Infosys cobalt

VIEW POINT

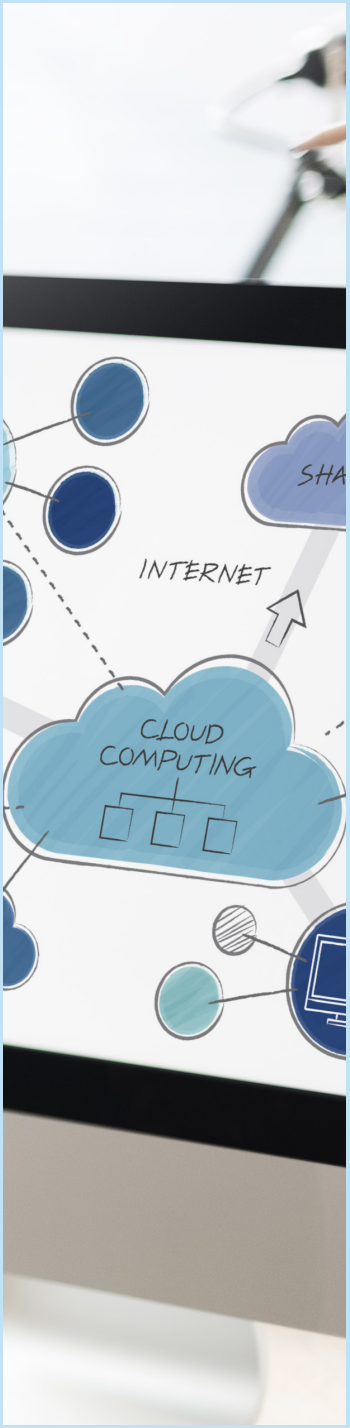

KPI-BASED COMPARATIVE STUDY FOR SELECTING AN ENTERPRISE CLOUD COMPUTING PLATFORM

Infosys®
Navigate your next

Table of Contents

1. Why is cloud important for business

Enterprises across the world are moving their business to the cloud and exploring the wide range of services offered by it. The obvious benefits such as reduced costs, easy access, security and sharing capabilities makes it ideal for businesses with growing data and nonlinear needs.

Cloud computing services fall into three main categories

- Platform as a Service (PaaS)
- Infrastructure as a Service (IaaS) and
- Function as a Service (FaaS)

With multiple players in the cloud technology space and intense competition most organizations find themselves in a dilemma while deciding on the cloud platform to be adopted. In order to determine the best cloud platform and make the right IT investment, an organization needs to analyze the IaaS, PaaS and FaaS capabilities of cloud platforms within their business context, evaluate the core components of each and compare the differences on the basis of metrics. Choosing the best platform is an important decision having a long-term impact, and this needs to be done after carefully considering both the enterprise needs and the capabilities of the cloud platform.

2. Why adopt cloud and how

The increased buzz around cloud can be attributed to its features such as increased flexibility and convenience, ease of management and scalability. Although this is true for all cloud platforms, there are specific benefits of adopting a certain cloud platform over others. This emanates from the fact that the leading players – Amazon, Microsoft and Google – have different approaches to cloud computing. These approaches are driven strongly by their respective background.

This document details the factors to be considered when evaluating and determining the best cloud platform for an enterprise. The platforms reviewed are Microsoft Azure and Amazon Web Services (AWS). The comparative evaluation and assessment are based on the IaaS, PaaS and FaaS capabilities with respect to common business scenarios.

3. How IaaS, PaaS and FaaS capabilities can be leveraged for different cloud migration needs

Once an enterprise has decided to adopt the cloud, it can leverage multiple approaches to do so. The migration strategy would depend on the environment, kind of data to be migrated, interdependencies and the complexity of existing architecture, and software licenses associated with the applications.

Some of the migration strategies that can be harnessed are:

- Lift and shift or re-hosting the entire application on cloud servers using native cloud capabilities
- Re-platforming by making a few cloud optimizations but retaining the core architecture
- Refactoring/ re-architecting by using cloud-native features to improve performance and scale compared to the on-premise environment, migrating to a service oriented or server-less architecture.
- Hybrid allows for partial retaining of existing architecture components with cloud service, which are native as well as other services offered by cloud vendors
- Augmenting the existing architecture cloud setup

Cloud offers capabilities that can cater to each of the above listed migration strategies.

Deep diving into IaaS, PaaS, and FaaS capabilities

Infrastructure as a service (IaaS) is a cloud computing service in which a service provider offers users access to underlying infrastructure such as servers, storage and networking. IaaS platforms eliminate the need to maintain in-house hardware and enable users to configure and manage infrastructure on the cloud. IaaS enables a few specific abstractions such as virtual machine nodes, software-defined networking and attachable storage. By obtaining access to these virtualized components, the enterprise can build its own IT platform.

As with all cloud technology, IaaS solutions are accessible over the internet from a cloud provider's data center. The cloud provider is responsible for the management and maintenance of the on-premises hardware (such as servers, storage and networking) and the virtualization or hypervisor layer. This includes a pool of servers and networks distributed across multiple data centers.

IaaS can offer many benefits particularly to enterprise customers who are seeking to expand or whose scale of operations varies. The key benefits of IaaS are the outsourcing of setup and maintenance operations, and automation of capacity provisioning along with the impression of near infinite scalability. It provides the flexibility to utilize the cloud resources on a need basis instead of purchasing, installing and integrating hardware. Typical features and benefits of IaaS are:

- Scale to expand in capacity with no fear of unused capability
- Time and cost savings
- Access on-demand and pay per use
- Experience independence by accessing the service from anywhere over the internet
- Physical security of data centers
- No single point of failure

Platform as a Service (PaaS) is a category of cloud computing services that allows customers to develop, run, and manage applications and services without having to maintain the associated infrastructure. PaaS services are hosted in the cloud and users can manage them via a web browser. Unlike IaaS, PaaS not only includes

infrastructure like servers, storage, and networking, but also middleware such as development tools, databases and business intelligence services. PaaS eliminates the need to buy and maintain software licenses and the underlying infrastructure and middleware. This makes scaling up or down easier as the user need not be concerned about the infrastructure. This elastic capacity makes PaaS a highly appealing solution for businesses. PaaS offers a few added advantages over IaaS due to its additional features like middleware, business tools, and software. Some of the main benefits of PaaS are:

- Pre-coded components such as workflow, security features etc.
- Pre-developed multi-platform applications – such as mobile
- Reduced cost with the pay-as–you-go model
- Support development teams spread across geographies
- Complete support for application lifecycle, from building, testing, deployment, and management within the same integrated environment.

Function as a Service (FaaS) - also referred to as 'server-less computing' via 'server-less architecture' - is a form of cloud computing services that allows users to write code without having to worry about the platform. It supports event-based architecture, for instance it is based on functions triggered by an event. While PaaS applications typically run all the time on at least one server, a FaaS application will run only when triggered. Developers can use it to deploy an individual function, action, or business logic. It enables the user to write a function without having to worry about server resources, deployment, scalability, etc. FaaS solutions are auto scalable and billing is based on real consumption.

Some of the key benefits of FaaS are:

- Complete abstraction of servers and infrastructure
- Scalability is handled entirely by the cloud provider
- Billing is based on executions and consumption alone, not on the size of infrastructure used
- A lot of innovation is still happening in this area which makes it a capability to watch for

FaaS is relatively new and has fewer support technologies than PaaS. Another limitation is that moving code written on one FaaS solution to another would require rewriting a good part of the code on the new platform.

All these have their own advantages and differences. The below chart gives readers a quick comparison to understand the differences between SaaS, PaaS, and IaaS:

IaaS		PaaS		FaaS	
You Manage	Cloud Manages	You Manage	Cloud Manages	You Manage	Cloud Manages
Data	Servers	Data	Servers	Data	Servers
Operating System	Storage Layers	Applications	Storage Layers	Application	Storage Layers
Applications	Networking	Database Administration, Application deployment infrastructure, Data Pipeline Execution Engines - Workflows, Reporting Engines, Analytics Engines	Networking		Networking
Environments/Runtimes	Security		Security		Security
Middle wares	Virtualization		Virtualization		Virtualization
Database Administration, Application deployment infrastructure, Data Pipeline Execution Engines - Workflows, Reporting Engines, Analytics Engines			Operating System		Operating System
			Environments/ Runtimes		Environments/Runtimes
			Middle wares		Middle wares
					Database Administration, Application deployment infrastructure, Data Pipeline Execution Engines, Reporting Engines, Analytics Engines

The cloud model provides different features and functionalities based on the organization's need and a suitable method can be adopted on the basis of different KPIs.

Let's take a simple example to understand how to select a service. If an organization wants to process some data in near-real time and visualize it, this solution can be built with IaaS, PaaS, FaaS services from the cloud, assuming Apache Spark would be the core for building the solution.

If the organization decides to deploy IaaS, it will also need to spin off Virtual Machine, install and write the Sqoop jobs to ingest the data from the Source System, install Apache Spark for processing, install the Hadoop distributions as the data lake, create the web app to showcase results, install its container etc. In this case, everything has to be managed by the organization. Here the cost aspects would be comparatively lesser but there would be a need for additional effort, a lot of skill sets and management of all these services, servers and technology become mandatory.

If we decide to leverage PaaS, the organization must deploy an HDInsight cluster in Azure. Everything described for IaaS will be available by default as part of the HDInsight cluster. The organization will be able to easily manage the OS, virtual machine etc. But the focus would be narrowed down to the use case and the services available as part of PaaS. Here only software like Spark, Hadoop and others need to be managed by the organization. The cost would be slightly higher when compared to IaaS based setup.

If we implement FaaS, the organization must deploy an Azure Databricks cluster. This will include all capabilities of IaaS and PaaS. Management of OS, VM, or even the management of software services like Spark is taken care of. These services will be part of the cloud feature and extend the Serverless nature. Here the focus would be narrowed down to use case and performance only. The cost would be higher compared to IaaS and PaaS.

4. How to determine the best cloud model?

Having discussed cloud migration strategies and capabilities, it is important to tie business objectives to these capabilities and identify the right cloud provider. Common business use cases for cloud computing could vary from simple data migration/archival to deploying enterprise applications on cloud. The choice of cloud platform to be adopted needs to be based on many factors such as time to provision, virtual machines and network connections, ease of data ingestion, environment security and data encryption, availability and performance of ML algorithms etc.

In order to perform a fair evaluation of both platforms, it is important to evaluate the performance of AWS and Azure on these metrics and evaluate these capability (IaaS, PaaS or FaaS) of cloud that can help enterprises to outperform

4.1 IaaS Evaluation

To evaluate the **IaaS** capabilities powered by AWS and Azure, let us consider a use case where in machine learning running on premise to be ported to the cloud. The cloud implementation would involve provisioning a virtual machine (Azure VM/AWS EC2) with suitable configurations, installing the required ML packages and setting up pipelines/jobs to run the same code. The virtual instance needs to be configured such that the code can be executed on it.

Both Azure Virtual Machine (VM) and AWS Elastic Compute Cloud (EC2) provides a variety of VM types that can be optimized to fit different use cases. Virtual instance or VM is composed of virtual CPU, virtual memory, temporary storage and networking capacity. A user can choose size, power, memory capacity, different availability zones, and regions from which to launch the instance. Both VM and EC2 also offer the option to store data in persistent storage called Blob Storage in Azure or EBS in Amazon.

The main factors to be considered in this comparison are time and cost to set up a virtual environment on cloud as well as the security of data.

Given below are the KPI or facts that can be considered to rate various cloud platforms:

- **Time to set up/deliver or time to provision VM:** A cloud platform can be rated 'easy' to set up if a guided wizard is available and all steps related to infrastructure setup can be done through the UI. If some amount of prior knowledge is essential for configuring the setup, we can call it 'medium' and if manual effort and detailed knowledge is required, the platform can be classified as 'complex' in terms of time to set up.

KPI that can be measured against the following are:

- Identification, exploration and analysis of additional services
- IAM, VPC, Subnet, Route Table with Inbound/outbound Rules
- Internet Gateway, Security Group / ACL/NAT, User Group
- VM Instance/EC2
- Data pipeline, Storage & Subscription Accounts
- **Complexity to set up network connections:** Most platforms provide guided wizards for setting up network connections which are 'easy' to set up. A platform that requires some amount of prior experience to navigate the UI is of 'medium' complexity and one that requires in-depth knowledge would be termed 'complex'.
  - Establish and test connection to VM from on premise and local system, various services like Putty Client required to connect, SSL network extender, and others can be calculated
  - Installation of dependent Python packages/software and managing the dependencies
- **Security:** Based on the features provided by the platform, security can be classified as basic, medium, and leading.
  - KPIs to measure Cloud level security, IAM, VPC, Subnet, Security Group, Gateways, SSH (PPK/PEM) Authentication to access EC2 instance on premise
  - Measure the KPIs for Encryption for data at rest and in transit
- **Cost:** Platforms can be evaluated based on hourly cost of virtual machines
  - Measure the KPI for cost breakdown based on:
    - Data Pipeline, EC2 Instance, S3 etc. or VM, Batch account, ADF, Blob storage

- Performance: Platforms can be evaluated based on compute time
    - Execution of the ML scripts to generate the output and latency
    - Accuracy of the model output

4.2 PaaS Evaluation

To evaluate the PaaS capabilities of AWS and Azure, we can consider a use case where an enterprise application needs to be built and deployed on the cloud. Here data will be fetched from various external services and needs to be transformed to make logical data sets. It also needs to be transformed and persisted/indexed in the cloud to provide an overall search functionality.

The main factors to be considered in this comparison are time and cost to set up a virtual environment on cloud as well as the security of data. Given below are the factors and how we can rate a platform based on these factors. Additional points that can be added with above observation (IaaS) are the following:

- **Time to set up/deliver or time to provision VM:** A cloud platform can be rated 'easy' to set up if a guided wizard is available and all steps related to provisioning services and components can be done through the UI. If some amount of prior knowledge is essential for configuring the setup, we can call it 'medium' and if manual effort and detailed knowledge is required, the platform can be classified as 'complex' in terms of time to set up. Along with above points mentioned for IaaS, the following KPI's also can be also considered
    - In AWS KPI against services like EMR, Lambda functions, Simple Queue Service, S3, Elastic Search Domain, API Gateway, Elastic Bean Stalk, IAM can be considered against the time to setup
    - Similarly, in Azure KPI's against services like HD Insight Kafka & Spark Cluster, Azure Search, Azure Storage Account, Batch Account, VPN, Azure Data Factory, Azure Web App, IAM can considered against time to setup
- **Complexity to operate:** If a guided wizard is provided for ingestion and integration, then the platform is rated 'easy' to set up. A platform that requires some amount of prior experience to navigate the UI is of 'medium' complexity and one that requires in-depth knowledge would be termed 'complex'.
    - Additional KPIs to measure in AWS: Spark code development with EMR, or Lambda function development, Web application deployment, etc.
    - Additional KPIs to measure Azure: Data factory pipeline, development of external data extraction and processing jobs, network connection and web app deployment
- **Performance:** Platforms can be evaluated based on performance time
    - Extraction of external services, parsing, ingesting and indexing the data
    - Overall latency can be evaluated (Possible to display the Indexed data using API Gateway component)

4.3 FaaS Evaluation

To evaluate the FaaS capabilities of cloud, let us consider a use case that involves using machine learning algorithms for predictions on a stream of incoming data. Implementation of the use case requires provisioning servers on cloud, storing the incoming data from on premise to cloud, processing the streamed data and writing it to the data lake. This data is then fed to a prediction algorithm and exposed via a Web service end point.

The main factors to be considered in this implementation are time and cost to setup the environment and performance of machine learning models on both platforms. Given below are the factors and how we can rate a platform based on these factors.

- **Model evaluation parameters** – Rated in absolute value based on Direct comparison of Model performance based on accuracy and error indicators like RMSE, R Square and others.
    - Availability of Machine Learning algorithms: Indicates the availability of algorithms for ML
- **Complexity of development** – If the complexity of developing a data pipeline and pre-processing is low and making changes to already developed code is easy, then it is rated 'easy'. If complexity of developing a data pipeline and pre-processing is medium and making changes is relatively difficult, then 'medium'; and if complexity of developing a data pipeline and pre-processing is high and making changes is very difficult, then 'complex'. Some of KPIS that can be considered are given below
    - Identify the KPI for complexity involved in developing the below components:
        - **AWS**
            - Lambda function development
            - Build of ML
            - Data pipeline orchestration
            - AWS API Gateway development
            - Web application UI development & deployment
        - **Azure**
            - History Load – Conversion of data to Blob
            - Data factory pipeline for incremental load and prediction
            - Development of Web app and rest service for search
            - Web App Deployment
    - Invoking the model or deploying the model with workflow or other applications which indicates how the model endpoint can be invoked from web service apps or standalone script, easily
- **Performance** – Rating is based on direct comparison of data processing time, model creation time, prediction time, and more
    - Model Training Time: Time taken in training a model algorithm
    - Percentage of accurate prediction: Indicates matching actual data value with the predicted data value. Percentage of correctly predicted values would be taken as accuracy parameter of algorithm
    - Train Data Loading, reading data from storage layer to ML layer
    - Time to Data Preprocessing like input of missing values, removing unnecessary values and additional features
    - Scalability and Model Deployment like Deployment and enhancing the scope of the application
    - Customized Data Transformation: Complexity of customizing the data transformation activity. Scaling and creating variables and principle component analysis etc.

About Authors

Vishnu Sankar is a senior architect with experience in building various products and solutions in the field of AI, ML , Real time and big data across multi cloud platforms. He is also an open source contributor from Infosys.

Christeena Sam is a Technology Lead with experience in building various solutions for database management, cluster monitoring and real time log analysis as part of COEs within Infosys.

References

1. https://docs.microsoft.com/en-us/azure/virtual-machines/linux/quick-create-portal
2. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EC2_GetStarted.html
3. https://aws.amazon.com/sagemaker/
4. https://docs.microsoft.com/en-in/azure/hdinsight/hdinsight-overview
5. https://azure.microsoft.com/en-in/pricing/
6. https://aws.amazon.com/calculator/calculator-assumptions/

For more information, contact askus@infosys.com

Infosys®
Navigate your next