A Study on Reducing Labor Costs Through the Use of WebSphere Cloudburst Appliance

IBM Software Group, Competitive Project Office

Scott Bain
Faisal Rajib
Barbara Sannerud
Dr. John Shedletsky
Dr. Barry Willner

With thanks to
Andrew Bradfield
M A Fellwock
Fehmina Merchant
John J. Thomas

Contents

Executive Summary

Annual IT operational costs continue to increase, with labor commanding a larger and larger share. For example, an IBM internal study of its own distributed infrastructure showed labor to be over 60% of the total operational cost per year[1], while industry analysts estimate labor costs can be as high as 80% of overall data center costs[2].

As a result, many customers are turning to private clouds, implementing such technologies as virtualization and consolidation, standardized workloads, and automation by way of self-service provisioning in an effort to reduce these costs. While only 12% of enterprises currently utilize some of these techniques, this number is expected to rise to 50% by 2012[3].

Quantifying the impact that private cloud technologies have on various aspects of labor, however, has proven elusive, resulting in slower adoption rates. Customers want to know, for example, just how much these solutions will affect the labor required to set-up and maintain both the physical as well as virtual infrastructure for a given deployment platform before committing resources to their implementation.

This paper describes an approach to help answer this question. We first looked at the impact of **virtualization** by constructing a labor model that calculated the total labor hours required to set-up and maintain the overall infrastructure for both stand-alone and virtualized x86-based server environments. Using actual customer data, the model allowed us to calculate the breakdown between the labor required for the physical and virtual infrastructure for a given number of workloads and specified time period. Next, we looked at the impact of **standardization**. We adjusted the model with a "clone" factor to reflect the fact that many companies implementing private clouds are using standardized workloads that can be easily copied or cloned to other virtualized servers to reduce labor costs even further. Finally, we took a look at **automation** by conducting a hands-on study to capture the time it took for an administrator to deploy an application manually onto a VMware-based server versus using the WebSphere Cloudburst Appliance (WCA). As the name implies, WebSphere Cloudburst is a hardware appliance that encapsulates 10+ years of best practices to automatically dispense pre-defined, customizable images of WebSphere Application Server to virtualized servers. Using the results from this study, we refined our labor model further to account for labor savings to be had through the use of automation with WCA.

Following this overall approach yielded the following observations:

- **Over five years, the labor associated with the management of a physical hardware server was more or less the same as that associated with managing one unique software image**
- **The greater the consolidation you can achieve, the lower you can reduce total physical server labor hours**

[1] IBM Internal Consolidation Project
[2] Source: Butler Group 2007 and http://www.itmanagement.com/blog/20070129/report-indicates-mainframe-adoption-continuing-to-grow/
[3] Internal IBM Cloud study 2009

- **The more images you can standardize and clone, the lower you can reduce software labor hours**
- **WebSphere Cloudburst Appliance can reduce software labor hours by up to 80% compared to manual deployment**

Labor and the Server Provisioning Lifecycle

Any discussion of labor needs to start with a process that describes the tasks associated with the acquisition, deployment and retirement of servers. Servers are first planned and acquired, then they are handed over to administrators to configure, set up and deploy. The operating systems software is installed, Hypervisors are configured, virtual servers configured, security profiles for users established, and the server is tested and deployed into production. Monthly maintenance continues including routine patches and fixes, and upgrades. The servers are ultimately cleansed and retired from service.

Figure 1 below depicts this provisioning lifecycle approach. It includes some procurement functions, set up and deployment functions, maintenance, troubleshooting and ultimate tear down. The labor categories included setup and tear down costs as well as the ongoing monthly maintenance and troubleshooting costs for physical servers and software virtual images.

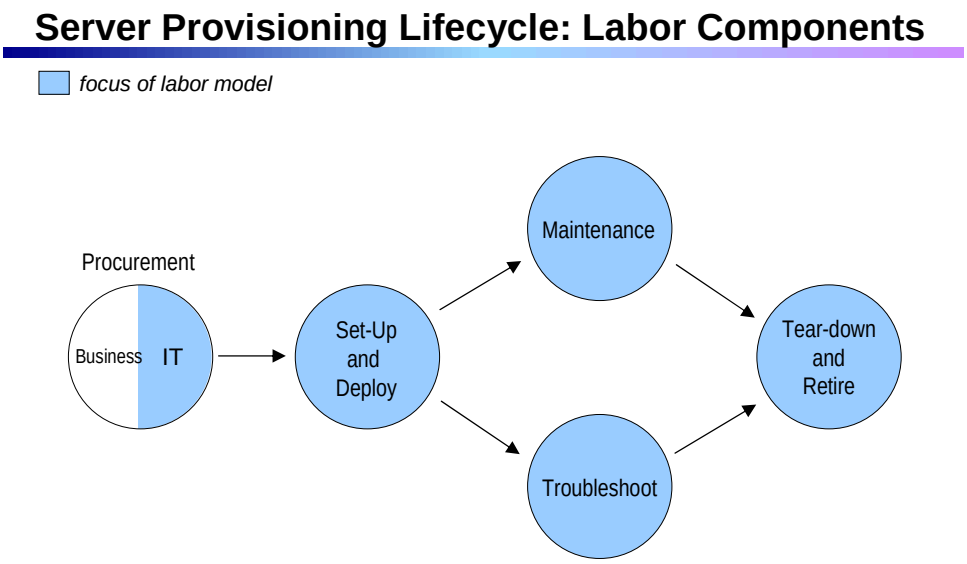

Figure 1

Quantifying the Impact of Virtualization on Labor

Virtualization plays a prominent role in the quest to reduce labor costs. To quantify its impact, we first devised a labor model based on the server provisioning lifecycle as shown in Figure 2 below.

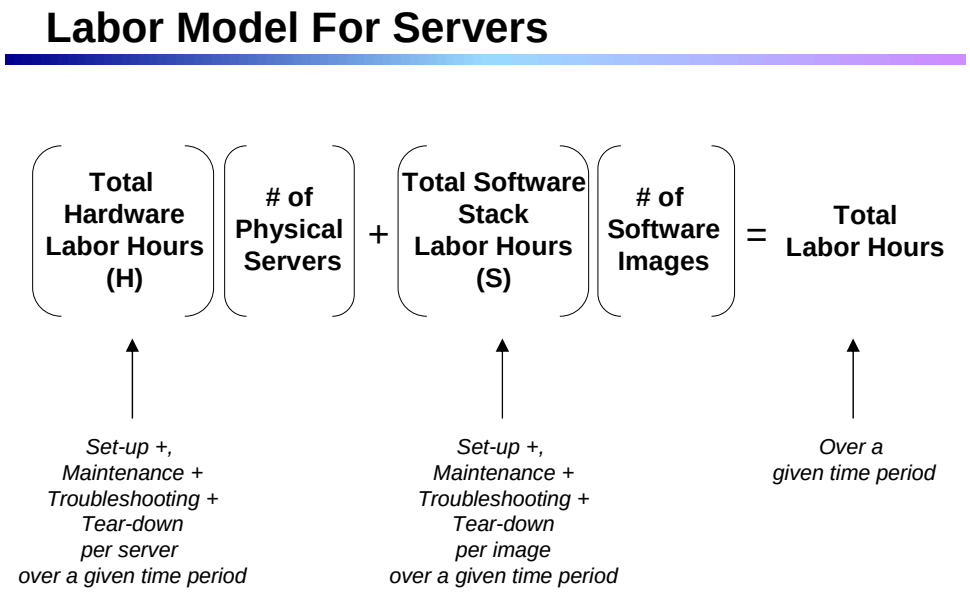

Figure 2

The formula represents the total labor hours ascribed to the management of a server environment as comprised of the hours spent managing a physical server over its lifetime plus the hours spent managing the software images over their lifetime. Total hardware server labor hours (H) include the set up and deployment hours representing one-time events such as sizing and configuring workloads, and testing of a physical computing element. They also include hours for scrubbing of servers, decommissioning, maintenance and troubleshooting for physical servers over the analysis period. Total software labor hours (S) include both the initial installation labor associated with the software stack or virtual images on the physical server along with ongoing maintenance and troubleshooting over the assessment period. These tasks include periodic patching and upgrades, associated testing functions, analysis of errors, debugging, fixes, testing and reboots.

Solving this equation for a stand-alone x86 environment gives us a picture of how much labor was required *before* virtualization. Similarly, solving the equation for the virtualized x86 environment gives us insight into the total hours needed *after* virtualization. Fortunately, we have data from customer case studies that can help us evaluate both equations.

The customer data we collected was based on the average number of servers managed per FTE (Full-Time Equivalent or administrator), which is a metric widely accepted in the industry. To come up with an overall average across all customers, we first grouped and ordered the data along the lines of "most efficient" (e.g. higher number of servers per

FTE) down to the “least efficient” (e.g. lower number of servers per FTE) as shown in Figure 3 below:

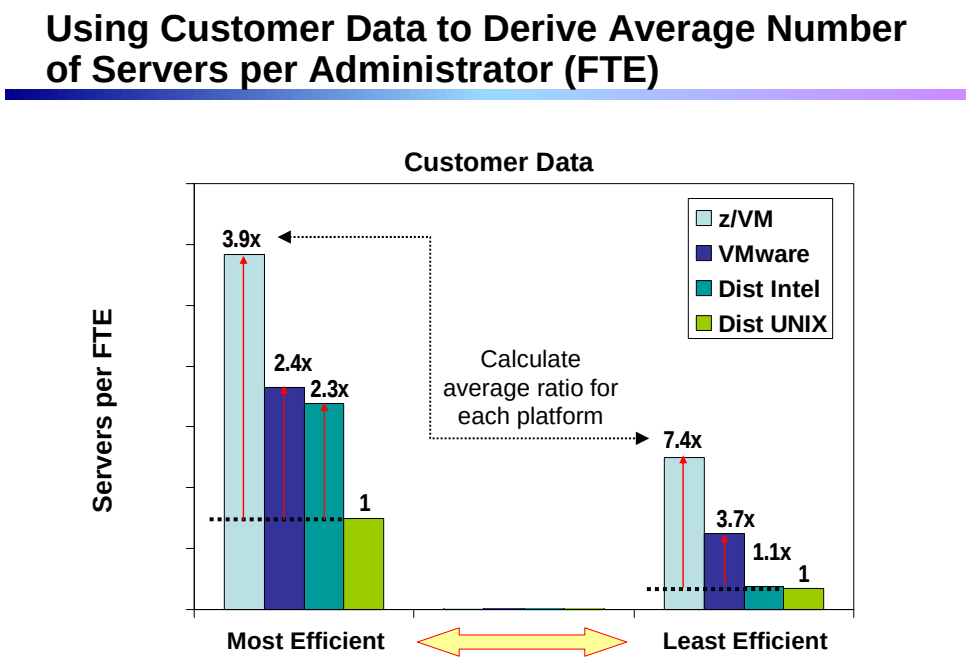

Figure 3

We then normalized the server/FTE value for distributed UNIX to 1 and calculated the relative server/FTE ratios across the different platforms within each ‘efficiency’ grouping. Next, we computed the overall average ratio for each platform and applied the results to the distributed UNIX baseline to come up with the average servers/FTE for all platforms as shown in Figure 4 below:

Using Customer Data to Derive Average Number of Servers per Administrator

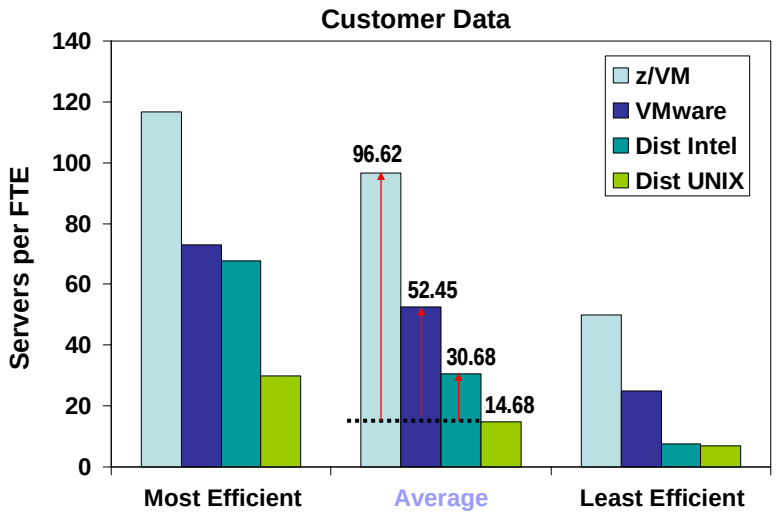

Figure 4

For the stand-alone x86 server case, this worked out to be 30.7 servers/FTE, while the virtualized x86 server case turned out to be 52.5 servers per administrator.

We then wanted to calculate the portion of *FTE labor needed to manage a server*.

Calculating the FTEs per server for stand-alone and virtualized x86-based servers

- **Stand-alone x86** data shows 30.7 servers managed per FTE,

  1/30.7= **.0326 FTE's** needed per server

- **Virtualized x86** data shows 52.5 virtual servers managed per FTE,

  1/52.5 = **.0191 FTE's** needed per server

Next, we wrote equations to represent the total FTE hours required to manage 100 Linux workloads over 5 years for both stand-alone and virtualized x86 platforms.

We assumed 10,400 hours or 52 weeks per year, 8-hour days for 5 years.

FTE hours needed to manage 100 workloads over 5 years:
Multiply **FTEs needed per server** * **total hours over 5 yrs.** * **number of software images**

- .0326 * 10,400 *100 = 33,904 hours needed for all stand-alone x86 servers
- .0191 * 10,400 *100= 19,864 hours needed for all virtualized x86 servers

On balance, this shows **a virtualized x86 environment requires 42% less total labor hours to manage 100 Linux workloads over 5 years than the stand-alone x86 scenario.** But what percentage of that time can be attributed to managing the hardware (H) vs. managing the software images (S)? To answer this question, we first need to determine the number of servers required to handle 100 Linux workloads. For the stand-alone case, this is straight-forward: you need one server for each workload, or 100 total physical servers. For the virtualized case, we need to determine how many of these workloads can be consolidated on a given server platform. Based on studies previously conducted by the SWG Competitive Project Office[4], we found that you could consolidate 8 stand-alone server workloads on a single 8-core x3950 system (Intel Xeon 3.5 GHz processors) as shown in Figure 5 below:

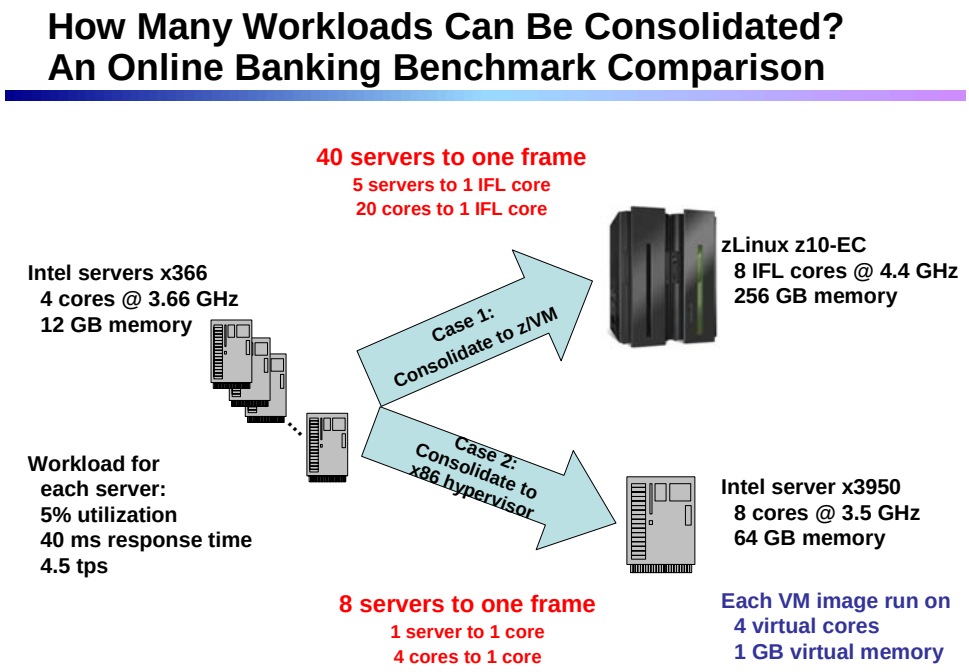

Figure 5

This means it would take 13 physical servers (100/8 = 12.8, rounded up to 13) to handle our 100 Linux workloads.

Given that both the stand-alone and virtualized scenarios are running on x86 platforms, it stands to reason that the hardware labor hours (H) and software image labor hours (S) are essentially the same for each alternative. Thus, we are left with the following equations:

(1) Stand-alone x86 100H+100S =33,904

(2) Virtualized x86 13H+100S =19,864

We have two equations in two variables, which we can then solve:

[4] "A Benchmark Study on Virtualization Platforms for Private Clouds", https://w3-03.ibm.com/sales/competition/compdlib.nsf/SearchView/1035807516E8C5E1852575F2005B01DE?Opendocument

Subtracting equation 2 from equation 1 to solve for H and S:

H = 161.38 hours

S = 177.66 hours

Therefore, over a 5-year planning horizon, the total hardware labor (H) to manage one server is 161.38 hours and the cost to manage a single software image (S) is 177.66 hours.

An immediate observation is that these values are *not* very different from one another; in other words, **the labor hours required to manage the hardware and software images over five years are roughly similar** when handling 100 Linux workloads on an x86-based server platform.

Quantifying the Impact of Standardization on Labor

As businesses begin to embrace virtualization, the variety of software images that need to be managed can quickly proliferate, resulting in higher labor costs if left unchecked. One way to address this problem is identify workloads that can be standardized and cloned. With standardization, much of the variability associated with deployment and maintenance of unique images is eliminated. This use of cloning dramatically reduces maintenance time, as the patches, testing and upgrades should be identical across cloned images. The question is, how can we quantify the material impact standardization has on reducing labor costs?

To estimate this, we applied a cloning factor to our original equation as shown below in Figure 6:

Use of Standardized Stacks Can Drive Down the Labor Hours for Virtual Images

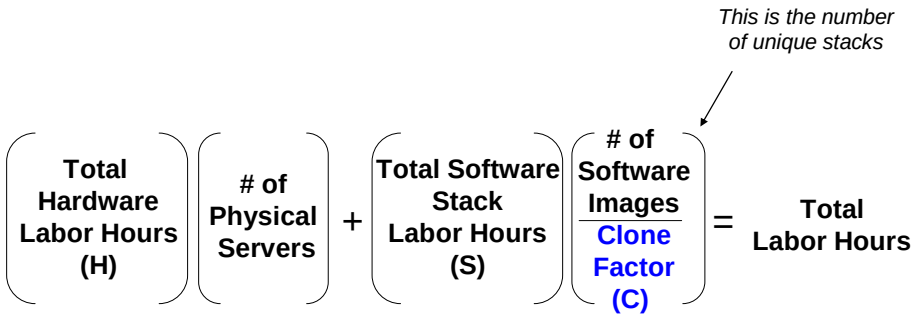

Where C = average number of copies deployed for each unique software stack (from 1 to 100 in our example)

Figure 6

Solving this equation for the virtualized x86 environment discussed earlier in the paper yields the following:

13H + 100(S/C) = total labor hours

Since we already know H and S from our previous calculations, we can substitute those values, resulting in the following:

13(161.38) + 100(177.66)/C = total labor hours, or

2097.44 + 177.66(100/C) = total labor hours

Expressing the formula this way allows us to play some "what if" games with the clone factor (C) to gauge the impact of standardization on total labor hours. For example, applying a clone factor of five would mean that out of 100 servers there are 100/5 or 20 unique images deployed, of which the rest are duplicates of the original five unique templates. This reduces the overall labor hours from the original virtualized x86 case of 19,864 to 5,654, a reduction of 72%!

The graph below in Figure 7 shows the labor savings to be had as you adjust the clone factor "C" between no clones (1) and 100 clones (100).

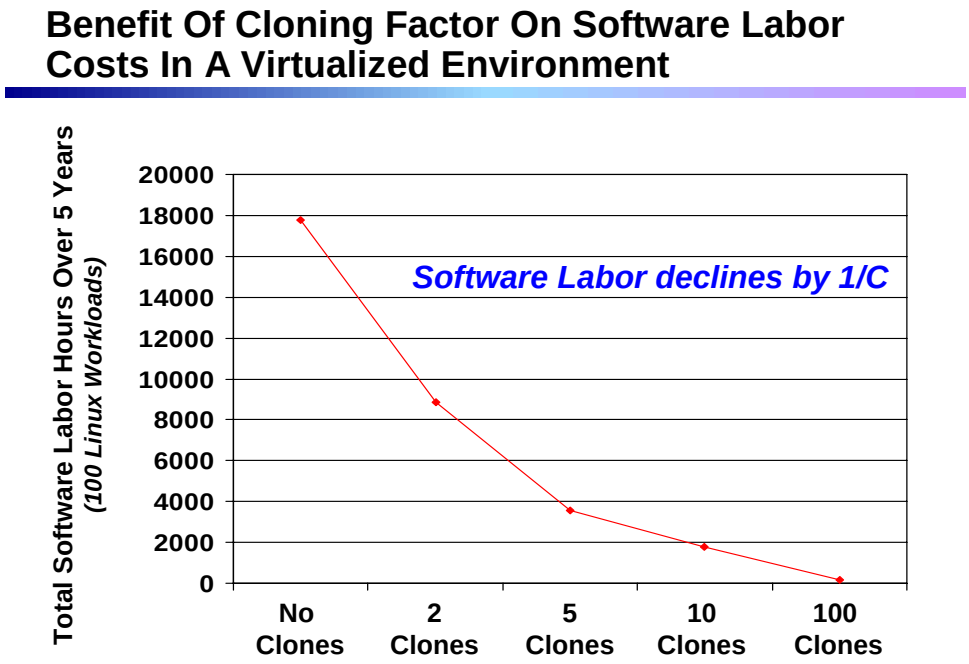

Figure 7

As you can see from the curve, total software labor hours decline by roughly the inverse of the cloning factor. Based on this revised labor model that takes into account the use of clones, we can make the following observations as shown in Figure 8:

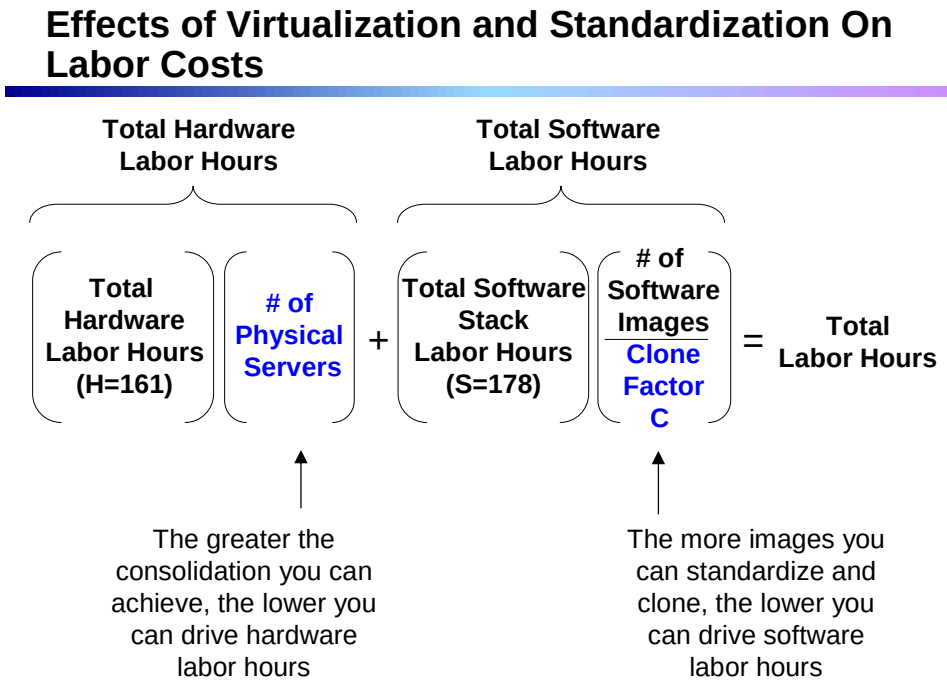

Figure 8

One of the levers in reducing labor costs is to reduce the number of physical servers you have to manage. Put another way, **the more workloads you can consolidate on a given platform, the more you can lower your labor costs**. This makes larger, more scalable systems like the IBM System p and System z ideal virtualization and consolidation platforms for implementing private clouds.

Another lever is the degree to which you can use workload standardization and cloning in your environment. Simply stated, **the higher the clone factor, the greater the reduction in labor costs associated with deploying and maintaining software virtual images**.

Quantifying the Impact of Automation on Labor

While virtualization and standardization can go a long way in reducing overall labor costs, the task of deploying a software stack as a VM image onto a virtualized server has historically been a highly labor-intensive task. For instance, one has to first deploy and configure the OS along with all requisite patches. After that, the administrator has to install and configure the application server and all its constituent components (e.g. HTTP server, etc.) as well as patches and other fixes. For applications requiring a database, that becomes yet another piece of middleware that needs to be installed and configured. Then there is the application itself. Collectively, deploying and testing a complete application manually can require days or weeks to accomplish depending upon its overall complexity. In a private cloud environment, this kind of turnaround is untenable.

WebSphere Cloudburst Appliance (WCA) is specifically designed to address this problem. Available as a hardware appliance, it takes 10+ years of best practices in WebSphere Application Server (WAS) deployments and encapsulates it into pre-defined, customizable images that can be dispensed to a variety of hypervisors used in virtualized servers. Its use of scripting and automation techniques greatly reduces the labor required to perform deployment tasks. Current configurations supported in the initial release of WCA include the use of WAS Hypervisor Edition running on the Linux OS that can be deployed to VMware ESX servers.

To help assess the extent to which the use of WebSphere Cloudburst Appliance can reduce labor hours, we conducted a hands-on study as shown on Figure 9 below:

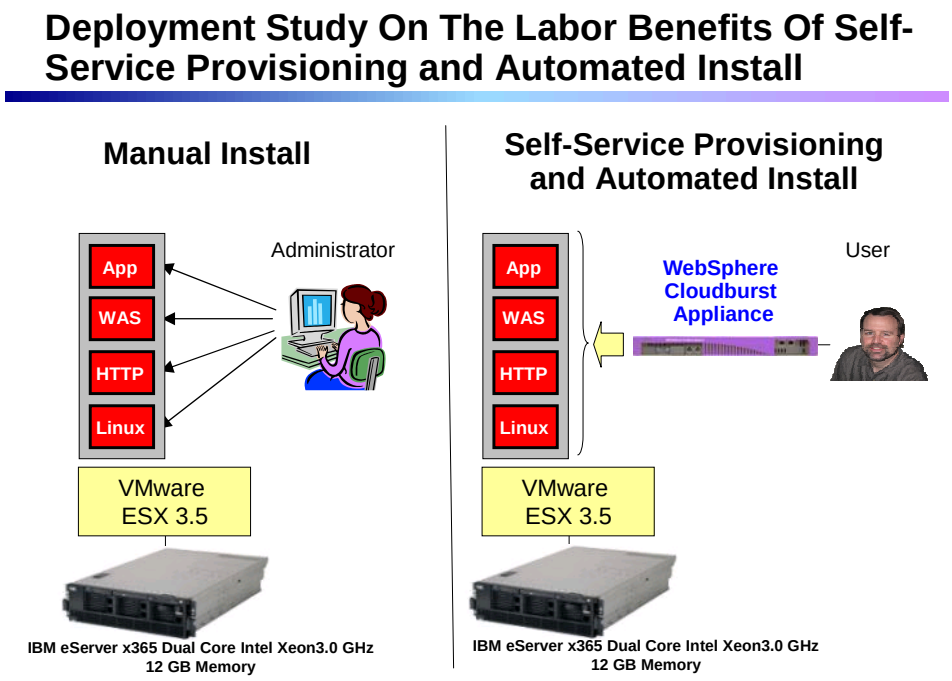

Figure 9

This study tracked the time it took to deploy and instantiate a WebSphere-based application on a virtual server using VMware. We captured metrics for doing this manually as well as using WCA. The results from this study show that the use of automation via the WebSphere Cloudburst Appliance can reduce software image labor hours by as much as 80%! (Figure 10):

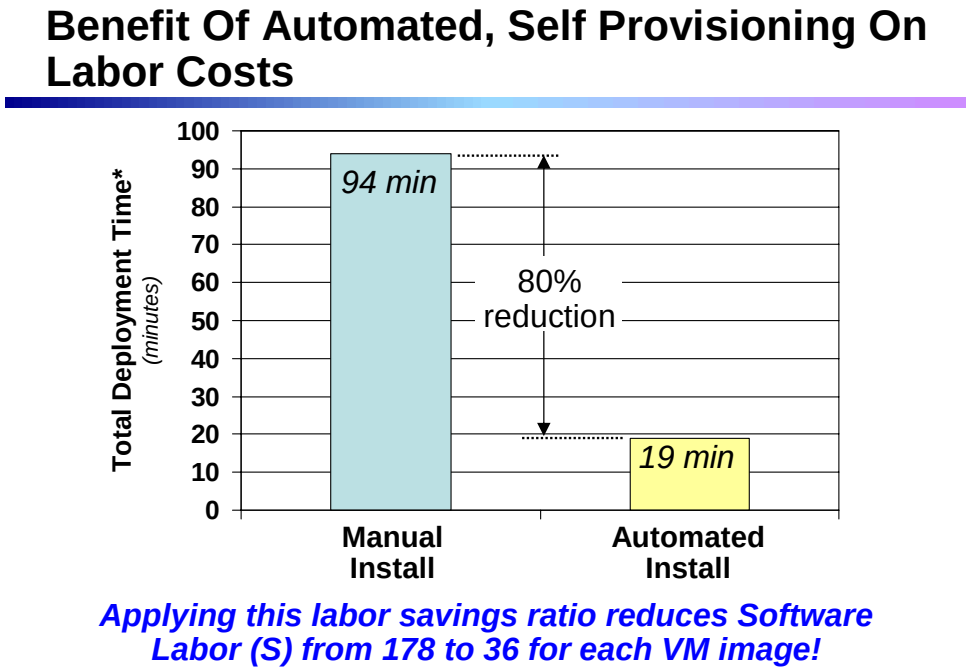

Figure 10

Putting It All Together

As our analysis shows, there are significant labor savings to be had through the use of virtualization, standardization, and automation. For our example of 100 Linux workloads over five years, virtualization by itself yields a 42% reduction while standardization alone reduces labor hours up to 72% with just a modest clone factor (C=5). Using WebSphere Cloudburst Appliance for automation results in a reduction of 80%. Taken collectively, companies can reduce their labor costs by up to 97% compared to a traditional stand-alone x86 environment and manual deployment methods (Figure 11):

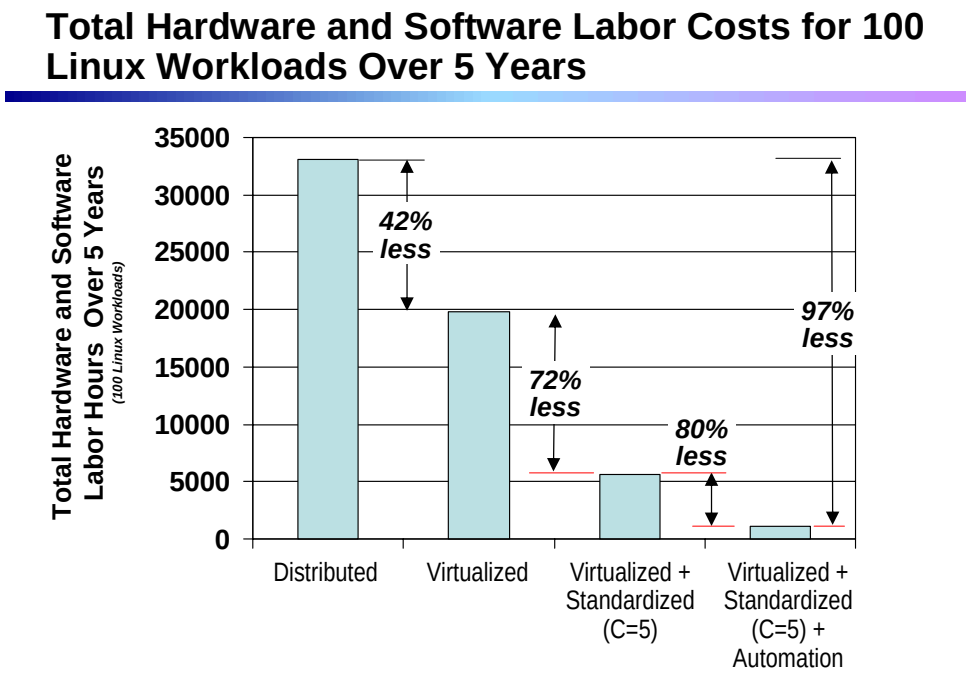

Figure 11

Summary

As overall labor costs continue to escalate, companies who embark on a private cloud strategy that uses virtualization, standardization, and automation to drive down these costs will find themselves in a competitively advantageous position. The labor model described throughput this paper can be used to estimate potential savings for a number of different deployment scenarios and technology choices. In our example, we chose to highlight the advantages of IBM's WebSphere Cloudburst Appliance as a means to achieve automation. Regardless of the implementation style, the labor model provides direction on the benefits that can be expected:

- **Over five years, the labor associated with the management of a physical server was more or less the same as that associated with managing each unique virtual image**

- **The greater the consolidation you can achieve, the lower you can reduce total physical server labor hours**
- **The more images you can standardize and clone, the lower you can reduce software image labor hours**
- **WebSphere Cloudburst Appliance can reduce labor hours for a unique software image by up to 80% compared to manual deployment**

Our conclusion is that tremendous reductions in operational labor are possible for those customers who pursue strategies of virtualization, standardization, and automation.

IBM Corporation
Software Group
Route 100
Somers, NY10589
USA

Produced in the United States