IOSR Journal of Computer Engineering (IOSR-JCE)
e-ISSN: 2278-0661, p- ISSN: 2278-8727Volume 16, Issue 1, Ver. IX (Feb. 2014), PP 113-119
www.iosrjournals.org

Best Practices in Implementation of Cloud

Navjot Kaur
(Department of Computer Science, S.G.T.B Khalsa College/ University of Delhi, India)

Abstract: *Cloud computing is one of the latest developments in the IT industry also known as on-demand computing. It provides the full scalability, reliability, high performance and relatively low cost feasible solution as compared to dedicated infrastructures. This technology has the capacity to admittance a common collection of resources on request. It is proving extremely striking to cash-strapped IT departments that are wanted to deliver better services under pressure. The hype around Cloud Computing is undeniable, especially regarding how it can result in huge savings in terms of I.T. costs for an enterprise. This has persuaded many enterprises to shift to cloud based software services. Cloud based software is the most widely demanded feature and most of the Information Technology providers are suggesting the cloud based solutions to their clients. But having many advantages for IT organizations cloud has some issues that must be considered during its deployment. The main concern is security privacy and trust. These issues are arises during the deployment of mostly public cloud because in public cloud infrastructure customer is not aware where the data store & how over the internet. Apart from these issues, there are many other challenges which should be dealt with. Aim of this paper is to device certain best practices for implementing cloud computing which would result in more stable performance when facing any kind of threat.*

Keywords: *Cloud computing, Information Technology, Security in Cloud Computing,*

I. INTRODUCTION

Cloud computing is the use of computing resources (hardware and software) that are delivered as a service over a network or Internet. The name comes from the use of 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. [10]

The latest developments in the field of Information Technology and Enabled Services offered people comforts and convenience. Cloud Computing provides the full scalability, reliability, high performance and relatively low cost feasible solution as compared to dedicated infrastructures. It is the application provided in the form of service over the internet and system hardware in the data centres that gives these services. This technology has the capacity to admittance a common collection of resources on request. It is proving extremely striking to cash-strapped IT departments that are wanted to deliver better services under pressure. [7]

The scope of this paper is to present an overview of Cloud Computing, various types of Cloud Computing services available and identify the various issues that the organizations have encountered during their implementation. After listing out all these issues, I will be focusing on deriving the best practices for implementation of cloud computing and then will apply these best practices for finding the resolutions for these issues.

The main purpose of this paper is to present a basic framework for evaluating the past implementation of cloud computing in various organizations and listing out the common issues faced by these organizations. My effort is motivated by the rise of Cloud Computing providers and the question when it is profitable for the organization to implement Cloud by eliminating the prospective issues in implementation. More and more companies already embrace Cloud Computing services as part of their IT infrastructure. However, currently, there is no guide to tell the prospective issues to take care of before the implementation of cloud. With my work I want to give an overview of the technical aspects that a valuation approach to Cloud Computing must take into consideration.

II. TYPES OF COMPUTING – DEPLOYMENT MODELS

Public Cloud: When the cloud is made available for the general customer on pay per use basis, it is said to be a Public Cloud. Some of the popular examples of public cloud service providers are Amazon, Google and Microsoft. They own and operate the infrastructure and offer access through Internet.

Private Cloud: Private cloud is a cloud infrastructure operated solely for a single organization, whether managed internally or by a third party and hosted internally or externally. Undertaking a private cloud project requires a significant level and degree of engagement to virtualize the business environment, and it will require the organization to reevaluate decisions about existing resources. When it is done right, it can have a positive impact on a business, but every one of the steps in the project raises security issues that must be addressed in order to avoid serious vulnerabilities.

Table 1: Difference between Public and Private Cloud

Difference between Public and Private Cloud		
	Public cloud	**Private cloud**
Initial cost	Typically zero	Typically high
Running cost	Predictable	Unpredictable
Customization	Impossible	Possible
Privacy	No (Host has access to the data)	Yes
Single sign-on	Impossible	Possible
Scaling up	Easy while within defined limits	Laborious but no limits

Hybrid Cloud: Hybrid cloud is a composition of private and public clouds that remain unique entities but are bound together, offering the benefits of multiple deployment models. By utilizing "hybrid cloud" architecture, companies and individuals are able to obtain degrees of fault tolerance combined with locally immediate usability without dependency on internet connectivity. Hybrid cloud architecture requires both on-premises resources and off-site (remote) server-based cloud infrastructure.

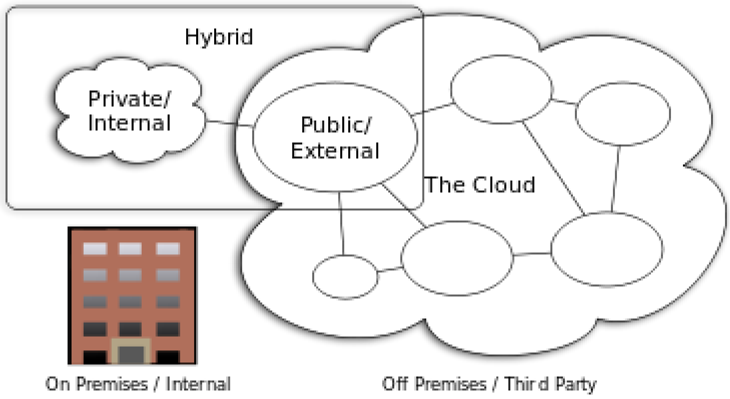

Figure 1: Types of Cloud Computing

III. MODELS OF CLOUD COMPUTING

There are many models for Public Cloud Computing. Some of the most widely used basic models for Public Cloud Computing are:

Basic Models

Infrastructure as a Service (IaaS)
Platform as a Service (PaaS)
Software as a Service (SaaS)

Infrastructure as a Service (IaaS) is a provision model in which an organization outsources the equipment used to support operations, including storage, hardware, servers and networking components. The service provider owns the equipment and is responsible for housing, running and maintaining it. The client normally pays on a per-use basis.

Amazon Web Services is the most widely used IaaS cloud computing platform today. Amazon provides a number of different levels of computational power for different pricing. The primary methods for data storage in Amazon Electric Compute Cloud (EC2) are Simple Storage Service (S3) and Elastic Block Storage (EBS). S3 is a highly scalable key-based storage system that transparently handles fault tolerance and data integrity. EBS provides a virtual storage device that can be associated with an elastic computing instance. S3 charges for space used per month, the volume of data transferred, and the number of metadata operations (in allotments of 1000). EBS charges for data stored per month. For both S3 and EBS, there is no charge for data transferred to and from EC2 within a domain (e.g., the U.S. or Europe)[13].

Platform as a Service (PaaS) is a way to rent hardware, operating systems, storage and network capacity over the Internet. The service delivery model allows the customer to rent virtualized servers and associated services for running existing applications or developing and testing new ones.

With PaaS, operating system features can be changed and upgraded frequently. Geographically distributed development teams can work together on software development projects. Services can be obtained from diverse sources that cross international boundaries. Initial and ongoing costs can be reduced by the use of infrastructure services from a single vendor rather than maintaining multiple hardware facilities that often perform duplicate functions or suffer from incompatibility problems. Overall expenses can also be minimized by unification of programming development efforts.

Software as a Service (SaaS) is a software distribution model in which applications are hosted by a vendor or service provider and made available to customers over a network, typically the Internet.

SaaS is becoming an increasingly prevalent delivery model as underlying technologies that support Web services and service-oriented architecture (SOA) mature and new developmental approaches, such as Ajax, become popular. Meanwhile, broadband service has become increasingly available to support user access from more areas around the world. [2]

IV. COMMON ISSUES IN IMPLEMENTATION OF CLOUD

Symantec’s 2013 Cloud Survey is a result of research conducted by ReRez in September-October 2012. The full study represents 3,236 organizations from 29 countries. Responses came from companies with a range of five to more than 5,000 employees. Of those responses, 1,358 came from SMBs and 1,878 came from enterprises. Below are some of the issues that were found in this survey[14]. I am listing down the identified issues as a result of this survey and other sources[1]. After these points, I’ll list down more issues that resulted as various implementations that have been done by IBM, Sapient and other similar organizations.

Rogue cloud implementations

According to the survey, rogue cloud deployments are one of the cost pitfalls. It is a surprisingly common problem, found in 89 percent of Indian enterprises and 92 per cent of Indian SMBs within the last year.

Among Indian enterprises who reported rogue cloud issues, 57 percent experienced the exposure of confidential information, and more than a third faced account takeover issues, defacement of Web properties, or stolen goods or services. The survey found that top rogue cloud issues for Indian SMBs include security, data protection and loss of confidential information. The challenge is escalating, with nearly half (48 percent) of Indian SMBs indicating that rogue cloud deployments are becoming more frequent. The most commonly cited reasons by Indian businesses for undertaking rogue cloud projects were to save time and money.

Cloud backup and recovery issues

Cloud is complicating backup and recovery. First, most Indian organizations use three or more solutions to back-up their physical, virtual and cloud data—leading to increased IT inefficiencies, risk and training costs. Furthermore, nearly two-thirds organizations have lost cloud data (60 percent of enterprises and 70 percent of SMBs), and most (80 percent) have experienced recovery failures.

Finally, most see cloud recovery as a slow, tedious process; 85 percent estimate it would take three or more days to recover from a catastrophic loss of data in the cloud.

Inefficient Cloud storage

Theoretically, cloud storage has the advantages of being quick to deploy, paying for only what is used and offering easy adjustment of capacity, all of which provide for high storage utilization rates. In practice, however, global cloud storage utilization is actually extremely low at just 17 percent. It’s even worse for SMBs, at just seven per cent. This is resulting in organizations paying for six times as much storage as they need. The problem is exacerbated because almost half of enterprises in India (48 percent) admit that little to none of their data is deduplicated, and 34 percent of SMBs indicate that half or more of their data is duplicate, all of which is leading organizations to pay for storage they don’t require.

Compliance and eDiscovery concerns

With growing regulatory and internal compliance frameworks, the survey revealed that two-thirds of Indian enterprises are concerned about not only meeting compliance requirements, but also proving it. However, nearly half (47 percent) have been fined for privacy violations in the cloud within the past 12 months. Organizations are performing equally poorly on the eDiscovery front: three-quarters of those who have received requests (76 percent) missed deadlines for delivering the requested information, potentially leading to fines or compromised legal positions. What’s worse, 34 percent never found the requested information.

Data in-transit issues

Organizations have all sorts of assets in the cloud – such as web properties, online businesses or web applications – that require SSL certificates to protect the data in transit whether it is personal or financial information, business transactions and other online interactions. The survey showed companies found managing many SSL certificates to be highly complex: Just 48 percent rate cloud SSL certificate management as easy and only 40 percent are certain their cloud-partner's certificates are in compliance with corporate standards.

The survey shows ignoring these hidden costs will have a serious impact on business. However, these issues are easily mitigated with careful planning, implementation and management:

- Focus policies on information and people, not technologies or platforms
- Educate, monitor and enforce policies
- Embrace tools that are platform agnostic
- Deduplicate data in the cloud

Symantec's 2013 Cloud Survey is a result of research conducted by ReRez in September-October 2012. The full study represents 3,236 organizations from 29 countries. Responses came from companies with a range of five to more than 5,000 employees. Of those responses, 1,358 came from SMBs and 1,878 came from enterprises[13].

Inexperience

For vendor management, new approaches and skills are required, and new ways of thinking are needed for strategic sourcing. Because cloud computing is so new, however, most organizations have yet to develop the necessary experience and skills. Practices that are well understood in terms of traditional outsourcing are unlikely to apply.

V. DEFINING BEST PRACTICES

Cloud reinforces highly scalable internet architecture which is built on some traditional concepts and some new and these new concepts entirely change the way applications are built and deployed. Implementation of cloud changes many processes, patterns, practices, philosophies but reinforces several service oriented principles. In this section some Best Practices are defined to build highly scalable applications[3].

Managing the Constraints: While implementing cloud in an organization, there will be cases when the IT managers will compare their system specifications with those being provided by the cloud. It should be noted that cloud provides abstract resources and they become powerful when they are combined with the on-demand provisioning model. For example, if the cloud does not provide them with exact or greater amount of RAM in a server, they should try using a distributed cache like memcached or partitioning their data across multiple servers.

Cloud service users should not be afraid and constrained when using cloud resources because it is important to understand that even if they might not get an exact replica of their existing hardware in the cloud environment, they have the ability to get more of those resources in the cloud to compensate that need.

In retrospect, when you combine the on-demand provisioning capabilities with the flexibility, you will realize that apparent constraints can actually be broken in ways that will actually improve the scalability and overall performance of the system.

Virtual Administration: The advent of cloud has changed the role of System Administrator to a "Virtual System Administrator". The cloud encourages automation of configuring servers and installing software because the infrastructure is programmable. System administrators need to move up the technology stack and learn how to manage abstract cloud resources using scripts.

Likewise, the role of Database Administrator is changed into a "Virtual Database Administrator" in which he/she manages resources through a web-based console, executes scripts that add new capacity programmatically in case the database hardware runs out of capacity and automates the day-to-day processes. The virtual DBA has to now learn new deployment methods (virtual machine images), embrace new models (query parallelization, geo-redundancy and asynchronous replication), rethink the architectural approach for data (sharding, horizontal partitioning, federating) and leverage different storage options available in the cloud for different types of datasets[5].

Design for failure and nothing will fail: In particular, one should assume that the hardware will fail. Assume that outages will occur. Assume that some disaster will strike the application hosted on cloud. Assume that one will be slammed with more than the expected number of requests per second someday. Assume that with time the application software will fail too. By being a pessimist, one ends up thinking about recovery strategies during design time, which helps in designing an overall system better.

If one realizes that things fail over time and incorporates that thinking into their architecture, build mechanisms to handle that failure before disaster strikes to deal with a scalable infrastructure, one ends up creating a fault-tolerant architecture that is optimized for the cloud.

Appropriate mechanisms needs to be built in to handle any kind of failure. For example, the following strategies can help in event of failure:

1. Have a coherent backup and restore strategy for the data and automate the same
2. Build process threads that resume on reboot
3. Allow the state of the system to re-sync by reloading messages from queues
4. Keep pre-configured and pre-optimized virtual images to support (2) and (3) on launch/boot
5. Avoid in-memory sessions or stateful user context, move that to data stores.

Good cloud architectures should be impervious to reboots and re-launches. If any particular instance on which controller thread was running dies, it should be brought up and resume the previous state automatically, as if nothing had happened.

Designing with an assumption that underlying hardware will fail, will prepare well for the future when it actually fails. This design principle will help in designing operation-friendly applications, as also highlighted in Hamilton's paper. If this principle can be extended to pro-actively measure and balance load dynamically, this might be able to deal with variance in network and disk performance that exists due to multi-tenant nature of the cloud[6].

Decouple the components: The key is to build components that do not have tight dependencies on each other, so that if one component were to die (fail), sleep (not respond) or remain busy (slow to respond) for some reason, the other components in the system are built so as to continue to work as if no failure is happening. In essence, loose coupling isolates the various layers and components of the application so that each component interacts asynchronously with the others and treats them as a "black box". For example, in the case of web application architecture, it is possible to isolate the app server from the web server and from the database. The app server does not know about your web server and vice versa, this gives decoupling between these layers and there are no dependencies code-wise or functional perspectives.

Questions that needs to be asked during the implementation:

1. Which business component or feature could be isolated from current monolithic application and can run standalone separately?
2. And then how to add more instances of that component without breaking my current system and at the same time serve more users?
3. How much effort will it take to encapsulate the component so that it can interact with other components asynchronously?

Decoupling the components, building asynchronous systems and scaling horizontally become very important in the context of the cloud. It will not only allow scaling out by adding more instances of same component but also allow to design innovative hybrid models in which a few components continue to run in on-premise while other components can take advantage of the cloud-scale and use the cloud for additional compute-power and bandwidth. That way with minimal effort, it is possible to "overflow" excess traffic to the cloud by implementing smart load balancing tactics.

One can build a loosely coupled system using messaging queues. If a queue/buffer is used to connect any two components together, it can support concurrency, high availability and load spikes. As a result, the overall system continues to perform even if parts of components are momentarily unavailable. If one component dies or becomes temporarily unavailable, the system will buffer the messages and get them processed when the component comes back up.

Implement Scaling: Scaling can be implemented in the following three ways:

1. Proactive Cyclic Scaling: Periodic scaling that occurs at fixed interval (daily, weekly, monthly, quarterly)
2. Proactive Event-based Scaling: Scaling just when you are expecting a big surge of traffic requests due to a scheduled business event (new product launch, marketing campaigns)
3. Auto-Scaling based on demand. By using a monitoring service, the system can send triggers to take appropriate actions so that it scales up or down based on metrics (utilization of the servers or network i/o, for instance)

To implement Scaling, one has to first automate the deployment process and streamline the configuration and build process. This will ensure that the system can scale without any human intervention. This will result in immediate cost benefits as the overall utilization is increased by ensuring the resources are closely aligned with demand rather than potentially running servers that are under-utilized[8].

Automate the Infrastructure: One of the most important benefits of using a cloud environment is the ability to use the cloud's APIs to automate the deployment process. It is recommended to take the time to create an automated deployment process early on during the migration process and not wait till the end. Creating an automated and repeatable deployment process will help reduce errors and facilitate an efficient and scalable update process.

The following steps should be undertaken to automate the deployment process:

- Create a library of "recipes" – small frequently-used scripts (for installation and configuration)
- Manage the configuration and deployment process using agents bundled inside an image
- Bootstrap the instances - On boot, instances should grab the necessary resources (code, scripts, configuration) based on the role and "attach" itself to a cluster to serve its function[12].

Think Parallel: The cloud makes parallelization effortless. Whether it is requesting data from the cloud, storing data to the cloud, processing data (or executing jobs) in the cloud, as a cloud architect, one needs to internalize the concept of parallelization when designing architectures in the cloud. It is advisable to not only implement parallelization wherever possible but also automate it because the cloud allows to create a repeatable process every easily.

When it comes to accessing (retrieving and storing) data, the cloud is designed to handle massively parallel operations. In order to achieve maximum performance and throughput, one should leverage request parallelization. Multi-threading the requests by using multiple concurrent threads will store or fetch the data faster than requesting it sequentially. Hence, wherever possible, the processes of a cloud application should be made thread-safe through a share-nothing philosophy and leverage multi-threading.

When it comes to processing or executing requests in the cloud, it becomes even more important to leverage parallelization. A general best practice, in the case of a web application, is to distribute the incoming requests across multiple asynchronous web servers using load balancer. In the case of batch processing application, master node can spawn up multiple slave worker nodes that processes task in parallel (as in distributed processing frameworks like Hadoop)[8].

Keep dynamic data closer to the compute and static data closer to the end-user: In general it's a good practice to keep all data as close as possible to your compute or processing elements to reduce latency. In the cloud, this best practice is even more relevant and important because one has to deal with Internet latencies. Moreover, in the cloud, the payment is for bandwidth in and out of the cloud by the gigabyte of data transfer and the cost can add up very quickly.

If a large quantity of data that needs to be processed resides outside of the cloud, it might be cheaper and faster to "ship" and transfer the data to the cloud first and then perform the computation. For example, in the case of a data warehousing application, it is advisable to move the dataset to the cloud and then perform parallel queries against the dataset. In the case of web applications that store and retrieve data from relational databases, it is advisable to move the database as well as the app server into the cloud all at once.

If the data is generated in the cloud, then the applications that consume the data should also be deployed in the cloud so that they can take advantage of in-cloud free data transfer and lower latencies. For example, in the case of an e-commerce web application that generates logs and click stream data, it is advisable to run the log analyzer and reporting engines in the cloud[11].

Conversely, if the data is static and not going to change often (for example, images, video, audio, PDFs, JS, CSS files), it is advisable to take advantage of a content delivery service so that the static data is cached at an edge location closer to the end-user (requester) thereby lowering the access latency. Due to the caching, a content delivery service provides faster access to popular objects.

VI. CONCLUSION

If there is one thing experts agree on, it is that cloud computing is here to stay. SMEs and start-up companies are already starting to reap the benefits of this new domain, and there is consensus that an increasing number of companies will replace local technology with cloud solutions in the years to come. There is also a dramatic increase in focus on digital security among SMEs. Many of the current systems are poorly protected and the adoption of cloud computing only further increases the complexity of securing data flows. Although many companies wish to priorities sustainability as part of their IT strategies, the tough financial times have forced them to focus on other things. This also raises the question as to what type of policy mix is needed to drive an agenda within the IT industry and how the short/medium term priorities can be addressed.

Continued growth of cloud computing will require vendors and firms to overcome its challenges together. Just as personal computers and servers shook up the world of mainframes and minicomputers, or as smartphones and tablets revolutionized the mobile commerce industry, cloud computing is bringing similar far-reaching changes to the licensing and provisioning of infrastructure and to methodologies for application development, deployment and delivery. But as cloud computing is in its infancy state, lack of standardization is hindering the usage of cloud, implementation of above defined best practices will not only help Cloud Service providers but also be expedient to cloud service users.

REFERENCES

Journal Papers:

[1] M. Vouk. Cloud Computing – Issues, Research and Implementations, Journal of Computing and Information Technology – CIT 16(4), 235 – 246, 2008.

[2] L.R.Rewatkar and U.A.Lanjewar, Implementation of Cloud Computing on Web Application, International Journal of Computer Applications (0975-8887), Volume 2 -No.8, June2012, pp 28-32

[3] J. Varia, "Architecting for the cloud: Best practices", May 2010

[4] M.Sharma, A.Mehra, H.Jola, A.Kumar, M.Misra and V.Tiwari, Scope of cloud computing for SMEs in India, Journal of Computing, Volume 2, Issue 5, May 2010, ISSN 2151-9617.

[5] B.D. Payne, M. Carbone, M. Sharif, and W. Lee. Lares: An architecture for secure active monitoring using virtualization. Security and Privacy, IEEE Symposium , 0:233-247, 2008

Books:

[6] Toby Velte, Anthony Velte, Robert Elsenpeter, Cloud computing, A practical approach(NewYork, McGraw-Hill, Inc. 2010)

Proceedings Papers:

[7] Alvi, F.A. , "A review on cloud computing security issues & challenges", Proc. 1st International Conference on Mobility for Life, 2012

[8] Li, Peng, and Lee W. Toderick. "Cloud in cloud: approaches and implementations." Proc. 2010 ACM conference on Information technology education, pp. 105-110. ACM, 2010.

[9] Zhang, L. J., Zhou, Q., 2009. CCOA: Cloud Computing Open Architecture. Proceedings of the 2009 IEEE International Conference on Web Services, pp. 607-616.

Online:

[10] NIST. Available: http://www.thecloudtutorial.com/nistcloudcomputingdefinition.html

[11] IBM. Available: http://www.ibm.com/ibm/cloud/cloudburst

[12] Security Guidance for Critical Areas of Focus in Cloud Computing V2.1‖, CSA, Dec 2009, Available at: https://cloudsecurityalliance.org/csaguide.pdf

[13] Amazon EC2. Available: http://aws.amazon.com/ec2.

[14] Symantec's 2013 Cloud Survey http://www.symantec.com/about/news/release/article.jsp?prid=20130115_01