
Response Time in Seconds
1.20
1.00
0.80
0.60
0.40
0.20
0.00
0 10 20 30 40 50 60 70 80 90 100
Number of Virtual Machines
Cloud TCO L&L
25

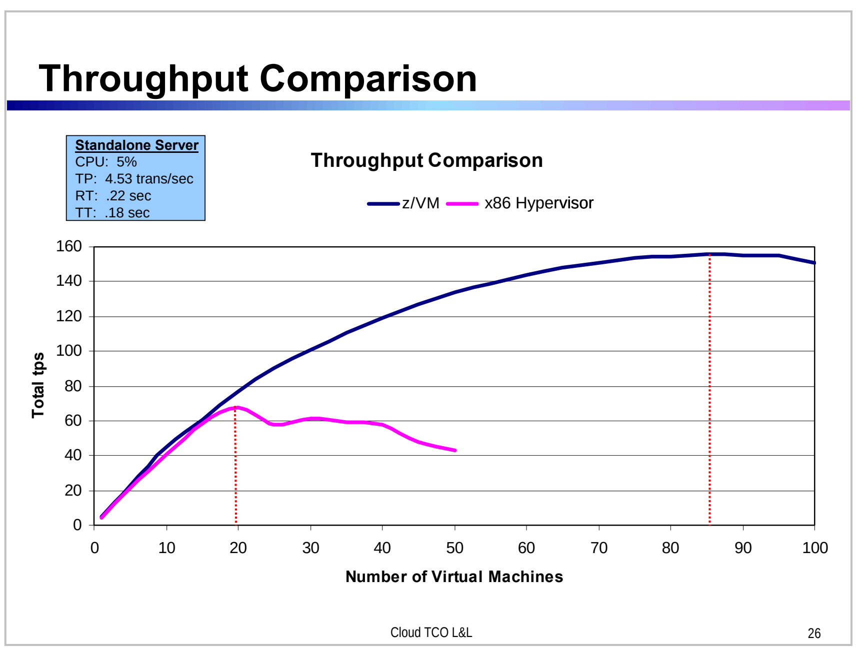
Throughput Comparison
Standalone Server
CPU: 5%
TP: 4.53 trans/sec
RT: .22 sec
TT: .18 sec
Throughput Comparison
z/VM
x86 Hypervisor
Total tps
160
140
120
100
80
60
40
20
0
0 10 20 30 40 50 60 70 80 90 100
Number of Virtual Machines
Cloud TCO L&L
26

For the x86 hypervisor case, the maximum number of VM images that could be supported with acceptable response time was approximately 30. The z/VM case, on the other hand, showed a maximum of about 75 images. Using maximum throughput as the metric reduced the maximum number of VM images to less than 20 for the x86-based hypervisor, whereas it increased the z/VM number to about 85.

Some customers use CPU utilization as their SLA metric. It should be noted, however, that these benchmark results were achieved by driving the machines to 100% utilization with no variation in workload. In practice, real world workloads have variability in demand. This variability has an effect on what CPU utilization can be tolerated on the target consolidation platform. Ideally, you want to be able to run utilization levels high enough to achieve the highest consolidation ratio, but still keep it less than 100% to allow for peaks caused by this variance in workload demand.

One way to gauge the impact of workload variability on consolidation ratios is to use statistical models. If we were to track average utilization over time of workloads with varying demand, the results would tend to exhibit a "bell curve" type distribution pattern as shown below. Theory tells us that 95% of all values fall within two standard deviations from the mean. Thus, if we had a Service Level Agreement (SLA) that called for us to be able to handle 95% of all workloads that occur on the system over a given time period, we would need to size a machine with a capacity of the mean + two times the standard deviation (sigma).

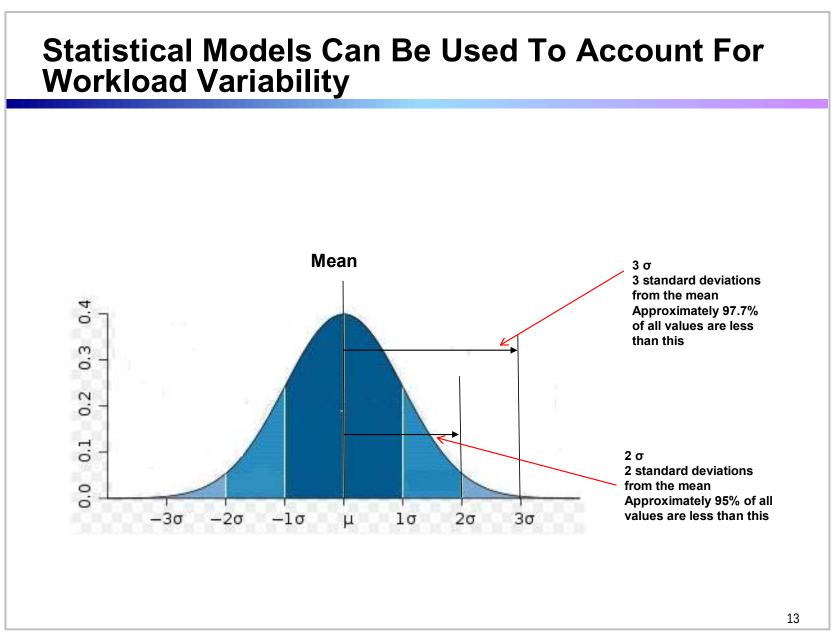

It turns out that when you combine workloads with variable demand, the overall variability of the combined workloads gets smaller. Thus, larger machines capable of accommodating more workloads can run at higher utilization levels than smaller ones. When we apply an SLA of 95%, the maximum number of VM images that can be supported by the x86 hypervisor system is reduced to 8, whereas the z/VM case can only handle a maximum of 40 images. More details behind these calculations will be forthcoming in a separate whitepaper.

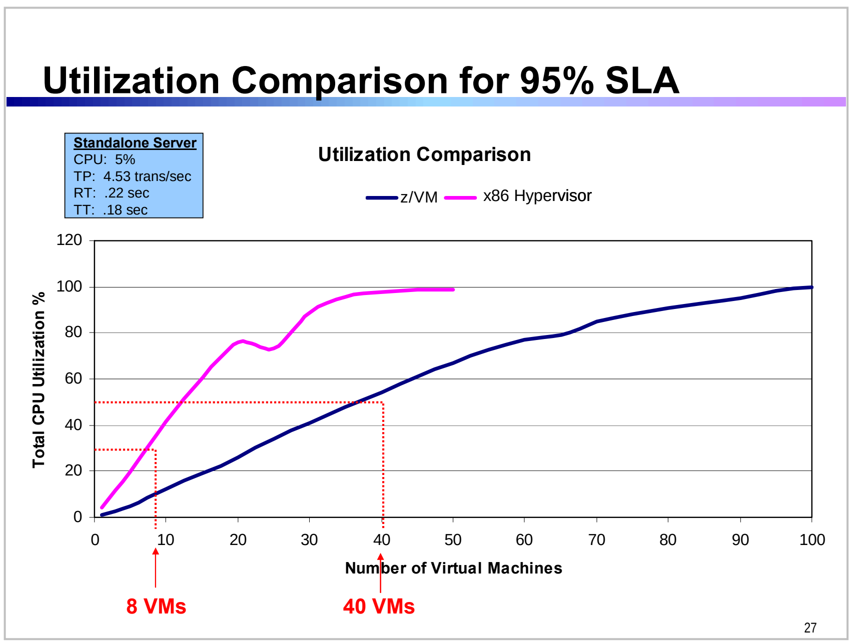

One other factor needs to be considered. These revised consolidation ratios assume "flat-out" operation. In practice, not all servers/images are used all the time. Experience from customers and public cloud providers have shown typical usage patterns of 14 hours per day (59%). This means in theory, we should be able to run 1.7X more images on each platform in order to maximize usage, thus increasing the consolidation ratio to almost 14 for the x86 hypervisor case and 68 for the z/VM case.

The next step in the analysis called for estimating the TCO over 5 years for running 100 Linux images using four different platforms to see which one delivered the lowest cost per image or workload:

1. Buy stand-alone x86 servers (running one image/workload on each)
2. Rent Amazon EC2 instances (running one image/workload on each)
3. Buy large x86 servers and provision virtual servers using an x86 hypervisor (private cloud)
4. Upgrade an existing z10® EC machine and provision virtual servers using z/VM (private cloud)

TCO components included hardware, software (Linux, WAS, DB2, and IBM Tivoli® Composite Application Manager for Applications), maintenance, facilities (power/cooling), and administration and assumed 24x7 operation. Administrative costs were derived from IBM RACE and other internal studies.

The results of this TCO comparison appear below:

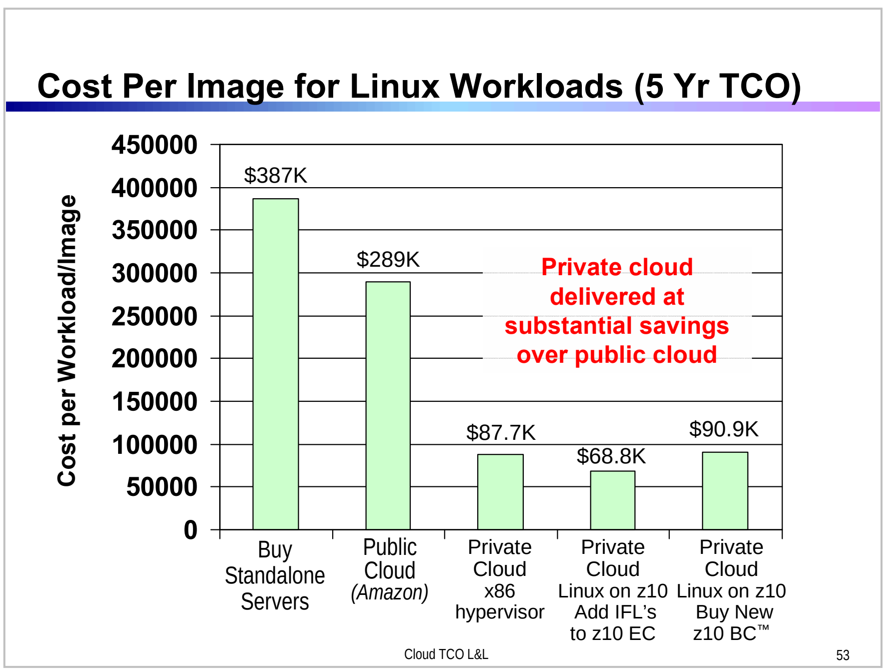

As expected, doing "business-as-usual" and buying stand-alone (or rack) servers is the most expensive option ($387,000 per image/workload). Upgrading an existing z10 EC system and provisioning your own virtual servers is the least expensive alternative ($68,800 or 82% less), while doing the same provisioning on larger x86 servers with a hypervisor resulted in an $87,700 charge per image (77% less). What may be surprising to some people is the fact that the public cloud option using Amazon EC2 instances was so expensive ($289,000 per image). This is due to several factors. First, using Amazon EC2 instances continuously for 24x7 operation results in higher runtime platform charges (almost 3X) than acquiring physical servers and provisioning workloads in-house to achieve higher utilization rates. In addition, customers must purchase software on a per EC2 instance basis rather than being able to take advantage of on-premise multi-core systems that can support multiple images on a given hardware platform. Finally, although individual server management is eliminated with the public cloud, there is still significant labor costs involved in the administration of each running application instance.

Perhaps more impressive is the cost per image of a new IBM System z10 Business Class™ (z10 BC™) system for those who don't currently utilize IBM System z hardware. At $90,900 per image/workload, this system is only slightly more expensive than even x86-based servers running a hypervisor.

A breakdown of each of the cost components for each scenario appears below:

Detailed Cost Breakdown for Linux Workloads (5 Yr TCO)

	Buy Another Server	**Rent a Virtual Server**	**Provision Your Own (x86 hypervisor)**	**Provision Your Own (Linux on System z10)**
Runtime Platform	100 IBM x3250 with 4 cores each	100 Amazon Extra Large EC2 instances	8 IBM x3950 with 8 cores each	12 IFLs added to existing IBM z10 EC
Hardware Costs ▪ Server ▪ Storage ▪ Networking	$5,000,000	$2,880,000	$1,390,000	$3,000,000
Software Costs ▪ OS (Linux) ▪ Hypervisor (on x86, z/Linux) ▪ App Server (IBM WAS) ▪ Database (IBM DB2) ▪ Monitoring (ITCAM for Apps)	$22,660,000	$22,010,000	$3,987,000	$2,644,000
Facilities and Admin ▪ Power ▪ Floor space ▪ Maintenance ▪ Systems admin	$11,020,000	$4,020,000 *(admin only)*	$3,395,000	$1,240,000
Total Cost	$38,680,000	$28,910,000	$8,772,000	$6,884,000
Number of Workloads/Images Supported	100	100	100	100
Total Cost per Image	$386,800	$289,100	$87,720	$68,840

All told, customers electing to provision their own virtual servers in a private cloud setting will find it far less expensive than either conventional stand-alone servers or public cloud alternatives.

Request-Driven Provisioning Through Self-Service Portals

While virtualization, consolidation, and service management provide the basic underpinnings of a dynamic infrastructure, it must be accompanied by a self-service portal that enables users to request IT services on demand and have the request fulfilled in minutes/hours versus days/weeks/months.

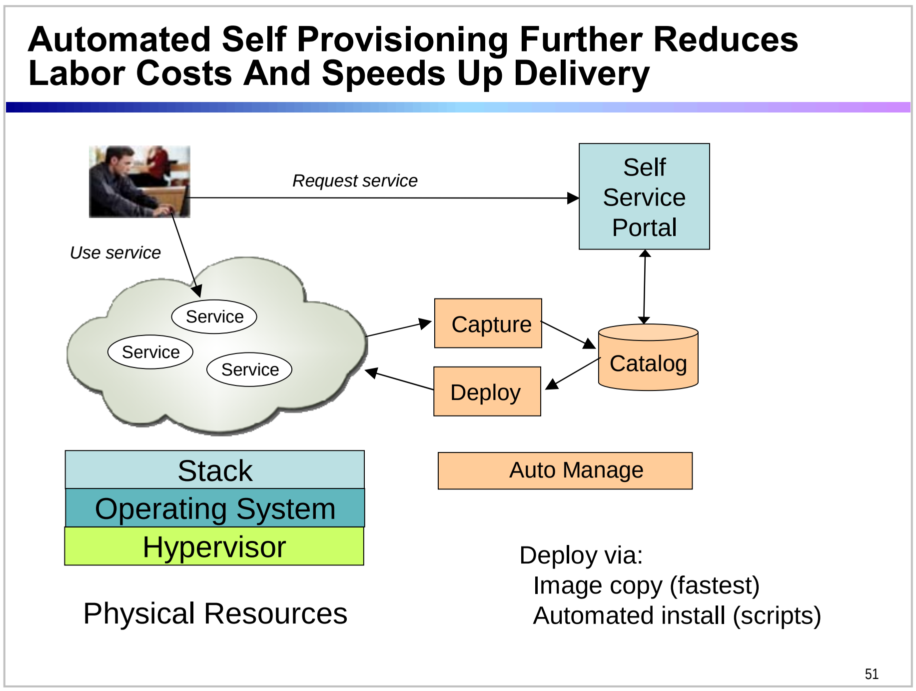

In this model, services are initially defined/created and stored in a service catalog. Requesters can then browse the catalog to find and select the desired service. After submitting the request, it gets routed for approval and then fulfilled by the underlying infrastructure. The software needed as part of the overall service is typically deployed in one of two ways: image copy (the fastest) or via automated install using scripts. When the service is no longer needed, the affected resources are freed up so that they can be claimed by other subsequent requests. In order for all of this to work seamlessly and transparently to the user, there needs to be automated management software that undergirds each step in the process.

IBM recently introduced Tivoli Service Automation Manager (TSAM) to manage this cloud services lifecycle and deliver request-driven provisioning for a private cloud environment. It leverages Tivoli Service Request Manager (TSRM) to provide a self-service UI for users to search against the catalog and select the desired service. It also utilizes Tivoli Provisioning Manager (TPM) to provision hardware and software resources according to best practices to satisfy the service request.

Summary

Escalating business requirements will continue to drive companies toward datacenter transformation. This includes pursuing ways to take costs out of their existing infrastructure, such as the use of virtualization and consolidation that reduce the number of physical servers needed and lowers energy consumption. Adding self-service, automated provisioning of IT services on top of this foundation to create a dynamic infrastructure allows IT to respond more quickly to the needs of the business. It also allows them to do so at lower cost than other alternatives (including public clouds) and avoid the threat of disintermediation. While multiple target platforms are available from which to deliver this private cloud environment, customers who have already made investments in IBM's System z platform will find it to be an attractive and cost-effective option.

IBM Systems and Technology Group
Route 100
Somers, New York 10589
U.S.A.
Produced in the United States of America,
07/2009

ZSW03126-USEN-00