January, 2012

Six Steps to Private Cloud Adoption

Key Considerations, Insights and Recommendations

By Shantanu Paknikar
Happiest Minds, IT Services Group

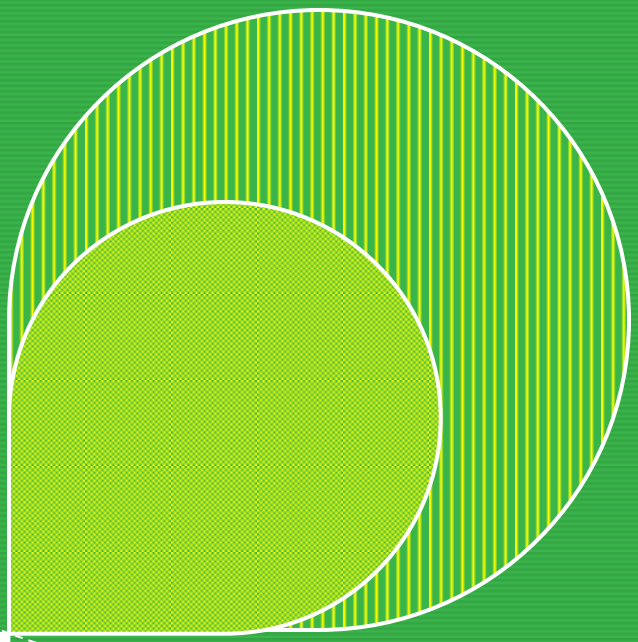

Six Steps to Private Cloud Adoption

In the very well-known book "The Tipping Point" by Malcolm Gladwell, the author defines a tipping point as a moment of critical mass, a threshold, a boiling point. The author also draws a parallel between the spread of certain ideas, products or concepts and the spread of viruses and epidemics – once they are past the tipping point. While the concept of *Cloud Computing* is not completely new, it certainly appears to have reached the tipping point.

The world over, across industry segments, across business and IT organizations, cloud computing is now being considered as one of the game changers. Organizations the world over have started moving to, or are planning to test the waters for the cloud. Cloud adoption appears to have acquired epidemic proportions.

Some time back, the "cloud" usually meant the public cloud scenario – software being delivered as services over the internet. Today however, there is an understanding that we could also have private (on-premise) or hybrid (combination of public and private) clouds. Private cloud implementations are happening and will continue to happen – for the simple reason that implementing them makes good business sense for large enterprises with significant investments in their own data centers, applications and infrastructure[i].

The six step approach described in this paper is primarily focused towards a private cloud scenario.

Business Overview

The overall business drivers for cloud adoption are straightforward and well known:

1. **Cost Optimization** – due to the increased utilization of resources [hardware and software] as a direct consequence of resource sharing. Optimization also results from the lower maintenance and upgrade costs that on-premise software would typically require.
2. **Increased Agility** – due to the concept of 'elasticity' or the rapid scale up and scale down of resources that is made possible by the cloud. With such agility, businesses can adapt rapidly to changes in their environment: For example, online retailers who see transaction volume spikes during the November – December months in the USA.
3. **Better Cash Flow Management** – due to the conversion of capital expenditure [e.g. software licenses, server and storage infrastructure] to operational expenditure [e.g. monthly rental].

A detailed discussion on the above is outside the scope of this paper; however there are some excellent references available for the same[ii].

The figure below illustrates the future (in some cases, present) scenario for a typical enterprise's private cloud implementation. The layers for infrastructure (IaaS), platforms (PaaS), Software (SaaS) and Business Processes (BPaaS) are well known from the "As a Service" terminology perspective.

Besides the IaaS, PaaS, SaaS and BPaaS layers, private cloud implementations will need to have layers for cloud security, cloud management, cloud integration, and cloud consumption. Cloud integration will need to happen both with on-premise applications that may not be moved to the cloud, as well as with public cloud applications.

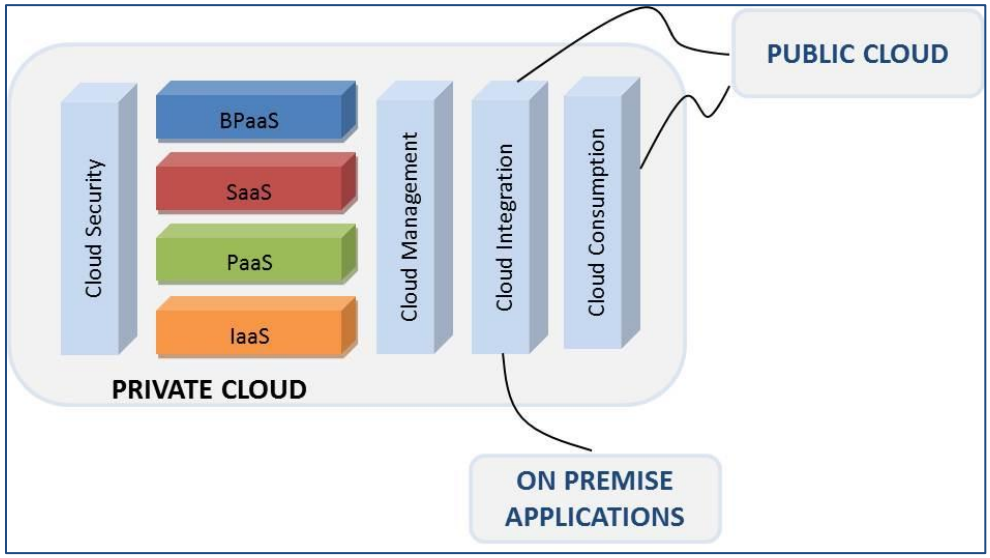

Figure 1: Cloud Scenario - Pictorial Depiction

One of the fundamental questions to be addressed by most organizations is their readiness for cloud adoption. This paper describes six steps and key considerations for each that would help organizations during the initial strategy and roadmap definition phases for private cloud adoption. These six steps are depicted in figure 2 below.

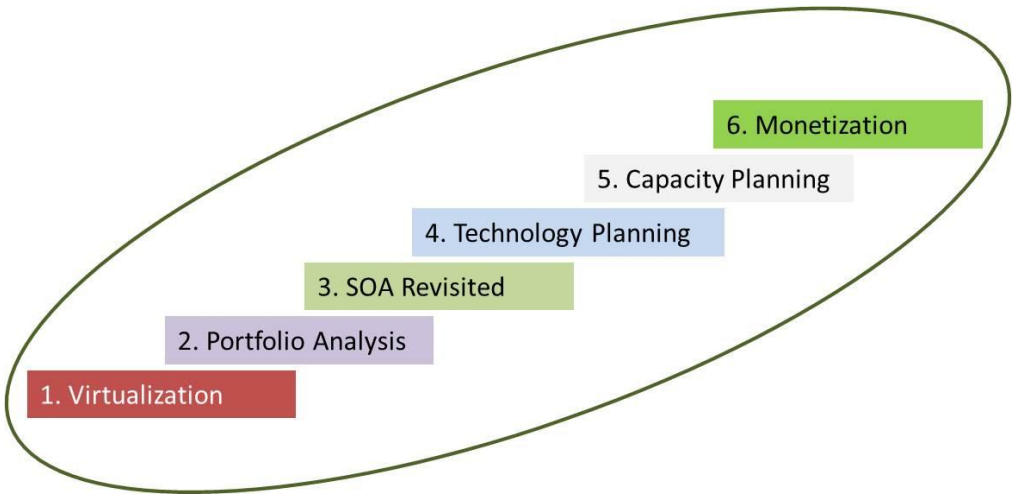

Figure 2: Six Steps to Private Cloud

Step 1: Virtualization

Virtualization Overview

The concept of Virtualization can be considered as the trigger or catalyst for the entire phenomenon of cloud computing – to the extent that a lot of people believe that the terms virtualization and cloud can be used interchangeably. While cloud computing today is much more than virtualization, the fact is that without virtualization there would have been no business case for cloud computing anyway.

Virtualization can be defined simply as "The conversion of physical assets into virtual ones". Assets could be hardware (such as computers, servers or disks), network, and software applications. Usually however, the initial and primary focus of virtualization is conversion of physical infrastructure (for example, desktops, servers, storage) to a virtual form. In particular, the focus in this article is on virtualization in the context of physical machines being converted to virtual machines.

The Business Case for Virtualization

The business case for virtualization is simple: over the years, data collected shows that hardware infrastructure has been grossly under-utilized, averaging around 30% of maximum capacity[iii]. With virtualization, each physical server can be converted (in this case for example) into three virtual ones, moving the overall utilization to around 90%. Equally important is the fact that virtual machines can be provisioned and de-provisioned on demand. Consequently, the same physical infrastructure can be shared across multiple customers, by the allocation of virtual machines as and when needed, and the de-allocation when not. *This results in the ability to charge the customer based on usage* – one of the cornerstones for the concept of cloud computing.

Key Considerations

The important factor to consider for virtualization is the analysis and selection of a software vendor. Today there are a large number of vendors to choose from in the market. The primary piece of software that enables the conversion of physical machines into virtual machines is the *hypervisor.* Therefore, vendor selection criteria include factors such as hypervisor robustness and performance, virtual machine density, supported platforms, and enterprise architecture suitability. Besides the hypervisor, there are usually other associated products offered by the vendor. The scale, number, and maturity of these products can also be used as additional vendor selection criteria.

Output

The output of this step is virtualized hardware or in the advanced case, a full-fledged Infrastructure-as-a-Service (IaaS) implementation.

Planning for the Future

The future is fairly easy to predict – it is very likely that (almost) every desktop and server machine in the enterprise will be a virtual machine, being accessed through mobile devices

such as netbooks, tablet computers and smart phones. *Therefore, the aim should always be for near 100% virtualization.* There needs to be a clear plan and roadmap for this with execution closely tracked.

Step 2: Portfolio Analysis – To migrate, or not?

Once the cloud adoption decision is taken, the next step after virtualization is in the area of *Portfolio Analysis* for the IT application landscape. The primary objective of portfolio analysis would be, at an application level, to decide on the migration strategy. Portfolio Analysis in itself is not a new area, being for example an integral part of approaches to Application Portfolio Rationalization initiatives which large IT organizations often execute. In this case, when the destination is a cloud, however, there are some unique aspects which need to be kept in mind.

Unique Aspects during Portfolio Analysis

1. Some applications will need to move to the *private cloud* in a "Lift and Shift" manner, with the applications themselves remaining largely untouched. Typically, these are likely to be custom applications with business dependencies such that any re-engineering would be detrimental to business interest.
2. Some applications or application functionality will need to move to the *public cloud*. There could be two scenarios here:
    a. Some functionality could be consumed from the cloud out-of-the-box. For example, some organizations may have custom applications (e.g. for CRM, ERP, and other areas) which can be retired, with the same functionality being consumed as SaaS functionality from established players such as Sales Force (SFDC).
    b. Alternatively, some functionality could be developed on public cloud PaaS platforms like the Google App Engine or Force.com.
3. Some applications might get re-engineered to SaaS applications that are provided to the external world. That is, the organization becomes a provider of SaaS services in some scenarios. An example is a bank offering payments services to other banks, or the more easily understood software vendor offering their software as SaaS.
4. Finally, not all applications and functionality will move to the cloud. Several applications could continue to remain alive as on-premise applications. For example, this could be for reasons related to criticality – high criticality implying a "do not touch" approach and low criticality implying a "to be retired / killed in the future" approach.

Key Considerations

The important questions that need to be addressed in the Portfolio Analysis step are

1. What is the split up of the above 4?
2. On what basis is the split up to be achieved?
3. How is the prioritization to be done? For example, which applications are cloud ready, and which require some effort to be made cloud ready?

Output

The output of the Portfolio Analysis step is a document capturing the current application landscape and showing the recommendations, per application, for the Cloud Migration / Adoption strategy.

Planning for the Future

While hardware virtualization delivers the benefits of infrastructure optimization, the next level benefits of cloud adoption will be realized only when platform and application functionality is moved to the cloud. Therefore it is clear that any enterprise cloud adoption strategy will need to have a clear, precise and well planned approach towards Portfolio Analysis in place. The aim for this step should be that of "Selective Adoption" - *Move platforms and applications in a phased manner and avoid a big bang / all or nothing approach!*

Step 3: SOA Revisited

The cloud computing world of the future will be the world of "Services". Every piece of application functionality on the cloud will need to be available as a service. The cloud will be the intermediate platform for service producers and service consumers to come together – fulfilling the vision originally outlined by the concept of Service Oriented Architecture (SOA). In fact, just as SOA in the previous decade was synonymous with *Web Services*, SOA in the current decade is likely to become synonymous with *Cloud Services.*

Key Considerations

Service Definition is a well-known topic in the world of Service Oriented Architecture (SOA). The concepts will be similar in the cloud services world as well, though implementations could be different. The input to the Service Definition phase will be the portfolio analysis document produced at the end of step 2. While the portfolio analysis document focus on the applications themselves, the focus in step 3 will now be on the application *functionality* that will need to be available in the cloud, but may *not* be migrated as-is. This essentially means two types of application functionalities:

a. Applications (e.g. Custom CRM) that will be retired on-premise and consumed from a public cloud provider.
b. Applications (e.g. Custom Business Related) that will be retired on-premise and re-engineered to be hosted on the private cloud.

During Service Definition, the work will be the identification of the application functionality that needs to be implemented as cloud services, and the definition of the cloud services. The service definition also needs to include the definition of service level agreements (SLAs) that the service can support.

Output

The output of the Service Definition phase is a document that provides, for the list of applications from the portfolio analysis document, the set of services that each application

maps to, the SLAs for those services, operations involved, input and output data, and a description of the service functionality.

Planning for the Future

Cloud services that replace on-premise application functionality should

1. Offer the same (or enhanced) functionality as the on-premise applications they are replacing
2. Offer nearly the same performance characteristics as the corresponding functionality on-premise.

In private cloud scenarios, it is likely that the SLAs for cloud services / applications will be the same as the on-premise applications. However in public cloud scenarios, there is a possibility of deviations – since the platform SLAs available might be different from the exact requirements. In either case, the key consideration in this step is to *aim for near zero SLA deviation with the on-premise scenario.*

Step 4: Technology Planning

The advent of Service Oriented Architecture (SOA) introduced a paradigm shift from the client server model to a service producer – service consumer model. Cloud computing is likely to take this paradigm shift to a logical conclusion. The enterprise (cloud) architecture of the future will need to be built keeping in mind computing available as a service – be it infrastructure, platforms, software, data, or business processes.

Key Considerations

Some of the things to keep in mind from a technology roadmap and architecture perspective would be as follows:

1. ***Vendor Selection****:* There are a large number of vendors at each layer of the cloud stack. During the technology planning activity, vendor selection should be an important focus area.
2. ***Private cloud scenario:*** For a private cloud, the architecture roadmap would need to cover most of the layers depicted in figure 1. It would also be necessary to keep in mind the perspectives of both the cloud provider [internal to the organization] and consumer [also internal to the organization]. For example, the provider perspective would need to consider application deployment and provisioning aspects, along with monitoring and management of cloud services.
3. ***Public Cloud Scenario:*** In this case, the architecture would need to address the perspective of the consumer only.
4. ***Hybrid cloud scenario:*** In this case, the architecture roadmap would need to elaborate on the interface and integration aspects. One other aspect to consider would be on how *"cloud burst"* situations would be handled.
5. ***Cloud Integration:*** As depicted in figure 1, in steady state an enterprise will have some applications migrated to the private cloud, some continuing as on-premise applications, and some applications retired with the corresponding functionality being consumed from the public cloud. Application integration in a cloud scenario

therefore becomes a key consideration in architecture planning. *"Service Oriented Integration",* followed by SOA practitioners, would be the recommended approach to follow in this case.

6. ***Management & Monitoring:*** In case of private cloud implementations, one of the important things to keep in mind is the management and monitoring aspects. These are sometimes overlooked because of the mindsets that are oriented towards public cloud, where services are merely consumed and the consumer does not need to manage and monitor them. Management and monitoring includes health checks for services, usage data collection, access management, managing scale up and scale down requirements, and so on.
7. ***Security Considerations:*** These are outlined in the section below.

Cloud Security

Cloud Security is a topic that evokes strong reactions - usually these are concerns on the security aspects of cloud computing. Consequently there is a lot of material, some of it hype and some excellent references, around the area of security and cloud computing. One of the best starting places is the Cloud Security Alliance web site[iv].

Key Considerations for Cloud Security

Private cloud security will be similar to on-premise security, with some unique aspects:

1. It is important to evaluate the vulnerability of the cloud enablement platform itself (e.g. hypervisor vulnerability).
2. You would also have access control policies for your on-premise applications. When moving these applications to your private cloud you might need to re-define these policies.

Public cloud security is where there are some apprehensions. However, in an age where we have so much of internet e-commerce, payment systems, card transactions, all happening online, the issue of public cloud security might not be as complex as it appears. When consuming services from the public cloud, you would need to evaluate the Cloud Provider's security policies and keep some of the following in mind:

1. How does the provider implement Identity and Access Management (IDAM)?
2. What data protection / encryption mechanisms are being used?
3. What type of logging, tracking and audit mechanisms are in place? Where is the audit data stored?
4. Where is business data stored?
5. What types of application security standards (e.g. Web application security) are being followed?

Output from Technology Planning

The output from the technology planning phase is an Enterprise Cloud Architecture document – this is the recommended blueprint for all the technology and architecture aspects of cloud adoption in the enterprise. This blueprint should outline the cloud scenarios applicable to the enterprise, a customized version of the cloud stack reference architecture,

the recommended technology platforms at each layer of the stack, integration aspects, security considerations, and management and monitoring approaches.

Planning for the Future

One of the challenges for cloud adoption is the fact that some of the technologies are not yet mature and there are limited reference architectures and implementations available. The industry is very fragmented with a large number of new vendors. Standards in cloud computing are in the process of getting defined. One way to deal with all this uncertainty would be to leverage the principles, learning and best practices from the world of Service Oriented Architecture (SOA) for Enterprise Cloud Architecture as well. For example, principles of SOA governance (e.g. registries, repositories, policy management) could be used as a base for coming up with cloud governance approaches as well.

Regarding cloud security - apprehensions on cloud security should by no means be allowed to become an inhibitor for cloud adoption. One way to address cloud security concerns is to separate out the public and private cloud security requirements. Security aspects for each scenario can then be addressed separately. In a private cloud scenario, the aim should be to replicate, as far as possible, the security policies currently in place. Of course, there could possibly be some refinements and additions for the private cloud context; such as hypervisor security. In a public cloud scenario, it is important to assess and evaluate cloud policies of the provider, and check whether those are acceptable from the organization perspective.

Step 5: Cloud Capacity Planning

The traditional approach to capacity planning has some fundamental limitations, as outlined below.

1. **Underutilization:** Capacity is usually planned for a peak load expectation. However, during the large majority of the time when the load is below the peak, the capacity lies underutilized.
2. **High upfront costs:** Secondly, capacity is expensive and involves high upfront capital expenditure, along with significant time and effort for acquisition and provisioning. In addition, shrinking or de-provisioning of capacity is very difficult.
3. **Minimal margin of error:** As a logical conclusion, any planning errors can have a huge impact – due to which the capacity planning exercise is usually very elaborate and time-consuming.

The promise of cloud computing is the potential to address all of the above limitations. Cloud computing can dramatically increase capacity utilization percentages. In fact, if servers were not underutilized, one basic pillar of the cloud business case would have fallen apart. In the cloud, capacity is less expensive and does not involve an upfront capital investment. The time for acquisition and provisioning is measured in minutes rather than days or weeks. Planning errors can be rectified easily due to the concept of elasticity – the ability of the cloud to increase or shrink capacity on demand.

Key Considerations

Cloud adoption implies a paradigm shift for capacity planning related activities – from a *dedicated resources and peak load model* to a *shared resources and average load model*. One other important aspect is that the focus for cloud capacity planning should be *across the layers* of the enterprise cloud architecture and not just at the infrastructure level. Some of the considerations would be as outlined below:

1. **IaaS Layer:** Key aspects of capacity planning will involve determining the number of virtual machines that can be accommodated on given physical hardware – depending on the CPU, RAM, and storage capacities needed for each virtual machine (and also keeping in mind hypervisor capabilities).
2. **PaaS layer:** Key considerations on capacity planning include determining the number of applications that are likely to be developed and hosted on the platform that needs to be made available as a service.
3. **SaaS and BPaaS layers:** Key considerations would include the number of users for each application or process, peak hours of usage, and so on.

One point to be noted is that capacity planning in the cloud is typically needed in the private cloud scenario, or when the enterprise is planning to be a provider of public cloud services.

Output

The output from this step is a Capacity Plan document, which outlines the capacity required [for example, number of virtual machines] to support the typical transaction volumes expected, interfaces that need to be built, management and monitoring requirements and security aspects.

Planning for the Future

With cloud adoption, the future of capacity planning looks much better than the past. One change in approach would be the consideration that traditional capacity planning is usually restricted to the infrastructure layer only. In the cloud scenario, besides the focus on the IaaS layer, it is important to have a holistic approach for capacity planning covering the PaaS, SaaS and BPaaS layers as well. It is also critical to keep in mind and plan for the dynamic allocation and de-allocation aspects.

Step 6: Monetization

Once cloud services are in place and available, they need to be monetized. The process of monetization of cloud services is often called *chargeback.* Usage based pricing is one of the fundamental building blocks of the cloud computing paradigm. Therefore approaches towards chargeback and monetization need to be planned and thought through carefully.

Key Considerations

The plans and strategies for chargeback approaches need to be thought through well before the cloud services go live. One obvious aspect here is to plan for usage metering and billing. There is a pitfall to be avoided though. Current approaches often look at chargeback

mechanisms from an individual cloud layer point of view rather than a holistic view. In other words, there could be an approach for chargeback of the IaaS layer, but none (or a completely disjoint one) for chargeback of the PaaS / SaaS layer. Ideally, what you would want is a single unified approach to charge the consumer for the usage of a service in the SaaS layer, which encompasses the implied usage of all the layers underneath.

The main recommendation here would be to ensure that the chargeback strategies include a focus across all the 4 layers of the cloud, from infrastructure to applications.

Output

The output of this phase is a monetization methodology note, which outlines chargeback strategies, pricing models, and approaches to invoice generation from metered usage data.

Conclusion

While cloud computing becomes increasingly widespread and gains mass acceptance, legacy investments in data centers, software and hardware infrastructure cannot be wished away. Private cloud implementations will become increasingly common to leverage such legacy investments, as well as reap the benefits of cloud adoption. In this paper, we have discussed six key considerations for private cloud adoption. The recommendations and insights provide an overall framework that can assist organizations during the planning and strategy phases of their cloud adoption journey. At Happiest Minds, we can assist you towards leveraging this six step approach for your cloud adoption initiatives.

References

1. *The case for Energy-Proportional Computing*, Luiz André Barroso and Urs Hölzle, 2007.
2. VMware white paper: Cloud Computing – An evolutionary approach to an IT revolution
3. Eucalyptus home page: http://www.eucalyptus.com/
4. Amazon Web Services home page: http://aws.amazon.com/
5. IBM Cloud Computing home page: http://www.ibm.com/cloud-computing/us/en/

About Happiest Minds Technologies

Happiest Minds, the Mindful IT Company, applies agile methodologies to enable digital transformation for enterprises and technology providers by delivering seamless customer experience, business efficiency and actionable insights. We leverage a spectrum of disruptive technologies such as: Big Data Analytics, AI & Cognitive Computing, Internet of Things, Cloud, Security, SDN-NFV, RPA, Blockchain, etc. Positioned as "Born Digital . Born Agile", our capabilities spans across product engineering, digital business solutions, infrastructure management and security services. We deliver these services across industry sectors such as retail, consumer packaged goods, edutech, e-commerce, banking, insurance, hi-tech, engineering R&D, manufacturing, automotive and travel/ transportation/hospitality.

Headquartered in Bangalore, India; Happiest Minds has operations in USA, UK, The Netherlands, Australia and Middle East.

To know more about our offerings. Please write to us at
business@happiestminds.com

About the author

Shantanu Paknikar (shantanu.paknikar@happiestminds.com) is General Manager, Solution Initiatives in the IT Services division of Happiest Minds. His expertise areas include middleware systems, SOA, BPM, enterprise integration, cloud computing, multimedia, and distributed computing systems. He is passionate about innovation, research, and learning. He believes that the IT services industry has turned the corner and needs to focus primarily on customer business problems and holistic, end-to-end IT solutions addressing them rather than products, feature sets, individual technology or expertise based offerings.

Footnotes

[i] At the same time, a lot of non-core functions are being consumed from the public cloud.

[ii] References for private cloud adoption can be found all over the net. Some good places to look for

[iii] include the leading vendor sites such as Microsoft, IBM, VMware and Amazon.

[iv] See reference 1 http://www.csa.org