Cloud Computing Payback

An explanation of where the ROI comes from

November, 2009

Richard Mayo
Senior Market Manager Cloud Computing
mayor@us.ibm.com

Charles Perng
IBM T.J. Watson Research Center
perng@us.ibm.com

Purpose of this Document

As companies throughout the world examine the business value of Cloud Computing, it is important to understand how Cloud can lower IT expenses. Understanding how long it takes until your business can recoup the investment it made in Cloud Computing is the "payback" period. In today's economic environment, this is a critical component of any investment analysis. When we speak with IT executives, regardless of the industry they are in, they want to have important answers to key cost saving questions about Cloud Computing. The intent of this document is to inform the reader about the five (5) key areas of costs savings that are associated with every Cloud infrastructure implementation. Each of the five areas will be defined in terms of their characteristics, the savings that your company would anticipate achieving based on the size of your environment and the data elements that need to be benchmarked and tracked to project and calculate the savings. The tracking and calculation of savings is a key element of any Cloud implementation program. We will also provide the list of underlying projects that comprise each cost saving area. These projects are the "action steps" that can be undertaken to obtain the savings for your organization.

The savings described in this document are based on the implementation of a cloud computing environment versus a traditional infrastructure. A traditional infrastructure typically has the characteristics of a single application per server with manual provisioning and a siloed management environment, which results in high costs and low productivity. An IBM study[1] has shown that moving from a traditional infrastructure to a public cloud can yield a modest reduction in costs. However, for significant cost reductions a private cloud infrastructure is required. This is due to the fact that the underlying cost savings of the public cloud accrue to the public cloud provider and are only partially past on to the public cloud user. Setting up a private cloud computing environment results in the largest cost savings based on today's public cloud pricing models.

[1] Advantages of a Dynamic Infrastructure: A Closer Look at Private Cloud TCO
IBM Software Group, Competitive Project Office**,** July, 2009
http://www-01.ibm.com/software/se/collaboration/pdf/CloudTCOWhitepaper_04212009.pdf

1. Overview

Today, organizations of all sizes are investigating cloud computing and the benefits it can bring to their company. Given the numerous claims of savings and productivity improvements, it can become difficult to understand exactly how these benefits might apply to your particular IT environment. The representations of savings in this guide are based on a 2009 Cloud Computing ROI study conducted by IBM Research. The study uses customer environment data and cost saving data from more than one dozen companies working with IBM on Cloud Computing. Input was taken directly from IBM clients as well as from IBM strategic outsourcing teams that manage these customer's environments. The results overview the financial benefits of implementing cloud computing and provide quantitative analysis of the payback. As you will see, the financial benefits of cloud computing are very real, the payback period is characteristically very short and the action steps to obtain these savings are very well defined.

2. The 5 key areas of cloud computing payback

The analysis of the customer payback for implementing cloud computing can be organized into five (5) key areas:

- Hardware
- Software
- Automated provisioning
- Productivity improvements
- System administration

All five of these areas can see significant reductions in cost and/or gains in productivity with the implementation of cloud computing. You probably could have named these five broad areas based on your own experience and knowledge of cloud; however, the significant size of the savings in certain sub-areas may surprise you.

One of the attributes and pre-requisites on the implementation path to a cloud computing environment is a highly virtualized environment. This "virtualization" will involve a consolidation of systems which will drive reductions in hardware costs. This is often the initial appeal of funding virtualization projects; however, the labor savings are even greater. Many companies still undertake the manual provisioning of IT systems, suffer long and costly delays while people wait for resources to become available, and distract highly skilled personnel from key project to focus on the mundane administration of systems. The automation of these tasks in a highly virtualized cloud environment can save significant labor costs while improving quality and productivity. The impact on the business in saved labor and use of time to accelerate projects can be significantly more than just the savings in hardware. The total savings substantially off-set the small incremental increase in software costs that are usually necessary to deliver virtualization and the service management component that are elements of every cloud computing environment. Let's look more closely at some views of the Cloud payback in more detail.

3. Examples of Cloud Computing client payback and ROI

The examples below show ROI projections from an internal IBM Research Cloud Computing study conducted in 2009. Each case is based on environmental data collected from a specific customer who preferred to remain anonymous in this publication; however their general industry and environmental size are shown. In this guide the size of the environment is based on the number of servers before consolidation, for small environments the range is 5 to 15 servers, for medium environments the range is 16 to 400 servers and for large environments it is over 400 servers. Regardless of size or industry, there is a very short payback time and very high projected annual and 3 year return on investment. As you will soon see, the payback period is shortest and the ROI is largest for the large environments, due to the economies of scale associated with the service management software. The cumulative cost comparison graphs shows that Cloud computing savings continue to grow over time.

The first example shows a banking customer with a large number of servers that are evolving to a cloud environment. In this example, the system administration costs are significantly reduced due to efficiency gains in administration driven by the implementation of service management software. The service management software enables an administrator to manage more systems and as a result fewer administrators are required. The provisioning costs are significantly reduced through the use of automation, which greatly reduces the time required to provision resources. This customer experienced a huge improvement in their productivity associated with their testing activities; however in overall percentage terms these savings are much smaller than the large savings in provisioning and system administrator costs. Productivity improvements can manifest themselves into increased revenues and happier customers over time. Those elements were not tracked as part of this study.

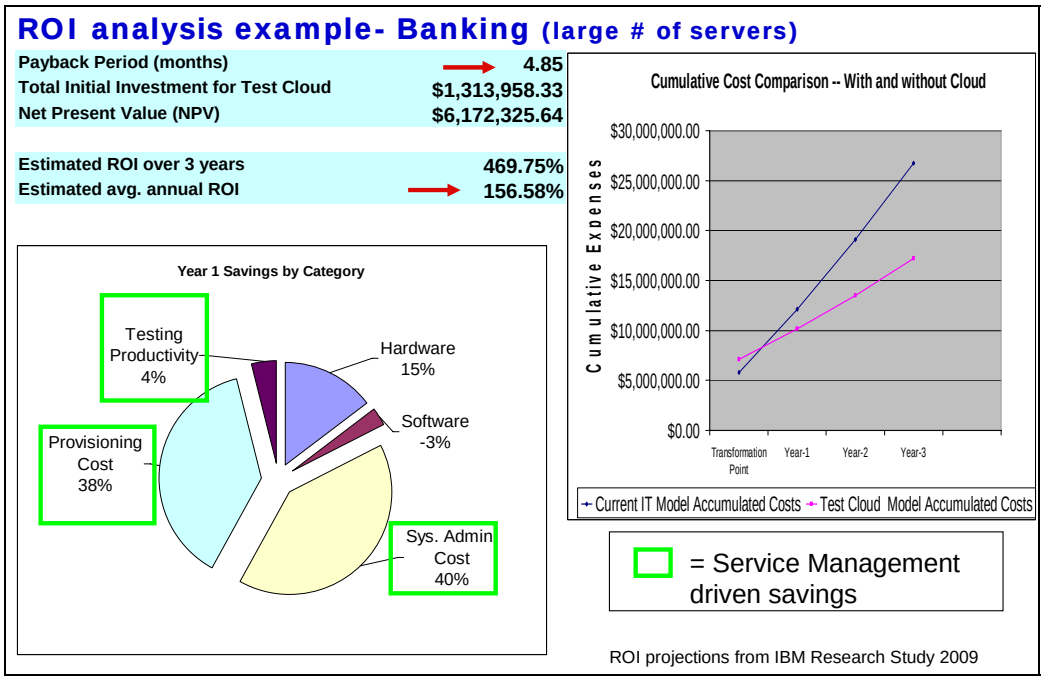

This second example is a manufacturing customer with a small number of servers that is currently an IBM strategic outsourcing customer. The study shows a large improvement in testing productivity based on a significant reduction in the tester wait time for standing up new and reconfigured systems from 70 hours to 1 hour. There was also a reduction in the number of system administrators that drove labor cost savings in the administration area. The provisioning savings appear smaller for this customer because the cost of maintaining the automation software was amortized over a small number of servers.

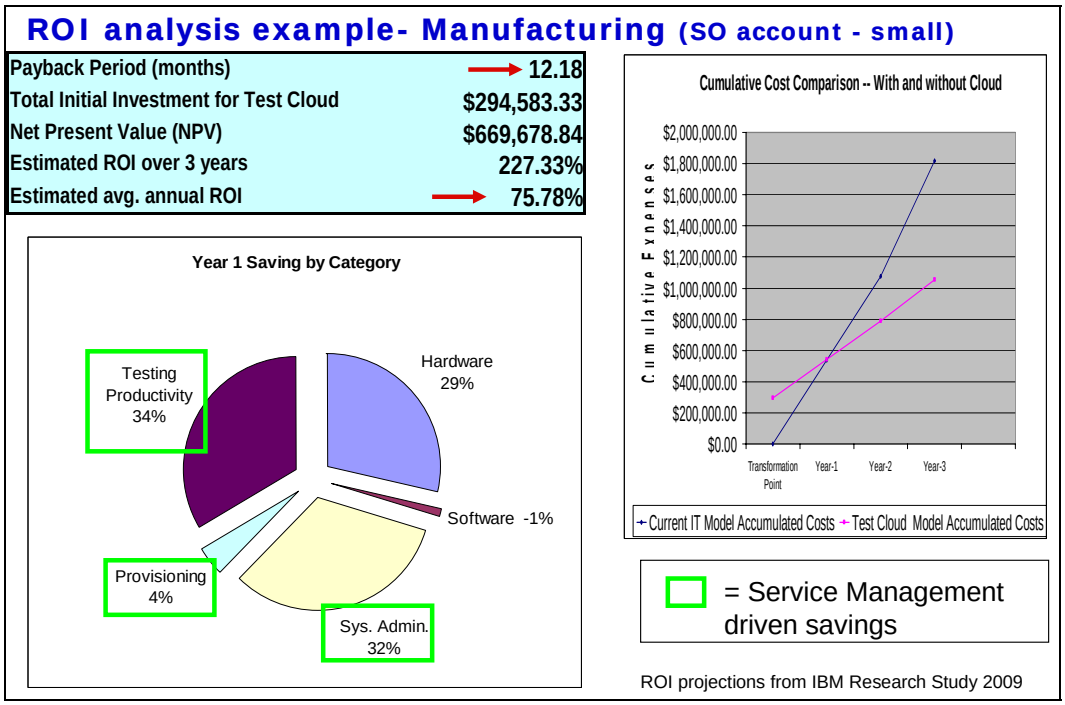

This third example is also in the banking vertical with a medium sized environment. In this sized environment, there is a sufficient number of servers to drive a balanced set of savings across the service management focus areas. The largest savings were driven by the system administration costs which saw a reduction in the number of administrators required and improvements to both testing productivity and provisioning costs.

ROI analysis example- Banking (medium # of servers)

Payback Period (months)	6.82
Total Initial Investment for Test Cloud	$302,958.33
Net Present Value (NPV)	$935,880.13
Estimated ROI over 3 years	308.91%
Estimated avg. annual ROI	102.97%

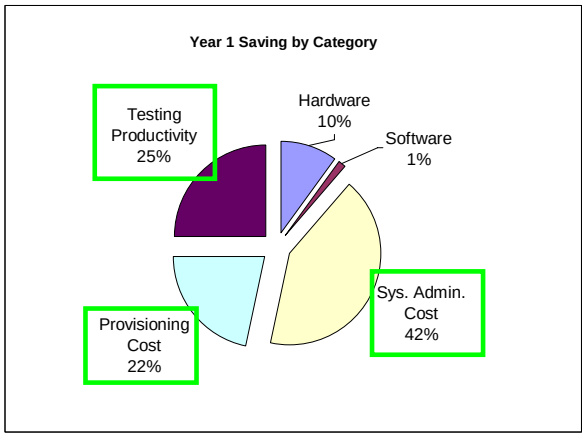

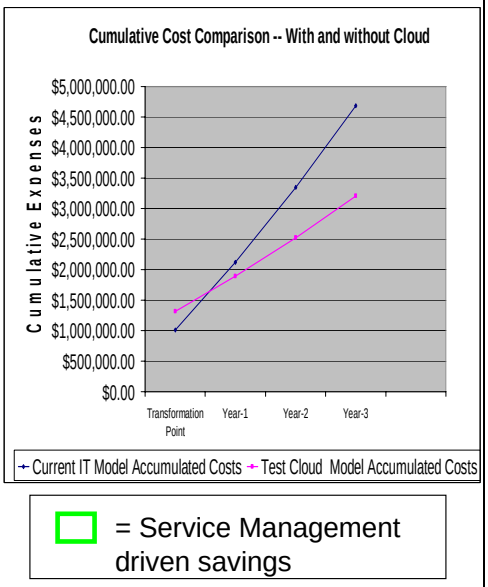

= Service Management driven savings

ROI projections from IBM Research Study 2009

4. Hardware payback

1.1 Overview

The hardware savings come from improving server utilization and decreasing the number of servers. The typical server in a datacenter is running a single application and is being utilized on average from 5% to 25% of its capacity. As systems are consolidated and virtualized in a cloud environment the number of servers required can drop significantly and the utilization of each server can be greatly increased resulting in significant savings in hardware costs today and the avoidance of future capital investment.

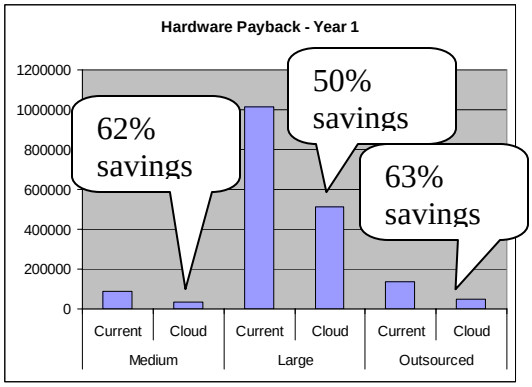

1.2 How hardware payback is achieved

There are two main areas of hardware payback. The first is physical server depreciation. Since fewer servers are required the depreciation expense can be reduced. The second area is comprised of energy and facilities costs. If there

are fewer servers using energy and requiring floor space this translates into direct bottom line savings. The graph at the right depicts these savings for each environment, in the outsourced case the savings are achieved through a reduction in physical server management charges being billed by the outsourcer. Typical savings in total hardware, energy and facilities can be in the range of 30% to 70%, based on your current size and annual spending. In the examples in this document the range is 50% to 63%. The cloud computing platform can also affect the size of the cost savings. Typically more mature platforms like the IBM zSeries can have larger savings due to their advanced virtualization and management capabilities which can support higher utilization rates.

1.3 Getting started with hardware

One of the first projects to execute is server consolidation. Server consolidation increases IT efficiency while reducing overall infrastructure costs through the elimination of underutilized servers. By moving to a virtualized and consolidated environment you can expect to see cost savings and also increased agility of the IT systems to meet rapidly changing requirements.

Whether you need help with a business case or start-to-finish consolidation of your IT environment, IBM can help you define the scope of your project and develop a road map for your server consolidation. And if you need more help, we can evaluate your workloads and infrastructure and provide a comprehensive business case, outlining total cost of ownership (TCO) and return on investment, with alternatives and a strategy that will meet the needs and challenges of your business for the next three to five years. For more advanced optimization initiatives, we can show you how to optimize your IT infrastructure using new technologies that enhance operational and cost performance by concentrating on:

- Power and cooling.
- Space utilization.
- Application performance requirements.
- Automation and virtualization.

Perhaps you have a particular focus relative to application. Whatever your needs, we can help you consolidate your workloads and upgrade databases and applications—including your customer in-house UNIX® applications—with technical expertise and services that span heterogeneous environments.

IBM also provides IBM CloudBurst which is an integrated hardware, software and services offering to deliver a complete virtualized private cloud environment.

5. Software payback

1.1 Overview

Software is the key to delivering the savings described in this guide. The two major components of software costs are virtualization software that enables consolidation and more importantly the service management software that enables visibility, control and

automation of your environment to ensure efficient service delivery. Without a service management system, the savings in the other areas can not be achieved. However, software does require a modest investment in order to delivery the large cost savings described in this guide.

1.2 Why software costs typically rise

The main drivers of increased software costs in a cloud environment are the cost of the virtualization software and the service management software to provide visibility, control and automation to the virtualized systems. Each consolidated system will require a license for virtualization software and additional service management software. These costs are partially offset by the reduction in the number of operating system licenses due to the number of systems being decreased as a result of the consolidation. As shown in the chart at the right the software costs generally increase but the overall percent increase is typically a small single digit percentage when compared to the overall savings achieved in other areas as shown in section three. In the case of the medium size banking customer, the client already had a virtualized environment established so they would be able to further reduce the number of virtualization licenses when they move to a cloud environment so their software costs would be reduced. The other two customers did not have virtualized environments at the start, so their software costs would increase.

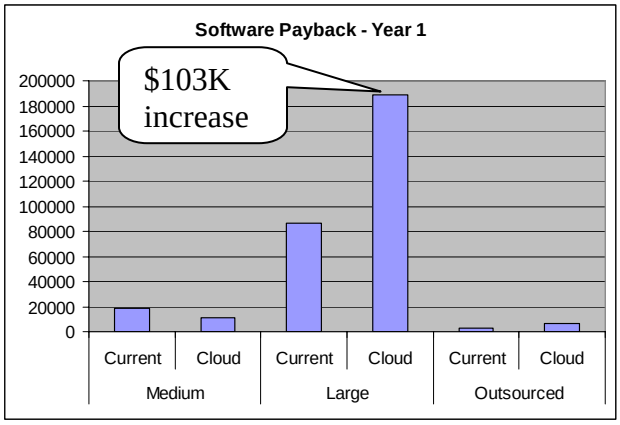

1.3 Getting started with software

To get started with software payback the first step is to install virtualization software. The virtualization software provides the basis for server consolidation and improving system utilization by running multiple workloads on a single server. The second step is to install a basic service management system for your cloud to enable the efficient operation and delivery of high quality services. The third step is to install additional service management software to focus on managing the virtualization aspect of your cloud environment and gain additional cost savings. The details of the second step are covered in the productivity payback section and the details of the third step are covered in the system administration section.

6. Automated Provisioning payback

1.1 Overview

Automated provisioning provides the ability to provision systems without the long and error prone manual process used in non-automated environments. Automated provisioning tools allow a skilled administrator to record and test provisioning scripts for deploying cloud services. These scripts can then be used by automated tools to consistently provision systems without the human error introduced through the use of manual provisioning processes. This greatly reduces the amount of time required to provision systems and allows skilled resource to be leveraged over many more systems, while improving the overall quality of the provisioning process. This provides the added benefit of supporting a policy based approach to provisioning. With a policy based approach IT, using automation, can decide when patches should be applied, what action to take when systems are no longer needed and insure that unused resources are quickly returned to the available pool to improve reuse.

1.2 How provisioning savings are achieved

The key savings metric for automated provisioning is the amount of time saved in deploying new systems which can translate into large cost savings. The chart at the right shows the savings when using automated provisioning in a medium and large cloud environment which is managed by the internal IT organization and a small out sourced cloud.

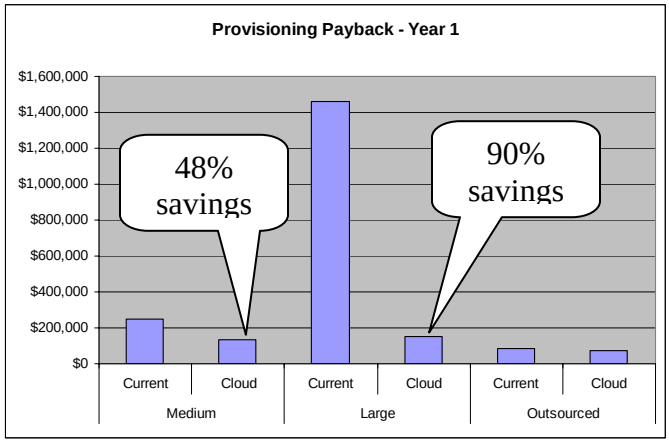

The use of automated provisioning tools drives the time required to provision each image from 40-70 hours to 30 minutes. As the number of images increases and the cost of training, deployment, administration and maintenance of the automation software is amortized over more images, the savings accelerate. In the large environment case the savings approach 90 percent.

Another important factor to consider is the use of standard images. 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. Simply stated, the higher the clone factor, the greater the reduction in labor costs associated with deploying and maintaining software virtual images.

1.3 Getting started with provisioning

To get started with provisioning requires the implementation of automated provisioning software to replace manual provisioning processes. With automated provisioning new systems and system updates can be provisioned quickly based on policies defined by IT. Automation allows skilled resource to focus on delivering new services versus mundane provisioning tasks and the use of policies multiplies the savings by insuring that IT can quickly de-provision and reuse resources.

Tivoli Provisioning Manager can help enable data centers to execute changes within the IT infrastructure more quickly, reliably and securely. Servers can be configured and provisioned from bare-metal through applications and maintained against standard hardware, software and security configurations. Computing resources can be deployed to meet current needs and then redeployed or repurposed as changes in demand occur. Tivoli Provisioning Manager supports server hardware, operating systems and virtualization technologies from a variety of vendors. Tivoli Provisioning Manager installs easily and delivers value quickly. Its intuitive interface helps users access the information and tools they need to be more productive. With pre-built or custom automation workflows, users can better manage IT resources and automate a broad array of tasks for data center environments.

Tivoli Provisioning Manager can be used to automate repetitive tasks, which can help minimize costly human errors. The software's Web Replay plug-in, for example, allows users to capture and rerun scenarios for even the most complex tasks. A skilled user can execute the individual steps involved in a procedure, save the recorded scenario and edit it as needed. The skilled user may also include help or instructional text on the captured screens. Other less skilled users can then use Tivoli Provisioning Manager to automatically execute these captured scenarios.

Tivoli Provisioning Manager is an integrated tool for provisioning resources across your environment and facilitating your data center automation initiatives. It can help automatically provision software and configurations to Windows servers and clients, as well as Linux and UNIX® physical and virtual servers. Because of the breadth of devices that it supports, its integrated discovery and inventory capabilities and its ease of use, Tivoli Provisioning Manager can be a valuable tool for driving IT operational efficiency to help minimize costs and maximize business flexibility.

7. Productivity payback

1.1 Overview

In the majority of customer environments today when new services are needed the process to get the services delivered is usually a paper based process that is not very responsive to the requestor. The new service being requested could be a development and test environment for a software developer, or it could be for a new application that the business needs to launch a new product. In the case of a developer requesting a new dev/test environment, the process usually starts with the requester filling out a paper form

and then this form is sent to various groups that need to approve the request before it gets sent to IT. Once it arrives in the IT shop they need to figure out if the necessary hardware is available and if the correct personnel are available that have the skills required to stand up the system. If it is determined that new hardware is required, lengthy delays can ensue.

With automated tools the requester chooses the environment from the services catalog which is accessed through the self service portal. Integral to the self serve portal is the service catalog that maintains an up to date list of available services and up-to-date system images. The tester specifies the start date desired, the environment, the software images and amount of time that the service is required at the time of the request. This eliminates the problems associated with environment quality and consistency as well as resources hanging in "limbo" because no one is sure when they are no longer needed. Policies can be defined to quickly reclaim systems when they are no longer being used so that less new hardware needs to be ordered and the existing hardware is used more efficiently. The self serve portal also enables automated workflows so that the necessary approvals can be quickly obtained and the automated provisioning process discussed in the previous section can be started. Environments can be stood up and activated in a matter of minutes as opposed to days or months.

1.2 How productivity savings are achieved

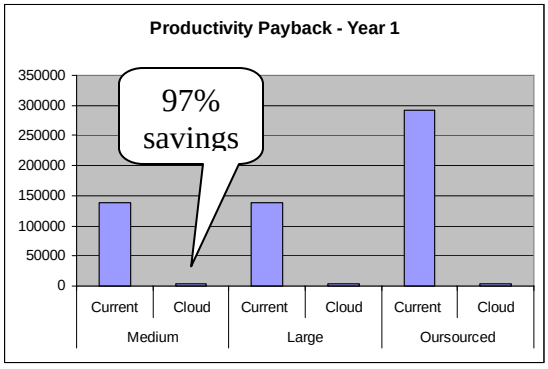

The key savings metric for productivity payback is the amount of time a requester wastes waiting for the resources they need to do their work. The chart at the right shows the savings when using a self serve portal, service catalog and automated provisioning to automate the service request process. The use of automation drives the idle time per project from tens of hours to 1 hour. In each case significant savings are achieved. Accelerating the completion of work has a dramatic impact on the agility of an IT organization to respond to the demands of the business and launch new services faster. This shortens time-to-market and fuels more innovation for revenue generation. These are key concerns of business executives across the globe. The financial impacts of these business benefits can be huge, but they were not studied as part of the current research.

1.3 Productivity

To get started with productivity improvements you should implement two projects; a user friendly self-service interface that accelerates time to value and a service catalog which enables standards that drive consistent service delivery. These two projects will enable end users to request the services they need from a service catalog and provide an automated approval process to enable the rapid delivery of new services.

Businesses that manually deploy and manage applications and IT services face a number of challenges. First, it can take weeks or months to receive the hardware and software needed to deploy many applications. Next, the cost associated with the wait for needed equipment as well as the cost to provide highly skilled IT staff members to manage the deployment process can be substantial. Lastly, significant time and financial resources are required to help ensure that the organization has audit processes integrated with process governance.

Tivoli Service Automation Manager provides adaptable and automated best practices for building and managing IT services. IT providers use templates to define service offerings that are available to line-of-business managers. Through these templates, IT providers take computing resources, such as virtualized operating systems and application middleware stacks, integrated with workflow processes and standardized configurations, and make them available as offerings to business operations staff members. Requestors using this one-stop interface do not need detailed information regarding the computing resource's requirements. They simply use the templates and pre-defined content to request the resources from IT. In turn, IT can provide these resources in a very short time, enabling users to respond rapidly to changing business scenarios.

Provisioning of services is automatically enabled without requiring system or database administrator involvement. Because these processes are automated and standardized, they make IT deployment repeatable, reliable and cost-efficient. In addition to lower server-to-administrator ratios and lower no-touch deployment, automated delivery of IT services also enables better server and energy utilization.

Tivoli Service Automation Manager and a cloud computing model facilitate standardization of the service request process. Through this standardization, service consumers gain a consistent and reliable quality of service when requesting and receiving resources from IT. This greatly improves the value of IT to the business. When a service is delivered, Tivoli Service Automation Manager manages the operational lifecycle components of that service. For instance, IT administrators can automatically implement monitoring and event management functions as well as recommended changes to deployment. They can also manage changes to deployment processes according to results obtained from monitoring statistics.

In order to accelerate the delivery of a managed cloud environment IBM also offers IBM CloudBurst. IBM CloudBurst is an integrated service management offering with network, servers, storage, quickstart services, and financing that delivers a pre-built cloud. IBM CloudBurst saves the time required to install and configure the hardware and software required to deliver cloud based services which greatly accelerates time to value.

Enhancing IBM CloudBurst with the IBM WebSphere CloudBurst Appliance provides a ready-to-go, cloud computing platform with specialized images for creating and deploying application environments based on IBM WebSphere software. The combination of the two offerings leverages the strengths of each to accelerate the delivery of new WebSphere based applications. IBM CloudBurst is a complete cloud computing

package that can deliver any type of virtual machine in support of multiple types of cloud workloads. The WebSphere CloudBurst Appliance stores and secures application environments consisting of patterns of IBM-supported WebSphere Application Server Hypervisor Edition virtual images, these environments can be dispensed into, managed and maintained in a private cloud. So a customer can deploy IBM CloudBurst to quickly deliver a cloud computing infrastructure and use the WebSphere CloudBurst Appliance to deliver the WebSphere based application environment into the IBM CloudBurst cloud.

WebSphere CloudBurst Appliance has been shown to dramatically reduce costs through improved hardware utilization providing for reduced power consumption, efficient license management and decreased labor costs. Additionally, WebSphere CloudBurst allows for faster time to value through reliable, consistent and repeatable dispensing/teardown of application environments, which can circumvent lengthy approval processes and improve upon security compliance. A hands on study was conducted to track the time it took to deploy and instantiate a WebSphere-based application on a virtual server using VMware. Metrics were captured based on a manual deployment and using the WebSphere CloudBurst Appliance. 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%. IBM offers a complimentary WebSphere CloudBurst Value Assessment to highlight the rapid ROI it can provide. The Value Assessment uses a flexible and proven financial model to capture and compare current and anticipated costs vs. benefits over a multi-year period. Use of the cost model allows you to better understand cost components, evaluate alternatives and develop project plans with greater confidence.

8. System Administration payback

1.1 Overview

As a cloud environment is implemented, hardware is consolidated and systems are virtualized to drive savings. However, one area that requires particular focus to ensure that costs don't actually increase is the area of system administration. In a cloud environment there are less physical servers to manage, but the number of virtual servers increases. Virtual systems are more complex to administer and this can lead to higher administrative costs. That is one of the reasons that a strong service management system is so critical to an efficient and effective cloud environment.

1.2 How system administration payback is achieved

The key savings metric for system administration payback is the efficiencies that can be gained by effectively managing the virtualized aspect of the cloud environment. With the proper service management tools for a highly virtualized environment, the administration savings can range

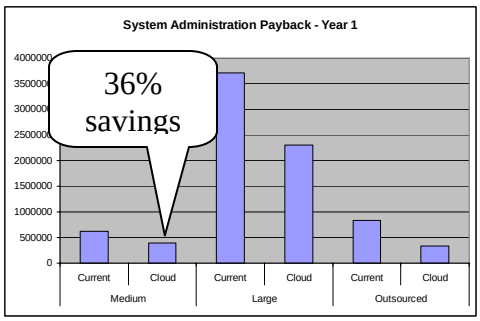

from 36% to 45%. The chart at the right shows the savings when using service management tools designed for virtualized systems.

1.3 Getting started with system administration

The automated provisioning of services using a self server portal and service catalog is a key first step as has been described. To achieve even high levels of savings and greater efficiencies three additional service management steps or "projects" should be implemented. The first project would map the cloud computing environment so that the relationship of the physical and virtual assets is understood so that reported problems and potential problems can be quickly isolated and remediated. The second project addresses the performance monitoring and management of the virtual servers so that trends can be established and watched so actions can be taken proactively. The third project is to manage the cloud at the services layer and ensure that the services are meeting their service level objectives. Let's look at each project in more detail.

As the use of virtualization continues to accelerate, organizations need to adapt their existing tools in order to fully exploit both virtual and physical resources. Many of the resources that work together to comprise a business service end up residing on heterogeneous and virtual environments requiring advanced capabilities and tools to manage the complex dependencies and relationships. IBM Tivoli software offers a range of service management solutions designed to help organizations better understand their virtualized environment — how to map virtual resources to physical resources, automate the various operations, calculate usage, and manage performance and availability.

As mentioned earlier, Tivoli Service Automation Manger provides the foundation for driving system administration savings by enabling standards and policies that drive efficiencies. In addition, Tivoli Service Automation Manager can manage changes to systems automatically based on monitoring data and provide audit trails of all action taken to improve governance. IBM also provides other system management software for virtualized environments that can drive even higher levels of savings.

Being able to visually trace the physical and virtual topology of the servers, associated infrastructure and composite workloads is critical. A high-level mapping of your virtual and physical servers and the workloads running across them provides a more complete understanding of both the physical and virtual infrastructure and the relationships between them — what physical components are used to build a virtual component and, equally, what virtual components share a given physical component. IBM Tivoli Application Dependency Discovery Manager provides visibility of interdependencies between applications, computer systems and networking devices, using agent-less, credential-free discovery and automated application maps.

Successful management of a virtualized environment relies on a comprehensive and cohesive view of the operational state of the heterogeneous resources that comprise the IT data center. Monitoring and visualization tools are a vital part of managing a virtualized environment, enabling organizations to better understand the current state of events to help predict potential system failures and bottlenecks. If a problem does occur,

the tools can help organizations rapidly isolate and analyze problems before service levels are threatened. IBM Tivoli Monitoring for Virtual Servers provides superior visualization to help identify trends, see impacts and take action, all from a centralized, easy-to-use portal.

Tivoli monitoring solutions also allow you to collect both traditional IT measurements and environmental measurements into a common dashboard that displays an integrated view of power usage, thermal data and application performance metrics. From there, you can feed the temperature data into data warehouses to share with other applications for real-time and historical trending analysis. For example, you might home in on a particular asset that consistently uses more energy than the others. By correlating this information with service level targets, you can take immediate action to slow the system down or move workloads between virtual machines without impacting SLAs.

You can also set policies to autonomously respond to certain preset thresholds or events to decrease heat generation. As workload increases and decreases, these automated power management policies can continually adjust the power by metering, controlling or capping consumption to save energy while maintaining response times.

IBM Tivoli Business Service Manager delivers advanced real-time visualization of services and processes to help organizations monitor and measure service availability and performance against defined revenue, growth and operational objectives. Detailed service dependency views can speed problem resolution through automated service-impact and root-cause analysis. In addition, Tivoli Business Service Manager offers service intelligence through real-time service scorecards that track key performance indicators (KPIs) and SLAs. Composite applications — which use business logic and data that span physical and virtual servers, integration middleware and mainframe systems — can be particularly challenging to monitor. IBM Tivoli Composite Application Manager (ITCAM) can help you avoid critical performance problems by proactively recognizing and isolating problems using robotic and real-time techniques. An intuitive user interface for powerful visualization and customizable views enable quick determination of degraded performance levels to help maintain customer satisfaction.

Virtualization goes a long way to helping organizations reach the goal of providing rapid near real-time response to customer requirements. When virtualization combines with automation, true flexibility, resiliency and responsiveness become a reality. For physical and virtual systems alike, the automation of critical operations is crucial for ensuring availability, providing efficient service and reducing manual errors, cost and complexity. Virtualization goes a long way to helping organizations reach the goal of providing rapid near real-time response to customer requirements. When automation is not widely applied to the tasks and workflows associated with processes such as deployment, fulfillment, and change and release management, operational efficiency and effectiveness are difficult to achieve. Tivoli software offers automation solutions designed to address key types of automation, including provisioning, workload, storage and physical automation.

9. Summary of savings and costs

The following table provides a summary of the savings in each of the five areas and the associated costs.

Area	Saving Metrics	Cost Metrics
Hardware	o Reduction in number of servers o Drives reduction in server depreciation cost, energy usage and facility costs	
Software	Reduction in the number of OS licenses	o Cost of virtualization software o Cost of cloud management software
Automated Provisioning	Reduction in number of hours per provisioning task	Training, deployment, administration and maintenance cost for automation software
Productivity	Reduction in number of hours waiting for images per project	
System Administration	Improved productivity of administration and support staff (support more systems per administrator)	

10. Tracking your Payback/ROI - What Data do you need to measure?

In order to determine your cloud computing payback, the first step is to gather the data required to establish a baseline for you pre-cloud environment. As the projects defined in this guide are executed, the savings can be measured against the baseline data to calculate your payback. Some estimates may need to be made for some of the data and as processes are put in place to better capture the data required more accurate results can be achieved. However, the lack of exact detailed data for each measurement should not inhibit you from getting started. Where exact data cannot yet be measured or calculated, make reasonable and logical estimates. This initial “Benchmark” is an important component of the overall program. As the results are tracked, data collection improves and the savings are achieved, the benchmark will provide insights to track against. The benchmark will also make it easier to justify the investments required to get started. Savvy IT executives have been known to use their savings projections to make the ongoing project "self funding". In those cases, there is an agreement in place with the finance team or CFO that some percentage of the annual savings will be returned to the IT organization to fund next years investments in further improvements. CFOs and IT CXO's love that concept of "self-funding"

1.1 Hardware metrics

- Number of existing physical servers
- Average total purchase price per server
- Average Hardware % utilization

1.2 Software metrics

- Number of existing non-Free OS (e.g. Windows Server) licenses
- Number of existing virtualization software licenses
- Annual virtualization software license cost per physical server
- Cost of non-free OS license (one time cost) per physical server
- Cost of non-free OS annual license maintenance per physical server

1.3 Provisioning metrics

- Number of images installed per year (provisioning requests)
- FTE hours required per provisioning without automation software
- FTE hours required per provisioning with automation software
- Costs for provisioning automation software (Tivoli Service Automation Manager)

1.4 Productivity metrics

- Total number of projects per year
- Average idle time caused by provisioning per project without automation
- Average idle time caused by provisioning per project with automation

1.5 System Administration metrics

- Number of existing physical servers
- Number of actively managed OS instances
- How many system administrators are managing the environment (excluding network and storage administration)
- Administration productivity gain due to cloud transformation
- Ratio of system administrators to support staff
- Annual labor cost per system administrator

11. For More Information

IBM offers a complimentary WebSphere CloudBurst Value Assessment workshop to organizations who are actively seeking to optimize their application server environment, as well as, improve upon current provisioning times and related costs. The assessment can be conducted remotely with IBM subject matter experts or on-site. It provides an improved understanding of your current application server environment, along with the financial and operational performance benefits resulting from using WebSphere CloudBurst Appliance. It also provides a business case for WebSphere CloudBurst deployment that's tuned to your business priorities and growth plans in support of project funding and budgeting. The WebSphere CloudBurst Value Assessment consists of:

- Review assessment objectives, agenda and deliverables

- Validate the current environment cost structure, growth plan data and other important factors
- Develop the WebSphere CloudBurst high-level architecture
- Understand administrative/business impact requirements and any related assumptions
- Finalize the financial model
- Develop and deliver executive presentation

Conducting a WebSphere CloudBurst Value Assessment can provide you with an improved understanding of your current application server environment, along with the financial and operational performance benefits resulting from WebSphere CloudBurst Appliance. It can also be the basis of a business case for WebSphere CloudBurst deployment that's tuned to your business priorities and growth plans in support of project funding and budgeting

Contact your IBM WebSphere account executive to learn more and schedule a value assessment.

IBM Corporation Software Group
Route 100
Somers, NY 10589
Produced in the United States of America
November 2009

DIW03009USEN-00