

An Approach to Cost-Effective Transformation of Workloads towards Cloud Delivery Models

Joydipto Banerjee
IBM
DLF IT Park
Kolkata

Biswajit Mohapatra
IBM
Tech Park One
Pune

ABSTRACT

Moving workloads on to cloud, promises to be an attractive investment for enterprises planning for a sustainable IT landscape. However, leveraging the existing IT investment of an organization and moving them over to the cloud environment with minimum business disruption and costs, seems to be the biggest challenge in adopting cloud based modernization. Enterprise workloads move into cloud delivery models with varying levels of transformation effort, return on investment and productivity gains. Like any transformation effort, migration to cloud can also be viewed as a phased transition/journey. In this respect, studying the various migration patterns, it is found that a common theme emerges that have been depicted in this paper. The authors have illustrated an approach that takes a portfolio view of applications and assesses highest degree of cloud affinity considerations that includes complexity of the selected workloads, target environment, organization maturity etc. The paper also discusses the taxonomy of cloud migrations and the overall migration model

Keywords

Cloud Computing, Workload Migrations

1. INTRODUCTION

The transition to cloud computing represents a major disruption for IT organizations. At the same time, market hype and diversity of offerings continues to confuse clients [1]. Some of the questions to ponder include - will the workload run in the target cloud environment, e.g. compatible infrastructure, middleware and operating system image? Will the target cloud environment comply with applicable security, privacy, performance, availability and regulatory requirements? Will benefits be realized from migration, e.g. lower ongoing operating costs, improved service levels or increased flexibility and responsiveness? To find answers to these questions one need to consider cloud migrations from an overall business modernization perspective and understand the key drivers.

To start with, it is important to understand that migrating applications to cloud is just one type of cloud adoption that an organization may aim at. Cloud migration of an application is defined as the process of relocating an application to cloud and making any necessary modifications that are required for the move.

The benefits of cloud migration include [2] –

- Better standardization of the applications within the enterprise.
- Improved operational efficiency like higher availability, better elasticity and faster response time.

- Reduced operational costs. More efficient utilization of infrastructure resources.

The various ways [3] of moving an enterprise's business applications to cloud can be broadly looked upon as –

1.1 Rewrite the application to exploit the cloud feature

These are 'born on the cloud' [4], applications that are written taking advantage of cloud and virtualization features. Though it is time consuming, yet most of the future green-field applications will fall in this category.

1.2 Replace with software as a service (SaaS) equivalent

For legacy application or those that are nearing end of life, it is sometime more cost effective to opt for similar services through SaaS model from cloud service providers.

1.3 Relocate the application to the cloud environment

This method refers to moving enterprise applications to cloud environment using various tools and techniques. This covers both traditional to cloud as well as cloud to cloud migrations.

Of the above three options, it is the last one which is mostly prevalent and which is being discussed here.

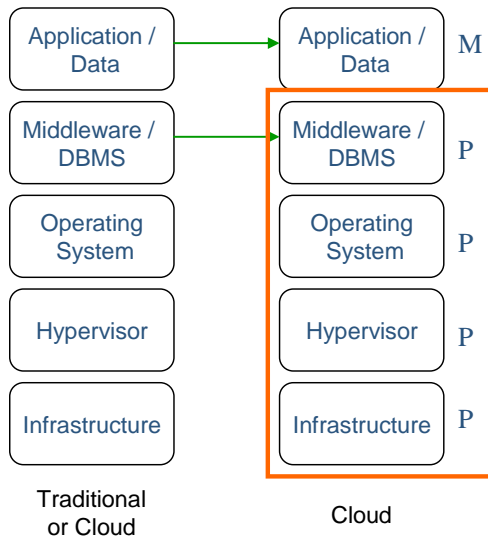
2. MIGRATION TO CLOUD TAXONOMY

The process to relocate applications onto cloud can itself be of several types. Migrations can be categorized into the following entities along with their major steps:

2.1 Data Migration / Application Data Migration Only

In this category, application packages, application databases, application shell scripts are migrated to a cloud provisioned environment. Key steps –

- Hardware resources (including Virtual Machines) are provisioned in the cloud
- Operating System is installed and configured in the target cloud
- Applications and databases are then installed and configured
- Data and application data are finally migrated over to the cloud environment



M = Migration P = Provision new

Fig 1. Data Migration / Application Data Migration

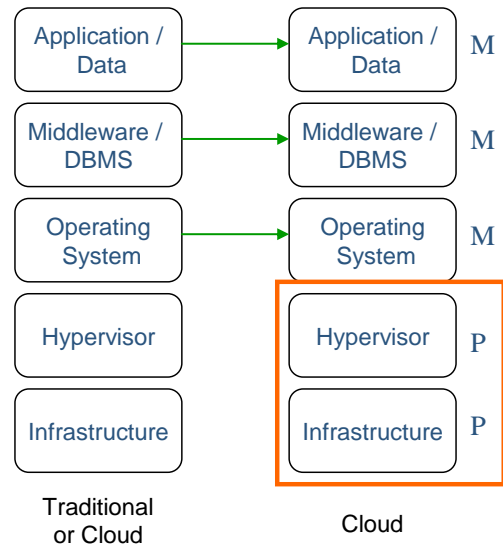
2.2 Application, Database, and Data Migration Only

Here, application packages, application databases and application shell scripts are migrated to a cloud provisioned environment. Any Operating System or middleware configurations and setup that are specific to the application are included as part of the workload migration. Key steps –

- Hardware resources (including Virtual Machines) are provisioned in the cloud
- Operating System is installed and configured in the target cloud
- Applications, application configurations, databases, and data are migrated over

This type of migration scenario appears to be the most complex one since it might involve moving legacy custom applications that must be re-compiled on the target or might even require porting and code remediation.

Applications written and compiled for one type of computer architecture will not run on another. Moreover, there could be issues with Little-endian and Big-endian formats. Engineers migrating code from a machine of one type of endianness to a machine of the other type must be aware of the possibility that problems will arise. Unfortunately the problems can be extremely subtle and not appear until deep in the testing phase. Even cross platform database migration usually involves using data migration methods like Export /Import etc.



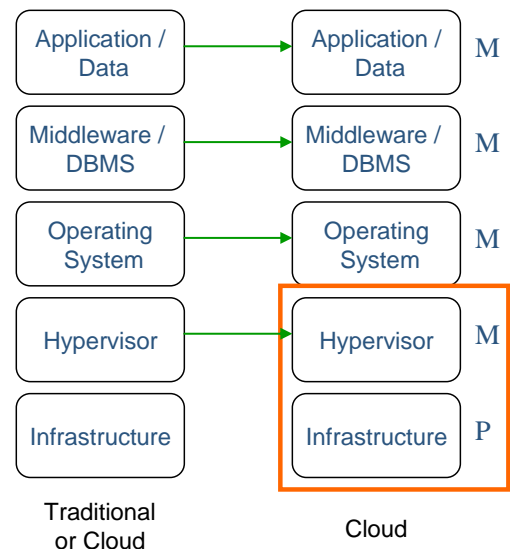
M = Migration P = Provision new

Fig 2. Application, Database, and Data Migration Only

2.3 Instance or Captured Image Migration

Application packages, application databases, application shell scripts are migrated to a cloud provisioned environment. The Middleware and Operating System are migrated along with their basic configurations. Examples of specific configurations include: application database configurations, data replication settings, custom application file paths, crontab entries etc. Key steps –

- Virtual Machine (and optionally an Operating System) is provisioned in the target cloud
- Operating System, applications, application configurations, databases, and data are then migrated over



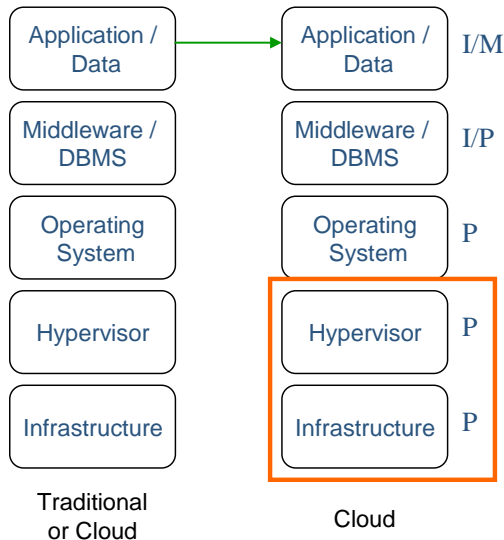
M = Migration P = Provision new

Fig 3. Instance or Captured Image Migration

2.4 Migration to a Managed Environment

This involves transformation of application and data layers that occurs with delivery onto a compliant managed environment (Hardware, Operating System, Middleware). This may require significant effort to refactor application and/or data layers to achieve a “cloud-ready” application. Key steps –

- Target cloud has compliance specifications
- Migration services and tools provide capabilities to comply with compliance specifications as part of migration



M = Migration P = Provision new

I = Install New

Fig 4. Migration to a Managed Environment

3. APPLICABILITY OF MIGRATION SCENARIO TO CLOUD TYPE

A comparative study of the various migration scenarios and their applicability to the cloud service models are shown in Fig 5.

Cloud Type Migration Scenario	IaaS	PaaS	SaaS
Data Migration / Application Data Migration Only	Y	Y	Y
Application, DB, and Data Migration Only	Y	Y	N
Instance or Captured Image Migration	Y	N	N
Migration to a Managed Environment	N	N	N

Fig 5. Mapping scenarios to cloud service models

4. CLOUD MIGRATION VENDOR CAPABILITIES

The various vendors / products that are currently providing solutions meant for the various categories of migration are shown in the Fig 6.

Data Migrations / Appl Data Migrations Only Aspera Tervela	Application, DB, and Data Migrations Oracle AppZero CohesiveFT CliQr AppLogic
Instance/Image Migration Vmware Converter Platespin Racemi Rivermeadow DoubleTake Ravello	Migration to Managed Cloud TnT Services Rapid Migration

Fig 6. Cloud migration tools and solutions

5. CLOUD MIGRATION STRATEGY

At the outset, it is very important to remember that there should be clear and compelling business reasons to move to cloud, otherwise the efforts and costs incurred might not deliver the desired results. Not all workloads are candidates for cloud migration – either technically or from a business standpoint or both. Before jumping headlong into any cloud migration activity it is advisable to address these key elements as described below.

5.1 Cloud Opportunity Identification

It is important to know if cloud is really a good fit for the customer across a number of characteristics. These include culture, organization, funding, risk model, financial targets and existing skills and capabilities (across both IT and the business). Also important is the group’s willingness to try potentially radically different delivery models.

Determining “fit” is based on many characteristics - culture, organization, funding, risk model, financial targets and existing skills and capabilities. Key questions and topics:

- Why is the organization interested in cloud? Is it ready to effectively leverage cloud?
- Is there a Return-On-Investment or quantified benefit the organization expects? Does it have a business case or cloud strategy?
- Are the IT department and/or the Line-of-business involved in the decision?
- What are the IT capabilities of the organization? Are they sufficient to meet business needs? Can cloud address the gaps?
- What is the organization’s end-state vision of cloud when complete?

All these go into identifying potential cloud opportunity areas, determining the desired cloud targets and accessing the potential cloud workloads

5.2 Current Business Processes and Potential Impact

Understanding organizational, cultural and process impacts are important considerations which may be obvious or hidden. They often are part of the motivation to implement cloud. Key questions and topics:

- What is the environment? What are the business objectives?
- What roles do people perform? What is the organizational culture?
- What is the structure of the organization?
- Who has funding authority?
- Does each group have complete ownership of its data, applications, and underlying infrastructure?
- What parts of the IT environment are in silo versus shared?
- Is there a willingness (e.g. culturally) to share services?

Reviewing the current IT and business environment at a high level helps the organization in taking stock of the overall readiness for cloud and also identify critical gaps in the process.

5.3 Identify Key Areas of Automations

Identifying if and where automation makes sense has a direct impact on implementing cloud. Finding out if and where it can be done, what part of the process is a good candidate, and potential impacts to processes, roles and financial models helps define the cloud project. Key questions and topics:

- Where does automation make sense?
- What existing processes have been defined?
- Who owns the process? Does it have more than one owner? Is self-service required?
- Will elastic scaling and provisioning be needed? Does it need to be automated?
- What are the underlying SLAs?
- What are the workflow and approval steps that will be automated? Will automation apply to some or all of a process?
- Are business processes, process engine, and/or workflow automation tools already in place?
- Will automation help or hinder the alignment with key regulatory and/or compliance requirements?

6. SYSTEMS ENGINEERING PROCESS FOR CLOUD MIGRATION

A generalized view is used as shown in Fig 7 for the purposes of understanding how migration patterns are applied to a project management style view and this view is used to set the initial context for a migration discussion. The end goal of these three phases is identification of cloud service and delivery model along with a set of candidate migration workloads.

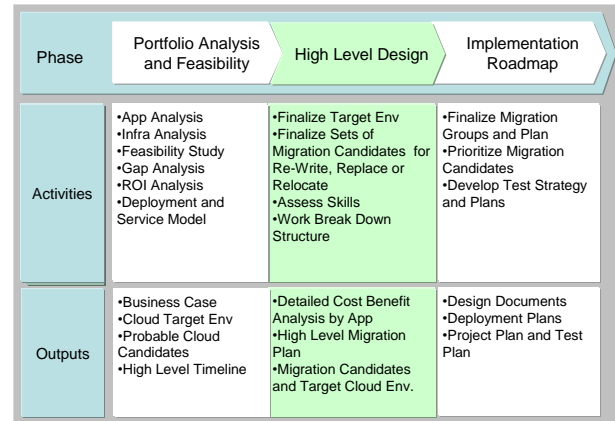


Fig 6. Cloud migration tools and solutions

6.1 Portfolio Analysis and Feasibility

6.1.1 Understand the Current IT environment

Once there is an understanding of the organization's starting point and the potential organizational and cultural impacts of cloud, it's time to start exploring the organization's current IT landscape – the kind of server systems (Intel, Power etc.), Operating Systems (Windows, Unix) and the software stack (Databases, Web servers). An initial analysis is done from both levels – infrastructure and application. This is done by interviewing key stakeholders as well as using scripts or automated tools to scan the existing portfolio of hardware and software. With the available information, an initial feasibility study is done to screen the workloads for their cloud affinity.

6.1.2 Select Cloud Service and Deployment Model

The following models are to be selected in this phase:

- Cloud service model(s) [5] – IaaS, PaaS, SaaS.
- Cloud deployment model(s) – public, private or hybrid.

Several deployment options further shape cloud computing delivery models.

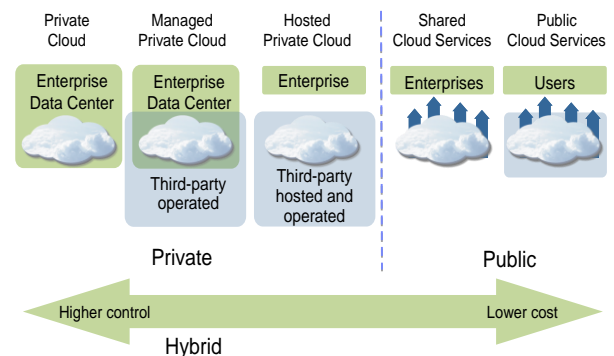


Fig 8. Various cloud deployment models

A private cloud naturally means a tighter control from the cloud consumer's perspective although it usually means a higher cost as well. On the other hand, as one moves toward public cloud environment, the level of control reduces with the corresponding cost advantage. The choice of deployment model is directly dependent on the risk appetite of the organization.

6.1.3 Cloud Use Cases and Actors

This involves determining the kind of business needs that would be fulfilled by the proposed cloud environment and the users who will be consuming those services. Several scenarios may exist depending upon the kind of cloud setup, viz. public, private, hybrid and the type of the users – internal and/or external. The situation becomes more complex if multiple cloud vendors are involved or if there are challenging integration issues with cloud and on-premise applications.

6.2 High Level Design

6.2.1 Architecture Overview

Taking an architectural approach to cloud helps the organization focus on technology and process requirements and treat IT as a utility. The organization can grow the use of the investment over time, building traction and obtaining constant value from cloud. Examples include choosing a single operating system and/or hardware platform or defining key processes spanning more than one business unit. A reference architecture, like that of NIST's Cloud Computing Reference Architecture [6] helps in incorporating all the essential elements that define the target cloud architecture.

6.2.2 Architecture Decisions

During this detailed transformation analysis phase, it is the cloud architect's responsibility to document the entire decision-making process, capture the architectural viewpoints, the tools and methods used and reusable assets used, if any. The following aspects of the overall solution are considered here:

6.2.2.1 Design



Fig 9. Design Elements

6.2.2.2 Transformation

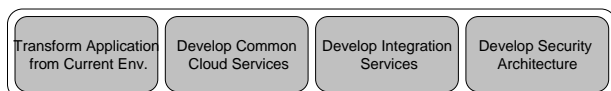


Fig 10. Transformation Elements

6.2.2.3 Deployment

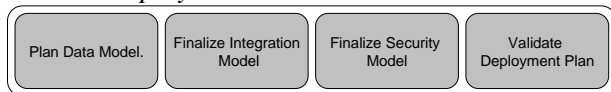


Fig 11. Deployment Elements

6.2.3 Candidate Workload Selection

Choosing the right workloads to move to cloud and understanding the effort required for migration are vital. To move a workload, one must know all of the components of the workload, the associated processes which depend on the workload and the effort as well as skill required to move the workload. The key questions and topics which are relevant in this context are:

- Which workloads will be moved to cloud?
- Are they composite or discrete workloads?

- Is individual component or groups of applications involved? Do they contain data and/or, applications?
- Do the workloads contain specific hardware configurations and equipment?
- Which applications need to be migrated and what level of effort is required?
- Which workloads will need to be fully converted / re-engineered?
- Which workloads can be moved with no modification?
- Can some workloads force consolidation of applications as part of the migration effort?

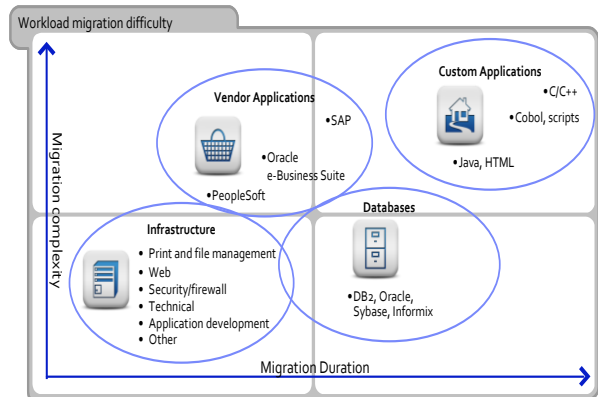


Fig 12. Migration Workload Positioning Matrix

When it comes to the difficulty or complexity of any migration engagement, not all workloads are equal. Fig 12 shows how it is preferable to divide complex migration opportunities into manageable workloads. The Infrastructure and Databases, positioned in the third and fourth quadrant respectively are the potential workloads that can be considered initially. Migration effort needs to match the client value it provides and thus it is advisable to start with the “low hanging fruit” that offer a quick return and establish credibility. The more complex candidate workloads consisting of third party or in-house custom applications can be taken into consideration later on.

Another view of comparing cost and risk of implementation is depicted in Fig 13. This also shows that it is best to start with development and test workloads followed by the more complex, production ones.



Fig 13. Comparing Cost and Risk of Implementation

6.2.4 Non-Functional Requirements

As in any transformation exercise, it is important to determine the non-functional requirements that the organization has defined like performance criteria, audit and control, backup requirements, security, response time etc. Most of these parameters depend on the business criticality of the selected workloads and the type of cloud service model proposed.

6.2.5 System Context

Understanding the system context helps in limiting the overall scope of the cloud migration candidates and also helps avoid any potential disruption in business functionality. Migration to the target cloud environment is usually done in phases, beginning with a set of workloads that are easiest to deploy. This is followed by the migration of the more complex set. During this exercise it is important to validate how the selected workloads are interacting with the other workloads currently and how would they be affected once migrated to cloud environment. A System Context diagram is thus an essential architectural artifact produced during the engagement.

6.3 Implementation Roadmap

6.3.1 Migration Groups and Plans

Migration groups are the sets of candidate workloads that have been identified in earlier phases, to be migrated together. Sets of candidate migration workloads need to be prioritized based on business or technical criteria. This sets the stage for which group of workloads will be migrated first and which should follow later on. A migration plan consists of details of the migration like a sequence of tasks that should be performed, their owners, expected start date etc. This plan forms the blueprint for the entire migration. A formal project plan, from a project management perspective is also usually prepared during this time.

6.3.2 Test Plans

At this time, a test strategy is determined. Various test plans, mostly for regression tests are prepared to check that the application behavior and business functions are intact post migration.

7. MIGRATION FRAMEWORK

Once the above cloud migration approach is established, the next phase is developing a detailed migration framework that would be applied for each of the identified sets of candidate workloads. Using inputs from many different team and sources within IBM who are focused on cloud migration, a normalized model showing the overall systems engineering and governance related approach of migration emerges. The purpose of this view is to illustrate the systems engineering approach supporting migration and to give some sense of those steps which occur sequentially and which can occur in parallel. It is important to note that the phases like analysis, planning and design that are depicted in this framework are applied to each of the selected workloads, rather than the entire portfolio. The Fig 14 illustrates the framework

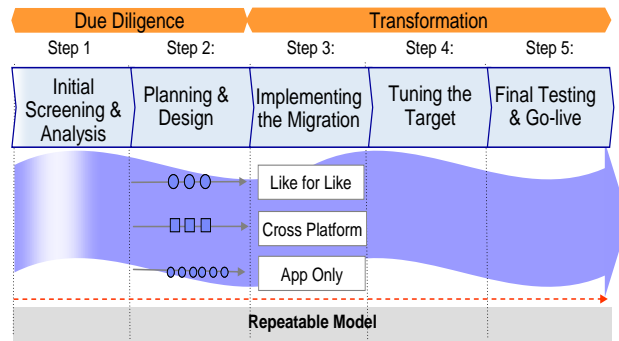


Fig 14. The framework for cloud migration

The detailed description of this framework is mentioned in [7].

8. MIGRATIONS – MYTH AND REALITY

There are many misconceptions and false expectations that surround migrations and it is important to realize the actual return on investment before proceeding with this initiative. Given below are the typical migration outcome types and how far they are valid:

8.1 Standardization

Standardization is one of the typical benefit of cloud computing and it is realistic to expect positive returns, especially in the operational expenses by moving to a virtualized environment with a standard set of hardware and software. Removing duplicate versions of the same software, thus requiring multiple licenses for the same product and maintenance costs, are just some of the advantages that standardization offers.

8.2 Consolidation

Mainly due to standardization, certain consolidation advantages can be clearly realized from a migration engagement. In most cases, the application portfolio assessments and the inventory analysis efforts go a long way in taking a stock of the organization's IT infrastructure and doing away with redundant, duplicate or unused resources. Moreover, the physical to virtual migration results in consolidation of servers, databases etc and reduction in data center floor space.

8.3 Scalability

It is important to bear in mind that migrating to cloud environment by itself does not guarantee more scalability. Infact, the application needs to be scalable in the first place to gain from the cloud model. This can only occur for those applications which have the affinity for scaling in the first place. Application Servers (such as Websphere Application Server) and Databases (including DB2 and Oracle) as well as analytics tooling like Hadoop have different levels of clustering capability. Cloud is best used in this scenario to instantiate additional instances of an application server which is pre-configured to work in a clustered (scalable) environment.

8.4 Reliability and Availability

Again, if the application in question is not inherently robust with a sound architecture, just moving it to cloud environment will not increase its reliability. Similar reasoning holds true for application availability as well. As an example, if a workload (application, data, system, etc.) has an affinity to be in a more available or scalable configuration (such as clustered configuration or a grid style environment), then using cloud's ability to rapidly deploy new instances, this can increase the availability of a system through the creation of new nodes in a cluster environment.

8.5 Performance

A commonly held belief is the migration to cloud environment will bring better performance on the part of the application. The focus of the cloud model is to deliver services in the most efficient manner possible without regard to the performance (or availability or scalability) of the workload. In fact, when moving a workload to highly virtualized environment, those workloads which do not have an affinity for a virtualized environment will likely have negative performance impacts from moving to a "bare metal" configuration to a virtualized configuration.

8.6 Speed to Market

Standardization and automation in a cloud environment does indeed bring about considerable improvement in deploying applications and speed to market. There is a strong dependency on standardization as a foundation for more rapid provisioning and for standardizing the development and testing environments used in new application development and maintenance of existing applications.

8.7 Error Reduction

Standardization and greater automation in cloud environment brings about a reduction in the number of errors in application development and deployment.

9. MIGRATION - ISSUES AND CHALLENGES

Most of the issues and challenges which hinder migration initiatives are actually similar to the generic barriers that exist for cloud adoption. The following are some of the key ones:

9.1 Data Security and Privacy

Enterprises have the requirement to comply with many of the regulations that require data governance. By moving the data into the cloud, enterprises fear they will lose some capabilities to govern their own data set. They would have to rely on the service providers to guarantee the safety of their data.

9.2 Service Quality

Service quality is one of the biggest factors that the enterprises cite as a reason for not moving their business applications to cloud. They feel that the service level agreements provided by the cloud providers today are not sufficient to guarantee the requirements for running a production applications on cloud especially related to the availability, performance and scalability.

9.3 True Cost Savings

Due to lack of understanding of the business reasons for adopting cloud and the many misconceptions surrounding it, companies often fail to realize the true cost saving potential after migrating to cloud. The cloud revenue model should include things such as indicators for comparing the availability, performance versus recovery time objective,

workload-wise assessments, Capex versus Opex costs benefits, utilization, etc.

9.4 Network Bandwidth

Poor network performance or lack of sufficient bandwidth is a cause of concern for organizations in their pursuit for cloud adoption. Moreover many cloud based applications are too bandwidth intensive.

9.5 Integration

Many applications have complex integration needs to connect to other cloud applications as well as other on-premise applications in a simple, quick and cost effective way.

10. CONCLUSIONS

Organizations want to migrate to the cloud to increase the business agility, flexibility, security and resiliency of their IT infrastructure while reducing costs. Moving workloads to "cloud environments" need to follow a multi-step process and which apparently does not look very different from a traditional cross-platform application migration. In summary, to accelerate an organization's cloud adoption strategy, one needs to –

- Identify those workloads that will best fit the target cloud while giving a good return on investment
- Have a structured approach to analyze each individual IT workload
- Have an alternative solution for workloads that are incompatible on cloud

This work tries to address the current gap in cloud migrations – the lack of a pre-defined approach and template. The authors have endeavored to put in place some framework or repeatable model that can accelerate an organization's journey to cloud.

11. ACKNOWLEDGMENTS

The authors would like to thank the Business Application Modernization department of IBM for providing the opportunity to research and work on this subject.

12. REFERENCES

- [1] Isom Pameka K, Kerrie Holley. 2012. Is Your Company Ready for Cloud: Choosing the Best Cloud Adoption Strategy for Your Business? Pearson Publication.
- [2] Galloway K, Garvin S. 2005. The Migration Advantage. IBM Systems Magazine.
- [3] Pettey C, Meulen R. Gartner Identifies Five Ways to Migrate Applications to the Cloud. 2011. Gartner.
- [4] Amies A, Sluiman H, Tong G. Q., and Liu N. G. 2012. Developing and Hosting Applications on the Cloud. Pearson Publication.
- [5] Savolainen E. Cloud Service Models. University of Helsinki.
- [6] Liu F, Tong J, Mao J et al. NIST Cloud Computing Reference Architecture. 2011. Special Publication 500-292. National Institute of Standards and Technology.
- [7] Banerjee J. 2012. Moving to the Cloud: Workload Migration Techniques and Approaches. In Proceedings of the 19th IEEE conference on High Performance Computing (Pune, India, December 2012).