Case Study:
Wipro's Pragmatic Use of Cloud Services

Gartner RAS Core Research Note G00174809, Frances Karamouzis, 15 April 2010, RA3 10162010

Wipro Technologies is a large IT services provider actively developing and piloting several cloud service offerings. Wipro embarked on a journey to practically test one of the potential solutions on its own internal mission-critical processes. Here, we present a case study on its internal execution of cloud services.

Key Findings

- Prior to launching the cloud services implementation, a definitive business case for savings will likely continue to be elusive. Wipro is a highly "metrics driven" culture, yet it struggled to quantify the business case in advance.
- One of the most surprising findings for Wipro is that, while significant savings were achieved, the most important result was not cost savings. The most important result was business-focused agility (speed of provisioning requests) and elasticity (ramp capacity requirements up or down on demand).
- The IT process (for provisioning of infrastructure) averaged 46 elapsed days prior to the use of cloud services, and the newly deployed process is 35 minutes.

Recommendations

- Sourcing managers should proactively evaluate some of the new industrialized IT services offerings, such as cloud services. Given the dynamic environment, it is a strategic imperative that CIOs and sourcing managers evaluate alternative delivery and acquisition models by applying Gartner's multisourcing best practices. Align sourcing decisions with the business's appetite for risk.
- Given some of the emerging technology involved, a proof-of-concept effort is imperative.
- CIOs and sourcing managers should use Wipro's lessons learned to explore new, industrialized IT services offerings (such as cloud services). There is potential to leverage the benefits far beyond some of the tactical drivers of traditional outsourcing, where cost savings often factor heavily in the sourcing decision.

WHAT YOU NEED TO KNOW

Wipro's implementation of a cloud services solution for its internal IT process for infrastructure provisioning reduced the average elapsed time from 46 days to 35 minutes using the "as a service" delivery model via a self-service portal environment.

Gartner

CASE STUDY

Introduction

Wipro Technologies is one of the largest IT services providers in India, with more than 100,000 employees. It delivered approximately US$4.3 billion in revenue in fiscal 2009. Wipro's core business is to provide application services (application development, maintenance and management), business process outsourcing (BPO) and infrastructure outsourcing to enterprise clients in 55 countries. For the purposes of this case study, one of the most critical business and IT processes behind the scenes within Wipro is the provisioning of infrastructure. The reason this internal process is important is that, on average, Wipro has 4,000 projects (distinct work efforts) being executed simultaneously for more than 700 clients. The teams are spread across multiple locations and have some of the most dynamic requirements for services, software and hardware. Additionally, various internal and client-facing teams are constantly piloting all sorts of solutions, prototypes for business development and R&D. Thus, the lifeblood of the company is the ability to quickly provision infrastructure to thousands of associates who are busy delivering solutions to customers and prospects. As such, while the provisioning of infrastructure is a behind-the-scenes process, it is a mission-critical step for an IT services provider like Wipro, where the underlying structure is essential for a consistent and predictable delivery of service levels to its clients.

The Challenge

The constant challenge for Wipro staff supporting associates who work directly with external clients and prospects is to get them the right technology configuration (servers, sizing, security and so on) in the shortest possible time frame. The Wipro staff team supporting this infrastructure faced a number of unique challenges that included the following:

- **Delays in provisioning hardware:** There was a lead time of about 46 days in addressing the infrastructure requests from the day of request to the time when the hardware is actually available in a usable state.
- **Ineffective utilization of hardware:** The average utilization of the servers was less than 10%, and the servers were sprawled across multiple locations. The spare capacity could not be used for new or alternative teams.
- **Support:** Distribution of the centers of excellence (CoEs) and solution teams across various locations made it difficult to provide the right level of support with a reasonable staffing level.
- **Diversity of technology:** Given the extensive Wipro client base, there is an extensive level of heterogeneous technology required. The teams' requirements had a high diversity of operating systems, platforms, tools and hardware.
- **Process management:** Distributed resources across multiple locations made it difficult to manage the lab infrastructure centrally, due to the dynamic aspects of the infrastructure.
- **Loss of data/artifacts:** End users found it difficult to back up critical files as server failures led to frequent disconnections from the backup network.
- **Ramp up and ramp down:** Project teams were not able to scale the computing environments up or down based on application workloads.
- **Short-term pilots:** Many project teams required high computing capacity for a short period of time, and it was not feasible to procure hardware and software for such projects.

Here, we summarize the overall process (see Figure 1). The process documented below is the "before" (prior to) the use of cloud services. This process normally required an average elapsed time of 46 days. Over the many years of work and process improvement efforts, many of the Wipro staff spent lots of time trying to make this as efficient and effective as possible. Nevertheless, there was some level of latency in the process.

Thus, the key business challenges were to increase the speed, reduce the cost of the provisioning of the infrastructure and ensure effective utilization.

Approach

The solution to these issues was to build a private cloud with capabilities of self-service provisioning to:

- Reduce process overheads.
- Support multiple types of virtualization technologies to support the needs of diverse operating systems and software.
- Support a wide range of existing and new hardware in the infrastructure.
- Enable effective management of assets, such as software licenses.
- Implement chargeback based on resource usage.
- Integrate with enterprise IT management tools for trouble ticketing, support and management.

Figure 1. Process Summary

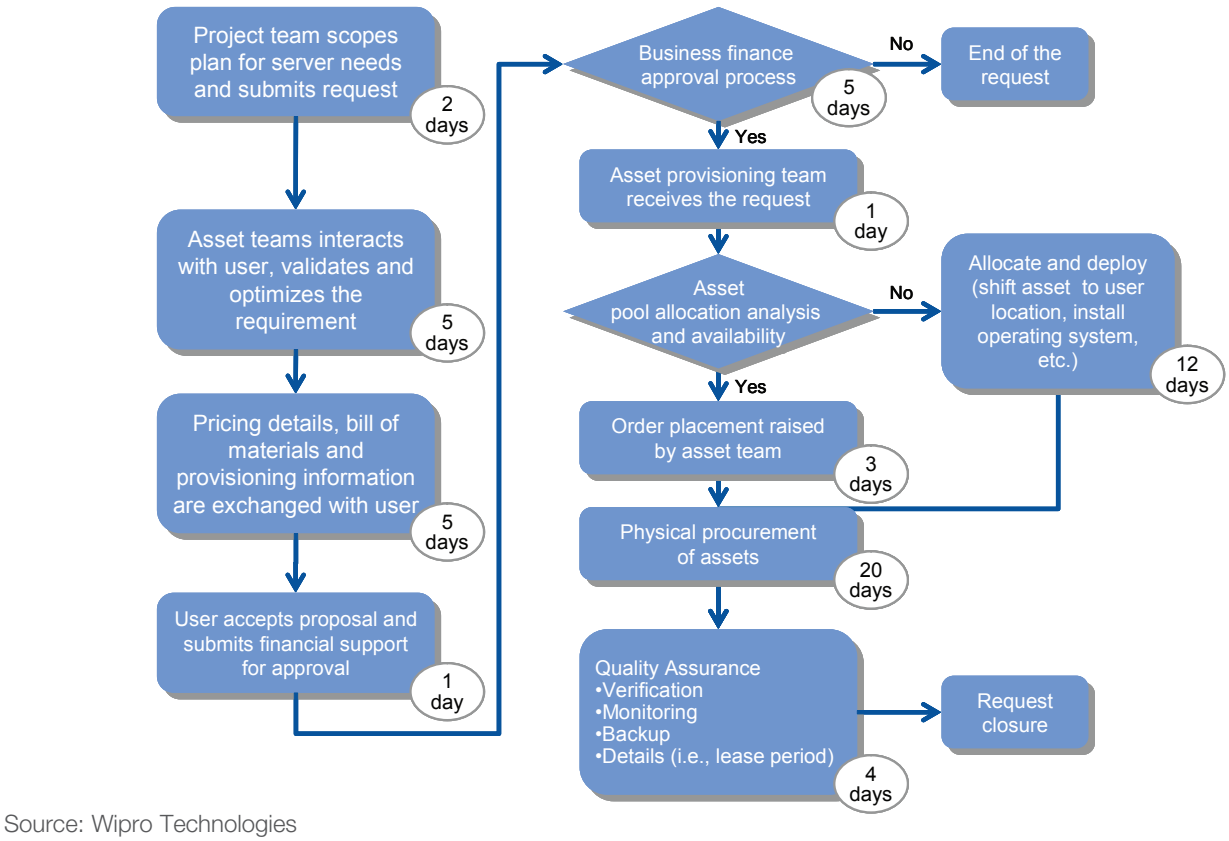

Source: Wipro Technologies

To execute, Wipro approached the project in the following manner:

Step 1 — Cloud project office and solution team incubation of a grid computing CoE

The Private Cloud project was initiated by Wipro's CTO office, with a core team of 14 engineers and expertise in virtualization, grid computing, security, networking infrastructure and Java Platform, Enterprise Edition. This team started identifying the technologies and developing the private cloud solution.

Step 2 — Evaluate cloud technology

Available products did not meet the requirements of managing multiple virtualization technologies, the ability to handle enterprise workloads, integration with enterprise IT management systems, implementing reverse chargeback and so on. Thus, the CTO office decided to build a private cloud solution using off-the-shelf components as much as possible and integrating/customizing open-source components to complete the stack.

Step 3 —Proof of concept

Given some of the emerging technology involved, a proof-of-concept effort was initiated to finalize the various elements of architecture and procure fully fledged equipment. The proof of concept took about three months.

Step 4 — Development of the Private Cloud solution

The solution development took more than a year (including incubation). The project was executed using the agile methodology with the objective of creating a productized solution, which can be used internally, as well as an infrastructure service for Wipro customers.

Step 5 — Deployment of Private Cloud for the test and development environment

In the first phase, 160 virtual servers were deployed and capacity was provisioned for the CoE and innovation teams.

Step 6 — Deploy internal IT applications

Post the successful deployment of the test and development environments. IT applications were deployed to test production worthiness and reliability.

Step 7 — Redesign CIO organizational processes

With some of the specific cloud technologies evolving in maturity and stability, coupled with the need for widespread deployment, the internal processes of the CIO organization were modified after a week-long Kaizen event to analyze the current process and optimize future processes.

Step 8 — Deployment across all locations

The Private Cloud solution is now a standardized platform for Wipro and is being rolled out across all locations. Currently, more than 1,000 servers are available in the Private Cloud, and there are three locations with Private Clouds.

Results

In summary, by employing cloud services and moving to the Private Cloud infrastructure provisioning, time has been reduced from an average of 46 elapsed days to 35 minutes using the "as a service" or "on demand"-type of provisioning via a self-service portal environment (see Figure 2).

The Private Cloud increased flexibility in allocation, as well as allowed for additional provisioning based on the usage patterns. This enabled an increase of the overall utilization levels for the hardware from less than 10% to more than 40%. The Private Cloud allows IT administrators to monitor the utilization and forecast of capacity. Additional hardware capacity is added before the available capacity falls below 20%. Dense packing of virtual machines on physical servers helped reduce the new server procurement costs by 30% in the first year of operation, with some additional capacity to spare.

In this schema, the Private Cloud is able to guarantee on-demand provisioning of virtual servers by maintaining additional small capacity. As an added benefit of consolidation, and better management of server power and cooling, it will save more than 30% in electricity costs.

The availability of a Private Cloud has made several innovations possible. The applied research group within the CTO organization has been able to incubate an e-discovery service based on private clouds. Several business units have started developing industry-specific applications as cloud services and are piloting on the Private Cloud.

Satisfaction levels from users of Wipro cloud services increased significantly due to its simplicity, self-service capability and management services, which include additional new services that were not available in the dedicated infrastructure model. More teams are now attracted to moving to Private Cloud infrastructure because of the benefits of cloud, other than the cost saving it offers.

Beyond the measurable business results depicted above, it was interesting to note that, prior to starting the project, it was difficult to determine a definitive business case for savings. It was extremely difficult to project the changes in behavior of such a diverse number of teams that tend to estimate conservatively and provide a great deal of cushion to absolutely ensure that deliverables for clients are completed on time, on budget and in scope. A number of findings from an analysis of the before and after include:

- Project teams wanted to have their own computer infrastructure and feared the loss of control of the timing or budgeting of the work effort. Most teams' behavior was defined by a determination that internal infrastructure provisioning would not result in the loss of valuable time on the project plan.

- Teams had concerns about the feasibility of using the cloud for their applications.

Figure 2. Summary of Benefits of Cloud Services Applied to Process for Infrastructure Provisioning

Self service portal for automation of server provisioning

Parameters	Before	After
Server Provision Cycle	• 46 days	• 35 minutes
Average Server Cost	• $2,000	• $800
Utilization	• <10%	• 40%
Reprovisioning of Assets	• Nil	• Reduce capex 30%
Silos	• Physical machines and application silos	• Removed silos with >500 virtual machines
License Management	• Highly inefficient	• Effective
Network Bandwidth	• Increased network bandwidth	• 1 network interface card cord for 32 servers
Energy Cost	• High	• Lower

Source: Wipro Technologies

- The consolidated upfront investment required to set up the cloud was very high, and none of the business units would individually step up to make such a high investment. The CTO and the CIO office invested hard to take a leadership role in this initiative.

One of the most surprising findings for Wipro is that, while significant savings were achieved, the most important business result was not cost reduction. The most important result was agility (the ability to respond to project team requests) and elasticity (ramp capacity requirements up or down on demand). Wipro teams can now service project requests faster with significantly higher levels of flexibility. The downstream impact is that the teams are provisioning fewer servers because they no longer "overestimate" or "plan for peak" (a conservative approach to provisioning, especially since in the past, it may have taken 46 days to get more infrastructure resources). Given that it now takes 35 minutes, and teams know that they can place a request and get the resources in 35 minutes, they no longer feel the need to overestimate. The result is that they only provision what they use.

Critical Success Factors

The critical success factors for the ongoing execution of these cloud services include:

- Automation of provisioning through a self-service portal
- Consistent and predictable provisioning on demand within 35 minutes
- Support for multiple virtualization technologies
- Ability to provision large capacities for short-term projects
- Catalog of standard application development and test environments
- Ability to execute a mix of transactional, collaborative and analytical applications
- Effective license management of software development tools
- Predictable, fair and measurable chargeback mechanism for business-unit budgeting

Lessons Learned

Through the proof of concept, followed by the extensive usage (more than six months) of the new provisioning process, the following is a summary of the lessons learned:

- Similar to traditional services, cloud services still require a level of hand-holding for service recipients. There was a wide disparity of comfort with the self-service portal. Thus, the help desk facility was still a critical requirement.
- Users would want various predefined configurations of infrastructure, as well as a catalog of software infrastructure components (operating systems, application servers, databases and so on). This optimizes the license management as well.
- Chargeback for cloud services was not designated as an initial priority. However, without a chargeback model, the result was a sprawl of servers on the cloud with very low utilization. The publication of data and information regarding charges dramatically altered behavior and resulted in users being more judicious about the use of the cloud services.
- Capacity gets consumed when there is a lack of time and capacity monitoring or metering; therefore, planning and forecasting become very crucial processes. Overprovisioning is the most important factor in getting a high return on investment when cloud utilization is low.
- It is not sufficient to just integrate IT infrastructure and provide a Private Cloud service. It is essential to integrate the processes of capacity planning, resource allocation, support, security and disaster recovery.
- Physical infrastructure should scale horizontally; vertical scaling of individual components is cost-prohibitive. It proved better to go with entry-level components scaling horizontally, rather than a single large component.
- Allocation of infrastructure into predefined units of small, midsize and large configurations with an automatic scale-up facility helped the end users' sizing process (at times, end users would be able to determine requirements in an open-ended structure).
- 24/7 support was needed, even for a development and testing environment, as a centralized infrastructure gets used by users around the globe.
- Support for multiple hypervisor technologies was critical. Different project teams working for customers would need different environments, according to customer preference.
- As the density of the VMS increased, it was determined that storage becomes a point of performance bottleneck. The Private Cloud failed initially with extremely slow storage performance as density increased. Integration of storage and computing needs is necessary to factor in application input/output (I/O) requirements, volume of data transferred and response times. Higher disk drive speeds, higher I/O operations per second (IOPS) and solid-state drives for caching are essential.