Data Migration Across The Clouds

Prashant Pant, Sanjeev Thakur

Abstract— Having an effective and efficient strategy for optimizing long distance data migration is essential for every data center. With the ever increasing demands for the IT needs of businesses it is also important for data centers to deliver data migration cost effectively especially when faced with the demands from remote office back up, outsourcing, data center movers and cloud computing.

Data management and migration are important research challenges of novel Cloud environments. While moving data among different geographical domains, it is important to lower the transmission cost for performance purposes. Efficient scheduling methods allow us to manage data transmissions with lower number of steps and shorter transmission time. In previous research efforts, several methods have been proposed in literature in order to manage data and minimize transmission cost for the case of Single Cluster environments.

This paper Explores the issues and method of Data Migration across the Clouds

Index Terms—Cloud Computing, Data Migration, Security Issue, Cloud Architecture.

I. INTRODUCTION

Cloud computing is the use of computing resources (hardware and software) that are delivered as a service over a network (typically the Internet). The name comes from the use of a cloud-shaped symbol as an abstraction for the complex infrastructure it contains in system diagrams. Cloud computing entrusts remote services with a user's data, software and computation as shown in figure1.1.


Fig 1.1:Cloud Computing

In an industry well-known for "the-year-of" hype around new technologies, many see the cloud computing buzz as an extension of previous "year-of" technologies, such as application service providers, software as a service (SaaS), and utility computing. In this case, the hype may be true. In an Information.

Week Analytic Survey, 46% of companies surveyed say they'll use or are likely to use cloud CPU, storage, or other infrastructure services; this is up from 31% just one year ago. Our mission as IT consultants is to understand: what the cloud is and what benefits it offers clients; what challenges and obstacles clients might have to overcome to tap into the cloud; and how their management of IT must change to secure and control their new cloud-driven infrastructure. When you migrate a client to the cloud, the issues you'll face fall into the following overall categories.

1.1 Cloud Computing Services

Main Cloud Computing Services are given below:
a. Infrastructure-AS-A-Service.
b. Platform-AS-A-Service.
c. Software-AS-A-Service.
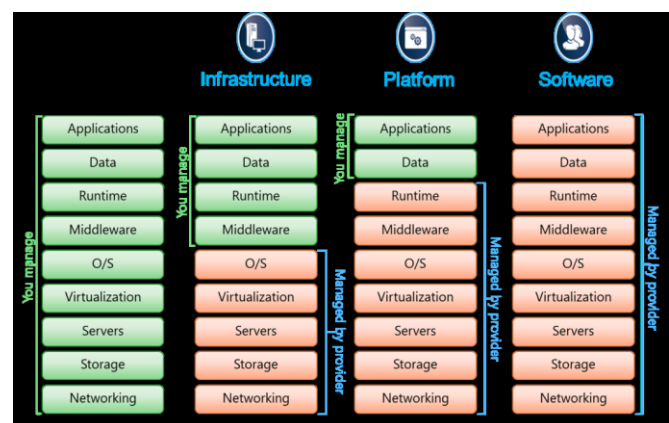Logical Cloud Computing Service model is defined in Fig: 1.2

Fig1.2: Cloud Computing Service Model

1.1.1   Software as a Service (SaaS)

In case of SaaS, you use the provider's applications on a cloud infrastructure with little to no control over the infrastructure, network, servers, operating systems, storage,etc. There are many examples of SaaS vendors – Salesforce.com, Google Apps, Ning, Cenzic, etc.

1.1.2   Platform as a Service (PaaS)

Customer deploys applications using an application development environment and middleware capabilities for specific languages like java, python, .net etc. and doesn't control infrastructure, servers, OS, or storage but has control over the apps. Some examples of PaaS vendors include Microsoft Azure, Amazon, Force.com

1.1.3   Infrastructure as a Service (IaaS)

Customer gets processing, APIs, storage, networks, and computing resources from the provider using his own OS, applications and may be some networking components. Some examples of IaaS vendors include Amazon, Rackspace, CloudFoundry.

1.2 CLOUD SECURITY ISSUES AND CHALLENGES

Cloud computing is a emerging technology with shared resources, lower cost and rely on pay per use according to the user demand. Due to many characteristics it has effect on IT budget and also impact on security, privacy and security issues .In this section all these issues are discussed.


Fig: 1.3 Cloud Securities

1.2.1 Physical Security

You want to make sure that physical security around the infrastructure is very tight – even tighter than in your environment because it's not your employees anymore.

1.2.2 Insider Abuse

When you "cloudize" your environment, you lose control over who's managing that infrastructure with your confidential information. Insider abuse is a common problem where information can be stolen and passed on to outsiders or they can collude with hackers.

1.2.3 Data Encryption

Cloud environments are shared and your data is in the same environment alongside data from other customers. Breaches can easily happen from one database to another.

1.2.4 Network Security

In the recent months, aggressive marketing by various Cloud providers have made it easier for hackers to get accounts and plant botnets. Cloud is also susceptible to a lot more Denial of Service attacks. Cloud Providers need to ensure that their perimeter is secure and barrier to attacks is high.

1.2.5 Virtualization Security

Almost all Cloud providers use virtualization to provide economies of scale and optimal distributed architecture. Virtualization has its own set of security issues.

1.2.6 Access Controls

Some of the big issues for Cloud services are around access control, authentication, user management, provisioning etc.

1.2.7 Application Security

With over 75% of attacks happening through Web applications, this becomes a critical piece in the overall cloud decision making process. Although the exposure is similar to what you would have in your own environment, it's on a massive scale and you may not have any control over it.

II. DATA MIGRATION ISSUES ACROSS COUDS

In an industry well-known for "the-year-of" hype around new technologies, many see the cloud computing buzz as an extension of previous "year-of" technologies, such as application service providers, software as a service (SaaS), and utility computing. In this case, the hype may be true. In an Information Week Analytic Survey, 46% of companies surveyed say they'll use or are likely to use cloud CPU, storage, or other infrastructure services; this is up from 31% just one year ago.

Our mission as IT consultants is to understand: what the cloud is and what benefits it offers clients; what challenges and obstacles clients might have to overcome to tap into the cloud; and how their management of IT must change to secure and control their new cloud-driven infrastructure. When you migrate a client to the cloud, the issues you'll face fall into the following overall categories.

2.1 Security

Security is an obvious threshold question; if the cloud is not secure, enterprises won't consider migrating their sensitive data to it, and the conversation is over. As Salesforce.com has proven, external SaaS providers can provide a level of security that will satisfy most customers. Migrating sales leads and prospect data is, however, quite different from outsourcing key competitive information, such as R&D specifications or corporate strategy documents.

As Craig Balding (author of the Cloud Security blog) noted in his speech to the security conference BruCON, many cloud infrastructures are composed of patchworks of open source code that may bring their own underlying vulnerabilities. He also states that since public clouds are multi-tenant, your application could be affected by the vulnerabilities or defects of your neighbors' code. You must ensure that you understand the underlying infrastructure of the cloud to which you migrate your clients; you must also advise clients to include security in their cloud SLAs and terms of service. If data center security is not your area of expertise, I'd advise engaging a security expert before migrating your client to the cloud, as service interruptions or vulnerabilities make for embarrassing and tense customer interactions. As Balding says, many cloud vendors are performing "on boarding audits" to reassure prospective customers that their level of security is sufficient.

2.2 Vendor management:

When the cloud is your IT platform, and it's in the hands of an outside firm, how do you ensure that their technical or business problems won't become yours?

Speaking of SLAs and terms, the experience of migrating to outsourced providers should give you a good starting roadmap, but the terrain in the clouds is different. Since the whole idea behind cloud computing is to offer a standardized, multi-tenant infrastructure, cloud vendors may not offer the same level of custom SLAs as IT managers are accustomed to. Some of the large vendors, such as Amazon.com and Microsoft, are integrating management dashboards into their cloud offerings. As you can see from Amazon's Service Health Dashboard, the level of information offered is pretty basic and might not be enough to satisfy many organizations. For clients who need more information, Amazon offers Amazon Web Services Management Console, a customizable monitoring interface with a more robust set of

data. Other cloud vendors are following, and cloud management startup firms are springing up to address this need. Still, you need to assist clients so the cloud services they dial up are manageable and can be monitored sufficiently to ensure they won't have interruptions and performance issues.

2.3 Technical integration:

The technical issues are also complex. Most firms that migrate to the cloud do so in a hybrid model, keeping certain key elements of their infrastructure in-house and under their direct control, while outsourcing less sensitive or core components. Integrating internal and external infrastructures can be a technical quagmire.

Go to the Web site of cloud vendors such as Joyent, and you'll find on-demand cloud services that can be purchased in real-time with only a credit card. While cloud services can be easy to purchase, does that mean they'll be easy to integrate into your current IT infrastructure? Cloud vendors expect customers to provide, or to develop jointly, a "virtual image" that specifies their basic server configuration, which is then built inside the cloud and offered as a service. (This implies that the organization *has* a standard configuration!) It also requires the IT team to have the skill set to create a VM template that includes the infrastructure, the application, and the security required by the enterprise. Force.com, Salesforce.com's cloud offering, is leading the way by offering integration as a service on top of its cloud offerings. You must help clients develop the "golden image" that will be the basis for their cloud server configuration, and then integrate that cloud into the "hybrid cloud" with their existing data centers and applications.

2.4 Process and culture:

There are also the ever-present political and cultural landmines. When anyone with a credit card can surf to the Web site of a public cloud vendor and dial up teraflops of cloud capacity, how does IT maintain control of its application architecture (or does it?).

We've seen that cloud services are available with a credit card. When IT power becomes cheap and easily accessed, IT's control over its internal customers can be diluted, as we saw with the initial explosion of client-server computing. When a corporate department wasn't getting what it wanted from IT, or wasn't getting it fast enough, they simply went out and bought a server and a cheap, shrink-wrapped application and stuck it under the desk. The nightmare this caused in terms of IT consistency, integrity, and management is nothing compared to the potential for disruption of a dial-up cloud.You must work closely with clients to ensure that when they bring the cloud into the enterprise it's done with all the required procedural safeguards in place. IT needs to be involved in the decision of which applications are cloud-eligible to ensure that sensitive data is protected and available. Most of all, you can help clients develop the processes that will keep them in control of their infrastructure while not becoming a roadblock to the innovation and cost savings that the cloud can offer.

2.5 Blind faith playing with fire

Any assumption that the "write once, run anywhere" adages of old will extend fully to the cloud's service-based model of computing delivery are nothing more than a combination of blind faith playing with fire.So is it simply a case of wrapping

a cloud-compliant shell around applications so that they can still function and breathe at 40,000 feet? Or is there a more pressing need to rewrite, recompile, remodel, restructure, re-architect and redeploy our applications to the SaaS model?We know that some cloud computing progression will not involve migration as such, ie the more flexible working model of processing and storage capacity creates an opportunity (an almost incumbent responsibility if you like) to 'innovate' rather than migrate – and it is at this point where new completely applications may be born unto the world.But what about integrating physical applications with cloud-based services – where should we start?Companies such as Virtustream talk about the possibility of taking physical applications to the cloud without rewriting them. "Over 80 per cent of the world's business applications are legacy-based today, written before virtualization and the cloud, and as a result fear has been preventing many businesses from migrating their applications to the cloud. Virtustream can use its patented micro-VM (μVM) technology to wrap around legacy application and even assure cloud performance, without the need for re-writing them," said the company's CEO & CTO Kevin Reid."Fear over where data is 'up there' and concerns of being in breach of compliance still holds people back, but we can identify where all data is, down to the node and spindle as well as offer commercial SLAs on based on workload performance," added Reid.But this statement sounds like a strange claim to make doesn't it? ie 80 per cent of the world's business applications are legacy-based. Surely 100 per cent of the world's business applications are legacy-based!

2.6 Deterioration of customer care & service quality:

Third party cloud infrastructure solutions present a risk to customer care and overall service quality for support managers, support engineers and customer care staff. Support managers and engineers are at risk of becoming dependent upon a cloud service provider which they have no control over and at risk of requiring additional resources to do the migration and deal with short term issues that arise subsequent to the migration (e.g. shortfalls in cloud operations knowledge resulting in tasks taking temporarily longer to complete). Support managers and engineers specifically risk becoming dependent upon a cloud service provider for resolving hardware and network issues. This is a risk as it could result in the deterioration of service quality that the support manager would not be able to control. Support managers also risk temporarily requiring more resources to cope with migration and also the relative lack of knowledge and experience held by support staff regarding cloud systems. This is a risk because staff may initially require more time to perform the same tasks due to having to learn how-to perform tasks in the cloud environment which could compromise service quality and customer

service. Customer care staff are also at risk of not being able to offer the existing levels of customer service as it may take longer to resolve customer queries as cooperation with external service providers may become necessary. This is a risk because response times to deal with customer queries may increase resulting in back-logs and cascades of additional work as customer call back for progress updates and will result in customer care staff dissatisfaction.

2.7 Uncertainty with new technology:

Third party cloud infrastructure implementations present a risk to the finance/business development staff as it may open the organization to long-term volatility derived from market forces associated with the costs of using a cloud and data transfer costs. This is a risk as the medium to long-term viability of a cloud solution versus an internal hosting solution are uncertain. Additionally, switching to external hosting decreases the certainty of customer lock-in interms of software support contracts as now the hardware is maintained externally and therefore the company can no longer make the case that it offers an 'all-in-one' maintenance contract which avoids having to deal with multiple contactors. Another consideration is the loss of in-house expertise resulting in additional barriers to bringing the system back in-house if the cloud provider is inadequate.

III. DATA MIGRATION ACROSS THE CLOUDS

The amount of time it takes to complete the actual migration of objects and data from one database is relatively less than the amount of time it takes to complete an overall migration from assessment to production rollout. Migrations of one relational database to another are comparatively easier than migrations of a non- relational database to a relational database, because the organization of objects in a relational database is quite similar compared to non-relational databases such as hierarchical and network databases. All major relational database vendors also offer tools that provide robust migration capabilities in an automated fashion. Regardless of the level of automation and success factor of any migration tool, however, sometimes manual intervention will be required when migrating from one database to another. Database migration tasks can be divided into the following categories:

- **Database schema migration**
- **Data migration**
- **Database stored program migration**
- **Application migration**
- **Database administration script migration**

Of all the migration tasks listed, the application migration task requires the most manual effort, although new tools and technologies are being developed to facilitate this task.

3.1 Database Schema Migration

Database schema migration essentially involves migration tables, indexes, and views in a database. Relational databases are similar in terms of how their data is organized in tables and indexes, but they are different in terms of additional extensions to these tables and indexes that are designed to improve performance and facilitate devel- opment. Most migration tools can convert the database schema relatively quickly and accurately. Target-specific database schemas can also be generated from modelling tools such as Erwin. These are the most important things to consider during database schema migration:

3.1.1 Ensuring completeness of the schema

It is necessary to ensure that all objects from the source database have been migrated over to the target database. It is very common to have multiple schemas and databases to support an application. Having circular dependencies among multiple schemas and databases may result in errors during schema creation on the target database, as some of these dependencies may not exist when a particular schema is

being migrated. After creating all the schemas in Oracle, all the objects that are marked as invalid need to be recompiled and verified to ensure that they are migrated successfully.

3.1.2 Tables with system functions as DEFAULT value clauses on columns

Many databases support having system functions as the DEFAULT value clauses on table columns. In almost all cases, these system functions do not exist in the Oracle database. As a result, some tables may not be created in Oracle, making other dependent objects invalid. It is recommended that you analyze the log resulting from the schema creation task, and isolate and rectify such errors.

3.1.3 Using clustered indexes

Clustered indexes in databases such as Sybase allow data storage in a physically sorted fashion to match the logical order (index). As data is added, it is sorted and stored in the order defined by the clustered index. This helps to reduce the time it takes to return the sorted data and to retrieve data by co-locating the index as well as the actual data in the same object. The Oracle database provides similar functionality with index-organized tables (IOTs). In IOTs, the primary key columns and the non-key data are stored in the same object. This helps users avoid having to look up data in tables separately, after index lookups, while executing a query in Oracle.

3.1.4 Creating database users and role assignment

Proper database roles and privileges on objects must be assigned to users. Schema and object-level privi- leges can be grouped into roles and assigned to users as needed. Creating roles and granting them to users can help in managing many object-level privileges.

3.1.5 Changing object names

Any changes to the database object names due to restrictions in the database, as discussed in the "Analysis and Design" section of this chapter, need to be identified and shared with all team members so that they can make suitable changes in their applications or other database components.

3.1.6 Partitioning database tables

Oracle allows large tables to be partitioned into smaller segments for management ease and for better performance due to the database query optimizer's ability to prune partitions during query execution, resulting in a reduction in the overall amount of data scanned. Based on data volume, performance, and manageability requirements, some tables may be chosen for partitioning. Although many relational databases support table parti- tioning, they implement this feature differently in terms of the methods allowed for partitioning, such as range, hash, and composite partitioning.

3.2 Data Migration

After database schema migration, some representative data from the source database is migrated to the target database to enable testing and to ensure that the data migration scripts or tools chosen for the task are configured properly. The most common approach for data migration is undoubtedly the use of scripts that execute database utilities to export data from the source database and import it into the target database (Oracle), because they are easy to use and are free.

Regardless of the tools and scripts used to perform data migration, migrations of very large databases require planning. When migrating very large databases (those with at least a few terabytes of data) it is important to have the right data migration strategy, have the appropriate tools, and, most importantly, use appropriate database features such as partitioning and compression. Migration of large databases is fraught with challenges, among them a narrow window of time and lack of system resources (e.g., staging areas for data files). The following data extraction and loading strategies can optimize the data extraction, transfer, and loading processes:

• Parallel extraction of data from the source database
• Loading of data into the target database in parallel
• Using multithreaded processes for data loading
• Avoidance of index maintenance during the data loading process
• Reduction of I/O operations and use of staging areas via named pipes for data transfer between source and target databases

3.3 Database Stored Program Migration

The task of migrating database stored programs includes migration of stored procedures, triggers, and views which, in many relational databases, are used for implementing critical business logic. In databases such as Microsoft SQL Server and Sybase, stored procedures and triggers are used extensively by developers to support simple functions (e.g., the CRUD operations CREATE, READ, UPDATE, and DELETE). However, using stored procedures exclusively for CRUD operations can result in inflexibility because the type of operation executed against a table is limited by the functionality implemented in the stored procedure.

Major tasks associated with stored program migration are:

3.3.1 Cleaning and optimizing code

Oracle SQL Developer and other migration tools support migration of stored programs very well. However, it is recom- mended that you test these converted stored procedures and triggers for accuracy and efficiency of the converted code. Developers can implement a simple business requirement in many ways, making it harder for tools to optimize all such coding techniques in the converted code. Stored procedures and functions with hundreds of lines of code or more should be verified and tested for efficiency in terms of database feature usage as well as optimized coding practices.

3.3.2 Handling errors in stored procedures and triggers

For applications that depend heavily on stored procedures and triggers, it is very common to see nested stored procedure calls. Automated migrations may not be able to handle error handling for nested stored procedure invocation. Therefore, it is necessary to pay close attention to error handling, especially for nested stored procedure invocations.

3.3.3 Using temporary tables extensively

Some database developers use temporary tables extensively to simplify queries and avoid writing a complex query involving several tables. Early versions of some databases also had restrictions on the number of tables that could be joined in a query efficiently. Therefore, migrating stored procedures with lots of temporary tables warrants a closer look so that they can be avoided and can result in simplified

code that leverages the native features of an Oracle database. Typically, migration tools maintain a one- to-one mapping of temporary tables during migration from one database to another. But important stored procedures which are executed very often and have demanding performance requirements should be examined thoroughly to elimi- nate unnecessary temporary tables in the new environment.

3.3.4 Converting stored procedures into functions

The Oracle database does not support returning results to callers using the RETURN verb in stored procedures. This verb is only allowed in Oracle stored functions and not in stored procedures. However, it is very common to find Sybase and Microsoft SQL Server stored procedures using the OUT parameter as well as the RETURN verb to pass values and data to the caller. Converting these stored procedures into functions in Oracle also results in a different call signature (i.e., the syntax for executing a stored proce- dure versus executing a stored function is different because stored functions in Oracle must return a value).

3.3.5 Determining the impact of stored procedures returning result sets on Java applications (JDBC)

The Oracle database returns result sets to caller programs via explicitly defined OUT variables in stored procedures. However, other data- bases return multiple result sets implicitly, without having to declare variables to do so.

This results in additional changes to Java programs when migrating toOracle, such as declaring additional variables, binding, and explicit access of these variables for result set data.

3.4 Application Migration

Application migration or porting can result from either migrating an application from one environment to another due to a complete rewrite, or simply from an underlying database platform that is being migrated to a new platform such as Oracle. Typically, application development falls into two categories:

3.4.1 Customized application development

In this category, applications are generally developed in-house, by IT organizations, to support business functions. These applications almost always try to leverage all the native features of the database platform, as well as other IT systems in the organization, to drive maximum performance and tighter integration. As a result, applications tend to be heavily dependent on the database platform in which they were initially devel- oped. As a result, any change to the database platform may result in changes to the applications. Features and functionalities leveraged by these applications also depend on the developer's skill set. Developers try to use the features they are most comfortable with. Once an application becomes obsolete due to a lack of the skills required to maintain its features, or due to the application becoming too brittle to add new features, the application is migrated to a new environment.

3.4.2 Generic application development (or packaged applications)

Typically, this category applies to independent software vendors (ISVs). ISVs develop generic application software that caters to a particular industry or a vertical market. They also tend to develop applications that do not depend heavily

on the database. In fact, major ISVs offer versions of applications based on a particular database platform. Migration of a packaged application from one database to another involves installing and configuring the new version of the packaged application and importing the data and all the customizations from the original application. This is by no means a trivial task, because thorough testing needs to be done after the migration. From time to time, ISVs are forced to add support for new data- bases to their application software due to customer demand. They are also under pressure to maintain a single or as few codebases as possible to reduce the effort involved in managing multiple codebases, each catering to a different database, because this means that if they have to implement a new feature, they will have to modify all the application codebases in a similar fashion and ensure consistency across them.From a migration perspective, customized applications are always migrated to new database platforms fully, because there is no need for them to support both the old and new database platforms in the long run. These applications can be changed to take full advantage of the new database platform. But ISVs need to support all existing database platforms, even as they add support for new databases. So, for them, it becomes a porting effort because they are simply adding more code to an existing application so that it will also work with the new database. ISVs try to reduce the application software codebase by using conditional coding practices such as conditional branches to a different piece of code, depending on the database platform on which it is deployed. Very large enterprise resource planning (ERP) software packages usually have separate codebases for each database.

As we mentioned when we were discussing the migration assessment phase,understanding the impact of database platform migration on applications is very important. Applications depend on the database platform in many ways:

3.4.3 Database-specific connection information

Every database requires certain information to establish a connection with it. In the event of a database change, this information has to be updated in the applications that connect to a specific database. If every single program in an application connects to the database directly, instead of relying on a central database access layer, this otherwise trivial task becomes a challenge. This task can be automated through the use of scripts from the operating system to search and replace appropriate connection strings in application programs.

3.4.4 Use of database

specific parameters ODBC/JDBC drivers for database vendors have different parameters to support different requirements, such as transaction control, date/timestamp formats, and so forth. The Oracle JDBC driver, by default, enables AUTO COMMIT on a connection. This might create problems, especially when calling a database stored procedure in Oracle which leverages global temporary tables. Having set the AUTO COMMIT by default, the data in temporary tables will be deleted after any data manipulation statement (INSERT, DELETE, or UPDATE). To avoid this scenario, AUTO COMMIT for a JDBC connection should be explicitly disabled. For example:

```
Conn.setAutoCommit(false);
```

• Use of database-specific SQL statements Using database-specific SQL statements with proprietary extensions requires changes when the database platform changes. It is a big challenge to identify how many application programs need to be changed because of their usage of SQL statements that do not conform to American National Standards Institute (ANSI) SQL standards or that are not supported by the Oracle database. In the assessment phase, there is a great deal of emphasis on identifying such programs and their database inter- actions in general (i.e., calling stored procedures, result set processing, embedded SQL usage, etc.).

• Invoking database stored procedures and functions that return result sets Applications using ODBC/OLEDB drivers generally do not need to be modified when the database is migrated to Oracle. However, as of the latest release of Oracle Database 11g R2 (11.2.0.1), Java applications using the Oracle JDBC driver invoking stored procedures returning result sets from the database need to be modified to accommodate Oracle-specific requirements in terms of including bind variables for result sets, processing of multiple result sets, and similar functionality. Hopefully, these changes will not be necessary in future releases of the Oracle database.

•APIs for manipulation of large objects There are differences in JDBC APIs used for manipulating large objects in Oracle as compared to databases such as Informix.

IV. POINTERS FOR A SUCCESSFUL DATABASE MIGRATION

4.1 Assess database size

Database sizing will determine what hardware is required, and how much storage and what instance will be needed after migration. This can be undertaken by the internal IT team itself.

4.2 Test applications before data migration

The applications the service provider uses to connect to the database have to be fine-tuned to the applications that will use the database. Applications running on the cloud database should also be compatible with cloud infrastructure, and provide better performance than the in-house set-up. The cloud datacenters may not be in the vicinity, and there may be high latency issues. Applications should be able to perform in such situations. Raise the issue with your service provider, and make sure you're both on the same page..

4.3 Data confidentiality is a deal maker

To begin with, you might want to migrate only those databases and applications which are not mission critical. First migrate those databases which can be hosted in environments that may not be trusted.

4.4 Design the service level agreement (SLA) document carefully

There are applications which will require 99.99% up-time. Make sure scheduled down-times don't interfere with your business needs.

4.5 Ensure scalability

The main attraction of a database migration to the cloud is immediate scalability. Services and infrastructure should ideally be scalable on the fly. Yes, that will have to be negotiated with the provider. Keep the service vendor in the loop about your business growth plans.

4.6 Mind your OS

Finding the operating system (OS) that works well with your databases is crucial. For example, Oracle is available for Linux as well as Windows. Although both serve the same purpose, there will be a huge difference performance-wise. Check for the same version of the OS on the cloud.

4.7 Eliminating garbage will reduce costs

Cleansing of data becomes very important as costing depends on the size of the data. As database size grows, costs will also go up. Make sure to eliminate garbage data from the database before migrating it.

V. WAYS TO OVERCOME HASSLES

During your cloud database migration you may have to deal with performance and security issues. Here is how this can be tackled with ease.

5.1 Security

Your public cloud host could potentially be untrusted. It can reside anywhere, and there is no control of the customer over this aspect. One way out is to implement a private cloud. Factor this into your SLA. The provider's job is to provide the infrastructure, make the data available, and adhere to the security policies of the agreement. The data massaging or cleansing activity must be undertaken in-house because in principle, the provider should not view or process any data from your database.

5.2 Applications performance may vary on the cloud

Keep in kind that data will travel over a remote network and not just a LAN after the database migration. There may arise a need for re-writing codes. Some applications will already be cloud-compatible, while others may not work at all. For example, Oracle has partnered with Amazon, but Oracle does allow for other service providers to host its databases. Know where your provider stands on knowledge of the various applications and databases that are in migration.

5.3 Multiple database migrations

Moving multiple databases can be a challenge if any application sdepend on all of them. In such a scenario, the entire structure will have to be migrated to the cloud. The difficulty lies in finding a vendor who will host the multiple database set-up. In general, the migration of one or two databases to the cloud is more feasible than migrating many.

VI. CONCLUSION

Cloud migration is the process of moving data, applicationsor other business elements from an organization's onsite computers to the cloud, or moving them from one cloud environment to another.

Cloud migration sometimes involves moving data or other business elements between cloud environments, which is known as cloud-to-cloud migration. The process of transitioning to a different cloud provider is known as cloud service migration. In any case, successful migration to a service provider's environment may require the use of middleware, such as a cloud integration tool, to bridge any gaps between the vendor's and the customer's (or other vendor's) technologies.

Transitioning to the cloud or between cloud environments presents the usual IT issues, but the problems are compounded by having data stored and managed remotely, by external organizations and often in multiple locations. Among these issues are special considerations for privacy, interoperability, data and application portability, data integrity, business continuity, and security.

ACKNOWLEDGMENT

REFERENCES

[1] F. A. Alvi1, B.S Choudary,N. Jaferry, E.Pathan.- A review on cloud computing security issues & challenges.

[2] B.Meena,Krishnaveer ,Abhishek Challa- Cloud Computing Security Issues with Possible Solutions.

[3] Dikaiakos, M.D; Katsaros, D.; Mehra, P.; Pallis, G.; Vakali, A.; (2010), "Cloud Computing Distributed Internet Computing for IT and Scientific Research".Vol.13 ,pp 10, Sept.-Oct. 2009.

[4] Shuai Z; Shufen Z; Xuebin C; Xiuzhen H; (2010), "Cloud Computing Research and Development Trend", 2nd International conference on Future Networks, 2010. ICFN ' 10. pp 23, 22-24 Jan 2010.

[5] Chang, L, Ti ; Chin L; Chang, A.Y.; Chun J, C;(2010), " Information security issue of enterprises adopting the application of cloud computing", IEEE 2010 Sixth International Conference on Networked Computing and Advanced Information Management (NCM),pp 645, 16-18 Aug. 2010.

[6] R. Maggiani; (2009), "Cloud computing is changing how we communicate," 2009 IEEE International Professional Communication Conference, IPCC 2009,Waikiki, HI, United states ,pp 1, 19-22 July.

[7] Geng L; David F; Jinzy Z; Glenn D; (2009), "Cloud computing: IT as Service, "IEEE computer society IT Professional", Vol. 11, pp.10-13, March-April 2009.

[8] Basit Ali; (2009), "Ufone Launches Uconnect", published in TelecomPK.Net,12 August 2009.

[9] Muzzammil Sheikh; (2011), "PTCL Launched EVO USB become Wi-Fi Hotspot", The Frontier Star (Northwest Frontier Province, Jan 26 2011 Issue.

[10] Grobauer, B.; Walloschek, T.; Stocker,E.;(2011), "Understanding Cloud Computing Vulnerabilities",5487489 searchabstrSecurity & Privacy, IEEE, Vol 9, pp 50.

[11] Gansen Z; Chunming R; Jin L; Feng Z; Yong T; (2010),,"Trusted Data Sharing over Untrusted Cloud Storage Providers",2010 IEEE Second International Conference on Cloud Computing Technology and Science (CloudCom), pp 97, Nov. 30 2010-Dec. 3 2010.

[12] Pearson, S.; (2009), "Taking account of privacy when designing cloud computing services",5071532 searchabstract CLOUD '09. ICSE Workshop on Software Engineering Challenges of Cloud Computing, 2009. pp 44, 23-23 May 2009.

[13] Kresimir P; Zeljko H; (2010), "Cloud computing security issues and challenges", MIPRO 2010, May 24-28, 2010, Opatija, Croatia.

[14] Minqi Z; Rong Z; Wei X; Weining Q; Aoying Z; (2010),"Security and Privacy in Cloud Computing: A Survey", Sixth international conference on Semantics Knowledge and Grid (SKG), pp 105, 1-3 Nov. 2010.

[15] Popovic K; Hocenski Z; (2010), "Cloud computing security issues and challenge", 5533317searchabstractMIPRO, 2010 Proceedings of the 33rd International Convention , pp 344,24-28 May 2010.

[16] Jensen, M.; Schwenk, J.; Gruschka, N.; Iacono, L.L.; (2010), "On Technical Security Issues in Cloud Computing", IEEE International Conference on Cloud Computing, 2009. CLOUD '09, pp 109, 21-25 Sept. 2009. 5708519 searchabstract

[17] Jianfeng Y; Zhibin C; (2010), "Cloud Computing Research and Security Issues", IEEE 2010 International Conference on Computational Intelligence and Software Engineering (CiSE), pp1, 10-12 Dec 2010.

[18] Jansen, W.A.; (2010), " Cloud Hooks: Security and Privacy Issues in Cloud Computing5719001 IEEE 2011 44th Hawaii International Conference on System Sciences (HICSS), pp1, 4-7 Jan. 2011.

[19] Tian L.Q; NI Y,LING; (2010) , "Evolution of user Behavior Trust in Cloud Computing", 2010 International Conference on Computer Application and System Modeling (ICCASM 2010),Vol. 7,pp V7-567, 22-24 Oct. 2010.

[20] Mathur, P; Nishchal, N.; (2010), "Cloud Computing: New challenge to the entire computer industry", 2010 1st International Conference on Parallel, Distributed and Grid Computing (PDGC - 2010), pp 223.

[21] Yuefa D; Wu B; Yaqiang G; Zhang Q; Tang C; (2009), " Data Security Model for Cloud Computing", Proceedings of the 2009 International Workshop on Information security and Applications (IWISA 2009)

[22] Dean and S. Ghemawat; (2010), "MapRduce: Simplified data processing large clusters", communication of the ACM, Vol.51, pages 107-113.

[23] Xue J; Zhang J.J; (2010),"A Brief Survey on the Security Model of Cloud Computing",2010 Ninth International Symposium on Distributed Computing and Applications to Business, Engineering and Science.

[24] *Migrating to the Cloud By Tom Laszewski P Nauduri.*