RESEARCH ARTICLE OPEN ACCESS

IT Infrastructure Setup using Cloud Services

Swapnil K. Gajbhiye[1], Manoj S. Chaudhari[2]
1(MTech Computer Science & Engineering
Priyadarshini Bhagwati College of Engineering Nagpur, India
2 (Assist. Prof. & Head of Information Technology Department
Priyadarshini Bhagwati College of Engg,Nagpur

Abstract
Cloud Computing is emerging today as a commercial infrastructure that eliminates the need for maintaining expensive computing hardware. Through the use of virtualization, clouds promise to address with the same shared set of physical resources a large user base with different needs. Thus, clouds promises to provide a flexible IT architecture, accessible through internet from lightweight portable devices. Cloud computing is a set of IT services that are provided to a customer over a network on a leased basis and with the ability to scale up or down their service requirements.
Usually Cloud Computing services are delivered by a third party provider who owns the infrastructure. Many industries, such as banking, healthcare and education are moving towards the cloud due to the efficiency of services provided by the pay-per-use pattern based on the resources such as processing power used, transactions carried out, bandwidth consumed, data transferred, or storage space occupied etc. In a cloud computing environment, the entire data resides over a set of networked resources, enabling the data to be accessed through virtual machines.
In this work we present an evaluation of the usefulness of the current cloud computing services for infrastructure setup. We analyze the performance of the Amazon EC2 platform. We also compare using long-term traces the performance characteristics and cost models of clouds with those of other platforms.
We will provide the guideline for using cloud services for setting up any IT infrastructure on cloud.
This research paper also analyzes the key research challenges present in cloud computing and offers best practices to service providers as well as enterprises hoping to leverage cloud service to improve their bottom line in this severe economic climate.

Keywords: Cloud Architecture, Cloud Computing, Software as a Service (SaaS), Platform as a Service (PaaS),Infrastructure as a Service (IaaS), Mobile Cloud Computing (MCC).

I. INTRODUCTION

Cloud computing has recently reached popularity and developed into a major trend in IT
Cloud computing is a complete new technology. It is the development of parallel computing, distributed computing grid computing, and is the combination and evolution of Virtualization, Utility computing, Software-as-a-Service (SaaS), Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS).
Cloud is a metaphor to describe web as a space where computing has been pre installed and exist as a service; data, operating systems, applications, storage and processing power exist on the web ready to be shared.
To users, cloud computing is a Pay-per-Use-On-Demand mode that can conveniently access shared IT resources through the Internet.
Users can use the IT infrastructure with Pay-per-Use-On-Demand mode; this would benefit and save

the cost to buy the physical resources that may be vacant

The advantages of using cloud computing include:

i) reduced hardware and maintenance cost,

ii) accessibility around the globe

iii) flexibility and highly automated processes wherein the customer need not worry about mundane concerns like software upgradation

Cloud Computing provides technological capabilities—generally maintained off premises

In the real scenario, they are renting the physical infrastructure, platforms and applications within a shared architecture. Cloud offerings can vary from virtual infrastructure, computing platforms, centralized data centers to end-user Web-Services and Web applications to enormous other focused computing services.

OVERVIEW

Cloud Computing is the combination of a technology, platform that provides hosting and storage service on the Internet. In such an environment users need not own the infrastructure for various computing services

Cloud Computing is a general term for anything that involves delivering hosted services over the Internet. Instead of a static system architecture, Cloud Computing supports the ability to dynamically scale up and quickly scale down, offering cloud consumers high reliability, quick response times, and the flexibility to handle traffic fluctuations and demand.

Cloud Computing also supports multi tenancy, providing systems configured in such a way that they can be pooled to be shared by many organizations or individuals [8]. Virtualization technology allows cloud vendors to convert one server into many virtual machines, thereby eliminating client-server computing with single-purpose systems.

This maximizes hardware capacity and allows customers to leverage economies of scale.

Benefits of Cloud computing are enormous. The most important one is that the customers don't need to buy the resource from a third party vendor, instead they can use the resource and pay for it as a service thus helping the customer to save time and money.Cloud is not only for Multinational companies but it's also being used by small and medium enterprises.

CLOUD COMPUTING BUILDING BLOCKS

I. DEPLOYMENT MODELS

In the cloud deployment model, networking, platform, storage, and software infrastructure are provided as services that scale up or down depending on the demand as depicted in figure 1. The Cloud Computing model has four main deployment models which are:

Private Cloud: Private cloud is a new term that some vendors have recently used to describe offerings that emulate cloud computing on private networks. It is set up within an organization's internal enterprise datacenter.

It is set up within an organization's internal enterprise datacenter. In the private cloud, scalable resources and virtual applications provided by the cloud vendor are pooled together and available for cloud users to share and use. It differs from the public cloud in that all the cloud resources and applications are managed by the organization itself, similar to Intranet functionality. Utilization on the private cloud can be much more secure than that of the public cloud because of its specified internal exposure. Only the organization and designated

stakeholders may have access to operate on a specific Private cloud.

Public Cloud: Public cloud describes cloud computing in the traditional mainstream sense, whereby resources are dynamically provisioned on a fine-grained, self-service basis over the Internet, via web applications/web services, from an off-site third-party provider who shares resources and bills on a fine-grained utility computing basis. It is typically based on a pay-per-use model, similar to a prepaid electricity metering system which is flexible enough to cater for spikes in demand for cloud optimization. Public clouds are less secure than the other cloud models because it places an additional burden of ensuring all applications and data accessed on the public cloud are not subjected to malicious attacks. Examples of a public cloud include Microsoft Azure, Google App Engine.

Hybrid Cloud: Hybrid cloud is a private cloud linked to one or more external cloud services, centrally managed, provisioned as a single unit, and circumscribed by a secure network. It provides virtual IT solutions through a mix of both public and private clouds. Hybrid Cloud provides more secure control of the data and applications and allows various parties to access information over the Internet. It also has an open architecture that allows interfaces with other management systems. Hybrid cloud can describe configuration combining a local device, such as a Plug computer with cloud services. It can also describe configurations combining virtual and physical, collocated assets - for example, a mostly virtualized environment that requires physical servers, routers, or other hardware such as a network appliance acting as a firewall or spam filter. An example of a Hybrid Cloud includes Amazon Web Services (AWS).

Community Cloud: Infrastructure shared by several organizations for a shared cause and may be managed by them or a third party service provider and rarely offered cloud model. These clouds are normally based on an agreement between related business organizations such as banking or educational organizations. A cloud environment operating according to this model may exist locally or remotely. An example of a Community Cloud includes Facebook

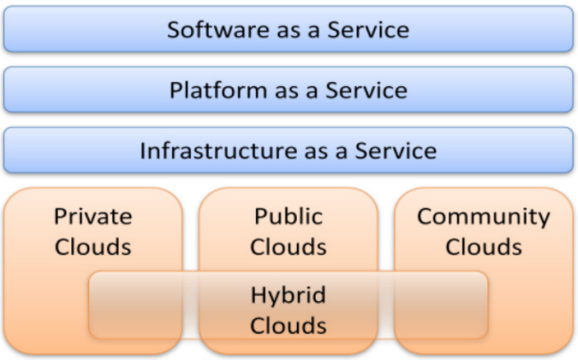

Figure 1: Cloud Deployment Models

I. SERVICE MODELS

Cloud service models are commonly divided into SaaS, PaaS, and IaaS that exhibited by a given cloud infrastructure. It's helpful to add more structure to the service model stacks: Fig. 2 shows a cloud reference architecture that makes the most important security-relevant cloud components explicit and provides an abstract overview of cloud computing for security issue analysis.

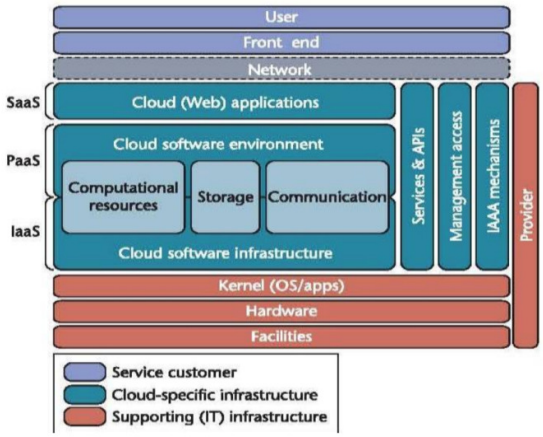

Figure 2: Cloud reference architecture

Software as a Service(SaaS): Cloud consumers release their applications in a hosting environment, which can be accessed through networks from various clients (e.g. Web browser, PDA, etc.) by application users. Cloud consumers do not have control over the cloud infrastructure that often employs multi-tenancy system architecture, namely, different cloud consumers' applications are organized in a single logical environment in the SaaS cloud to achieve economies of scale and optimization in terms of speed, security, availability, disaster recovery and maintenance. Examples of SaaS include SalesForce.com, Google Mail, Google Docs, and so forth.

Platform as a Service(PaaS): PaaS is a development platform supporting the full "Software Lifecycle" which allows cloud consumers to develop cloud services and applications (e.g. SaaS) directly on the PaaS cloud. Hence, the difference between SaaS and PaaS is that SaaS only hosts completed cloud applications whereas PaaS offers a development platform that hosts both completed and in-progress cloud applications. This requires PaaS, in addition to supporting application hosting environment, to possess development infrastructure including programming environment, tools, configuration management, and so forth. An example of PaaS is Google AppEngine.

Infrastructure as a Service (IaaS): Cloud consumers directly use IT infrastructures (processing, storage, networks and other fundamental computing resources) provided in the IaaS cloud. Virtualization is extensively used in IaaS cloud in order to integrate/decompose physical resources in an ad-hoc manner to meet growing or shrinking resource demand from cloud consumers. The basic strategy of virtualization is to set up independent virtual machines (VM) that are isolated from both the underlying hardware and other VMs. Notice that this strategy is different from the multi-tenancy model, which aims to transform the application software architecture so that multiple instances (from multiple cloud consumers) can run on a single application (i.e. the same logic machine). An example of IaaS is **Amazon's EC2, Amazon's RDS** etc.

Data as a Service (DaaS): The delivery of virtualized storage on demand becomes a separate Cloud service - data storage service. Notice that DaaS could be seen as a special type IaaS. The motivation is that on-premise enterprise database systems are often tied in a prohibitive upfront cost in dedicated server, software license, post-delivery services and in-house IT maintenance. DaaS allows consumers to pay for what they are actually using rather than the site license for the entire database. In addition to traditional storage interfaces such as RDBMS and file systems, some DaaS offerings provide table-style abstractions that are designed to scale out to store and retrieve a huge amount of data within a very compressed timeframe, often too large, too expensive or too slow for most commercial RDBMS to cope with. Examples of this kind of DaaS include **Amazon S3**, Google BigTable, and Apache HBase, etc.

Cloud IaaS service providers

Cloud infrastructure as a service (IaaS) is a type of cloud computing service, it parallels the infrastructure and data center initiatives of IT. Cloud compute IaaS constitutes the largest segment of this market (the broader IaaS market also includes cloud storage and cloud printing).
The market for cloud IaaS is dominated by two leading service providers. Other service providers have responded by launching new offerings, but customers must carefully manage the risks of adopting less-mature offerings.

Magic Quadrant for Cloud Infrastructure as a Service, Worldwide

Source: Gartner (June 2017)

Amazon Web Services Cloud Platform for infrastructure setup

The AWS Cloud provides a broad set of infrastructure services, such as computing power, storage options, networking and databases that are delivered as a utility: on-demand, available in seconds, with pay-as-you-go pricing. From data warehousing to deployment tools, directories to content delivery, over 90 AWS services are available. New services can be provisioned quickly, without upfront capital expense. This allows enterprises, start-ups, small and medium-sized businesses, and customers in the public sector to access the building blocks they need to respond quickly to changing business requirements. This whitepaper provides you with an overview of the benefits of the AWS Cloud and introduces you to the services that make up the platform.

AWS consists of many cloud services that you can use in combinations tailored to your business or organizational needs. This section introduces the major AWS services by category. To access the services, you can use the AWS Management Console, the Command Line Interface, or Software Development Kits (SDKs).

AWS Management Console:Access and manage Amazon Web Services through the AWS Management Console, a simple and intuitive user interface. You can also use the AWS Console Mobile App to quickly view resources on the go.

AWS Command Line Interface: The AWS Command Line Interface (CLI) is a unified tool to manage your AWS services. With just one tool to download and configure, you can control multiple AWS services from the command line and automate them through scripts.

Software Development Kits: AWS Software Development Kits (SDKs) simplify using AWS services in your applications with an Application Program Interface (API) tailored to your programming language or platform.

AWS Compute Services

Amazon Elastic Compute Cloud (**Amazon EC2**) is a web service that provides secure, resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers.
The Amazon EC2 simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances (called Amazon EC2 instances) to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers and system administrators the tools to build failure resilient applications and isolate themselves from common failure scenarios.

AWS Batch enables developers, scientists, and engineers to easily and efficiently run hundreds of thousands of batch computing jobs on AWS. AWS Batch dynamically provisions the optimal quantity and type of compute resources (e.g., CPU or memory-optimized instances) based on the volume and specific resource requirements of the batch jobs submitted. With AWS Batch, there is no need to install and manage batch computing software or server clusters that you use to run your jobs, allowing you to focus on analyzing results and solving problems. AWS Batch plans, schedules, and executes your batch computing workloads across the full range of AWS compute services and features, such as Amazon EC2 and Spot Instances

AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume—there is no charge when your code is not running. With Lambda, you can run code for virtually any type of application or backend service—all with zero administration. Just upload your code, and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services, or you can call it directly from any web or mobile app.

Auto Scaling helps you maintain application availability and allows you to scale your Amazon EC2 capacity up or down automatically according to conditions that you define. You can use Auto Scaling to help ensure that you are running your desired number of Amazon EC2 instances. Auto Scaling can also automatically increase the number of Amazon EC2 instances during demand spikes to maintain performance and decrease capacity during lulls to reduce costs. Auto Scaling is well suited both to applications that have stable demand patterns and applications that experience hourly, daily, or weekly variability in usage.

Elastic Load Balancing (ELB) automatically distributes incoming application traffic across multiple EC2 instances. It enables you to achieve greater levels of fault tolerance in your applications, seamlessly providing the required amount of load balancing capacity needed to distribute application traffic.

Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes offer the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes—all while paying a low price for only what you provision.

Amazon Simple Storage Service (Amazon S3) is object storage with a simple web service interface to store and retrieve any amount of data from anywhere on the web. It is designed to deliver 99.999999999% durability, and scales past trillions of objects worldwide

Amazon Glacier is a secure, durable, and extremely low-cost storage service for data archiving and long term backup. You can reliably store large or small amounts of data for as little as $0.004 per gigabyte per month, a significant savings compared to on-premises solutions. To keep costs low yet suitable for varying retrieval needs, Amazon Glacier provides three options for access to archives, from a few minutes to several hours.

AWS Database Services

Amazon Aurora is a MySQL and PostgreSQL compatible relational database engine that combines the speed and availability of high-end commercial databases with the simplicity and cost-effectiveness of open source databases. Amazon Aurora provides up to five times better performance than MySQL with the security, availability, and reliability of a commercial database at one tenth the cost.

Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, freeing you up to focus on your applications and business. Amazon RDS provides you six familiar database engines to choose from, including Amazon Aurora (p. 16), PostgreSQL, MySQL, MariaDB, Oracle, and Microsoft SQL Server.

Amazon DynamoDB is a fast and flexible NoSQL database service for all applications that need consistent, single-digit millisecond latency at any scale. It is a fully managed database and supports both document and key-value data models. Its flexible data model and reliable performance make it a great fit for mobile, web, gaming, ad-tech, Internet of Things (IoT), and many other applications.

Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory caches, instead of relying entirely on slower disk-based databases.

AWS Database Migration Service helps you migrate databases to AWS easily and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. The AWS Database Migration Service can migrate your data to and from most widely used commercial and open-source databases. The service supports homogenous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora or Microsoft SQL Server to MySQL. It also allows you to stream data to Amazon Redshift from any of the supported sources including Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle, SAP ASE, and SQL Server, enabling consolidation and easy analysis of data in the petabyte-scale data warehouse. AWS Database Migration Service can also be used for continuous data replication with high availability.

Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define. You have

complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. You can use both IPv4 and IPv6 in your VPC for secure and easy access to resources and applications.

Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service. It is designed to give developers and businesses an extremely reliable and cost-effective way to route end users to Internet applications by translating human readable names, such as www.example.com, into the numeric IP addresses, such as 192.0.2.1, that computers use to connect to each other. Amazon Route 53 is fully compliant with IPv6 as well.

Amazon CloudWatch is a monitoring service for AWS Cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources. Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by your applications and services, and any log files your applications generate. You can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.

Athena is easy to use. Simply point to your data in Amazon S3, define the schema, and start querying using standard SQL. Most results are delivered within seconds. With Athena, there's no need for complex extract, transform, and load (ETL) jobs to prepare your data for analysis. This makes it easy for anyone with SQL skills to quickly analyze large-scale datasets.

Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse that makes it simple and cost effective to analyze all your data using your existing business intelligence tools. Start small for $0.25 per hour with no commitments and scale to petabytes for $1,000 per terabyte per year, less than a tenth of the cost of traditional solutions. Customers typically see 3x compression, reducing their costs to $333 per uncompressed terabyte per year.

Amazon Machine Learning (Amazon ML) is a service that makes it easy for developers of all skill levels to use machine learning technology. Amazon Machine Learning provides visualization tools and wizards that guide you through the process of creating machine learning models without having to learn complex ML algorithms and technology. Once your models are ready, Amazon Machine Learning makes it easy to obtain predictions for your application using simple APIs, without having to implement custom prediction generation code or manage any infrastructure.

Amazon Pinpoint makes it easy to run targeted campaigns to drive user engagement in mobile apps. Amazon Pinpoint helps you understand user behavior, define which users to target, determine which messages to send, schedule the best time to deliver the messages, and then track the results of your campaign.

There are many other services in AWS which can be used for hosting and running any kind of application on cloud environment.

Benefits of Cloud Migration

There are many problems that moving to the cloud can solve. Here are some typical scenarios that will benefit from cloud migration.

- Your application is experiencing increased traffic and it's becoming difficult to scale resources on the fly to meet the increasing demand.
- You need to reduce operational costs, while increasing the effectiveness of IT processes.
- Your clients require fast application implementation and deployment, and thus want to focus more on development while reducing infrastructure overhead.
- Your clients want to expand their business geographically, but you suspect that setting up a multi-region infrastructure – with all the associated maintenance, time, human, and error control effort – is going to be a challenge.
- It's becoming more difficult and expensive to keep up with your growing storage needs.
- You'd like to build a widely distributed development team. Cloud computing environments allow remotely located employees to access applications and work via the Internet.
- You need to establish a disaster recovery system but setting it up for an entire data center could double the cost. It would also require a complex disaster recovery plan. Cloud disaster recovery systems can be implemented much more quickly and give you much better control over your resources.
- Tracking and upgrading underlying server software is a time consuming, yet essential process that requires periodic and sometimes immediate upgrades. In some cases, a cloud provider will take care of this automatically. Some cloud computing models similarly handle many administration tasks such as database backup, software upgrades, and periodic maintenance.

CONCLUSIONS

Cloud computing is changing the way IT departments buy IT. Businesses have a range of paths to the cloud, including infrastructure, platforms and applications that are available from cloud providers as online services.

AWS provides building blocks that you can assemble quickly to support virtually any workload. With AWS, you'll find a complete set of highly available services that are designed to work together to build sophisticated scalable applications. You have access to highly durable storage, low-cost compute, high-performance databases, management tools, and more. All this is available without upfront cost, and you pay for only what you use. These services help organizations move faster, lower IT costs, and scale. AWS is trusted by the largest enterprises and the hottest start-ups to power a wide variety of workloads, including web and mobile applications, game development, data processing and warehousing, storage, archive, and many others.

REFERENCES

1. https://aws.amazon.com/about-aws/global-infrastructure
2. https://aws.amazon.com/blogs/aws/
3. https://www.ec2instances.info/
4. https://aws.amazon.com/getting-started/projects/host-static-website/?c_1
5. https://docs.aws.amazon.com/aws-technical-content/latest/aws-overview
6. https://cloudacademy.com/blog/cloud-migration-benefits-risks
7. https://www.cloudcomputing-news.net
8. https://aws.amazon.com/documentation
9. https://aws.amazon.com/whitepapers/