

AWS CLOUD COST MANAGEMENT

Best practices to reduce cloud costs and keep them minimal.

There is a shift in the technology industry for companies to move towards more cost-efficient Cloud environments. Companies that move to the Cloud can see immediate benefits to their storage capabilities, computer resources scalability, user experience, security, and costs. If your organization is currently experiencing challenges with their technology stack, it might be time to consider moving to the Cloud.

WHAT IS THE CLOUD?

The Cloud is your organization's technology infrastructure managed over the internet through a qualified Cloud provider. Cloud technology allows organizations to manage their resource needs, increase security, enhance the user experience, and seamlessly integrate custom applications as well as off-the-shelf applications. The Cloud is limitless in resources available allowing users to quickly respond to overnight changes.

CLOUD SOLUTIONS OFFER:

Flexibility: The Cloud allows organizations to adjust resources' capabilities.

Cost-Effective: The Cloud allows organizations to keep start-up costs to a minimum, and only pay for what costs incurred from usage.

Scalability: The Cloud can handle an increase or decrease in workloads easily by configuring hardware requirements or adding additional/removing nodes.

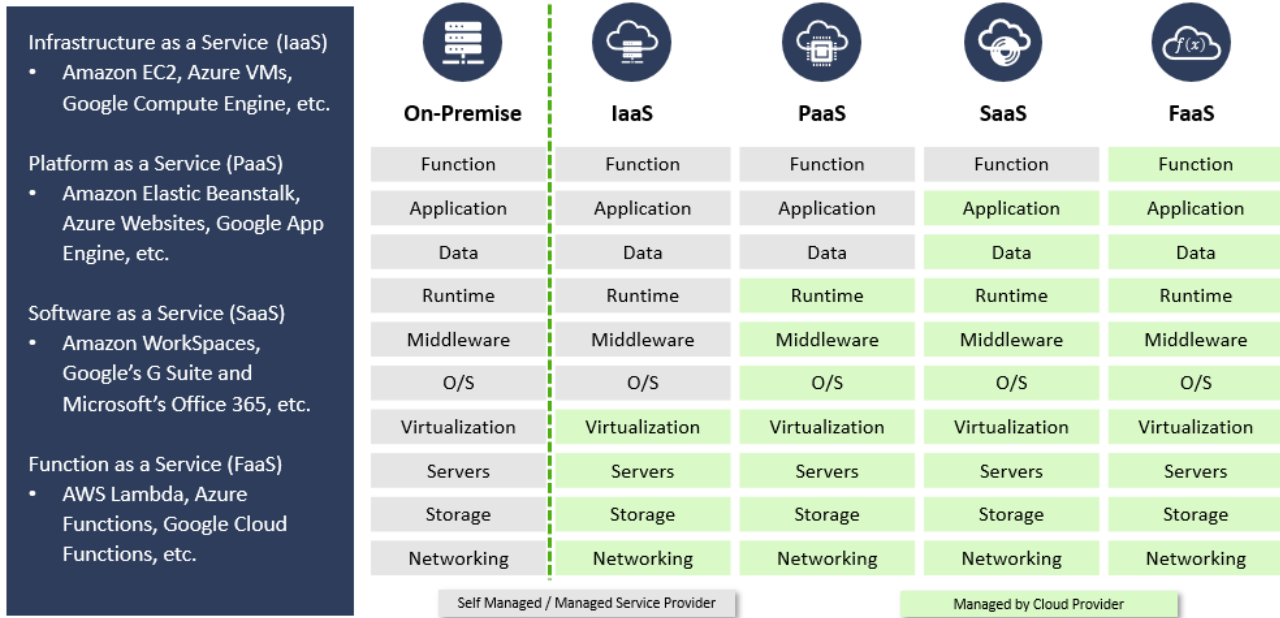
Elasticity: Cloud environments allow organizations to scale resources as needed, as resources demand. Allowing organizations to respond better to peaks and valleys in their business needs.

High-Availability and Fault-Tolerance: The Cloud provides organizations with minimal services interruption, instant backup, and creates data centers to ensure coverage across the organization's network.

Agility: Cloud environments allow engineers to rapidly develop, stage, test, and deploy applications quickly and efficiently.

Many organizations employing an on-premises server will enjoy many of the benefits immediately. Organizations may require or be reluctant to move their entire technology stack directly to the Cloud. The Cloud allows for phased deployment of infrastructure as well as hybrid on-premises and Cloud configurations.

There are four standard types of Cloud environments that can be deployed:



Before an organization should consider moving their infrastructure into the Cloud, they should take an assessment of their organization. Organizations should assess how much they are spending on on-premises servers, including Server, Software, Storage, Network, and Facility costs. Organizations should understand the peaks and valleys of their business capacity requirements. On-premises force organizations to plan for the highest peak times, resulting in resources that are unnecessary and costly.

By shifting to the Cloud, organizations can see immediate savings using agreed shared licensing with the Cloud provider. Savings plans are available through most Cloud providers allowing organizations to manage their costs as their resources increase or decrease. Cloud providers allow for storage of data at extremely reasonable rates, while allowing organizations the control to determine how long data is stored for and how it can be recovered.

Once organizations have decided to move to the Cloud, the first question is always, what is this going to cost me? Moving to the Cloud, costs must be managed properly otherwise they can become a nightmare for accountants and executives who must answer why their costs have increased consistently without any changes to roles or responsibilities. Managing Cloud costs is paramount to successful implementation and transition to a fully functional Cloud environment.

CLOUD COSTS CONSIST OF THREE PRIMARY GROUPS:

Compute: Configure the server resources you require and pay for the usage of the server compute power as you need it.

Storage: Pay for the storage of data in the Cloud. Storage in the Cloud is convenient and secure.

Data Transferred out of the Cloud: Cloud service providers prefer that you keep data housed within the secure Cloud networks created. Moving data out of that secure area will cost a minimal fee.

Managing Cloud costs can be accomplished with some simple standards and quality initial up-front assessments. During an organization's assessment phase, they should determine what type of resources they require. Does the organization manage large data streams, perform big data analytics, manage their own website, and/or create their own applications? This will assist in developing proper resource needs in the Cloud. Organizations should consider how much data they are storing. Are they storing 100 GBs of data, or 100 TBs of data? Good data storage practices will help alleviate storage costs that become increasingly larger. Organizations implementing the Cloud will need to work with users to ensure better standard storage practices. Finally, monitoring and the administration of Cloud costs should be performed daily by a dedicated user.

SOME BEST PRACTICES TO EMPLOY:

Choosing your Cloud Compute Nodes. This is the engine that will run your script, house your website, or manage your data storage. Depending on the needs of the organization, you should consider the following:

General Purpose Nodes: Typically used as web servers or code repositories. They offer a good balance between compute, memory, and networking usage. They can be deployed for most use cases.

Compute Optimized Nodes: Typically used for compute intensive tasks that require high performance processors. Typical users would deploy this type of node for Batch processing workloads, high performance web servers or scientific modeling, and machine learning.

Memory Optimized: Primarily used for processing large amounts of data in memory. Use cases include high performance relational (and non-relational) databases, distributed caches, in memory databases for BI usage, and real time processing of unstructured (raw) data.

Storage Optimized: Typically used in database deployments requiring high, sequential read and write access on large data sets in local storage. Use cases include OLTOP systems, relational and noSQL databases, in-memory databases (such as Redis), Data Warehousing applications, and distributed file systems.

Data Storage is paramount to any organization but managing that data storage can be a challenge. Consider how much data storage you need. How much data is required to be kept? How long should data be stored? How often is the data accessed? How many duplicate copies of the same data set exist for different users? Organizations who implement a few techniques to manage their storage costs will be more successful.

Centralized Data Warehouse: Implement a single point of truth. Reduce multiple copies of the same data set in different database locations. Store reference tables, raw data, and historical vendor data. Allow users appropriate access to central tables.

Data Governance: Organizations should implement data ownership. Data should be owned for ingestion and deployment. Having a single responsible employee will reduce questions when data is deployed to downstream users.

Storage Practices: Many Cloud providers allow for the use of storage tiers, where data can be stored in reduced cost storage areas but incur higher times to retrieve data. Store data that is not needed regularly in lower tier

storage. If data is not required to be kept, delete the data. Maintain a periodic delete schedule to remove data sets no longer in use. Set a standard time frame for data usage of 1-3 years.

Database Tables and Views: When databases are in use, data Views should be deployed to keep a central single table as the point of truth. Copying and creating the same table in different databases is just incurring a cost to store the same table multiple times. Reduce those costs by using views of the master table.

Script Practices: When transitioning to the Cloud users may have some challenges, however implementing good script practices can reduce Cloud costs considerably.

Creating Tables: When creating tables use indexes and keys. The Cloud uses a distributed file system to store data and indexes and keys can be used to reduce compute resources needed, at the same time increasing the query response time. When creating new tables or subsets of data, employ the use of Views.

Implement Views: Users should be selective when using Views and should consider the size of the new table. Users should develop new scripts on small data sets as well as testing and developing on data sets that will not incur a high rate of usage.

Compute Needs: Determine use cases and create data sets for use with scripts before deploying to production. This will reduce the resources needed and address any potential deployment concerns. Before deploying data or applications to a production environment, stage the data in a different environment. Perform load and balance tests. Ensure that resources are deployed and configured correctly for the much larger production deployments.

Temporary Tables: Users should employ the best practice of removing temporary tables and checkpoint tables. Users running large scripts monthly that require check point tables, should be cleaned up after running. Storage of unnecessary data is costly and can become very costly quickly.

Monitoring/Administration: Once users have converted to the Cloud, they will need to manage their costs associated. There are some simple techniques that can be used to ensure costs are managed properly.

Tagging: Cloud providers allow for tagging of resources. Develop a standard set of tags and tag values. Use this standard setup for all Cloud infrastructure in use.

Department – Engineering
Project – Data Pipeline Ingestion
Owner – Name of Manager
Environment – Development / Production
Feature – Name of Feature Cost (S3/Redshift/EC2/etc.)

Cost Tracking Application: Implement a Cost Tracking application to monitor costs. There are many Cloud applications available to monitor Cloud costs. They allow you to connect directly to Cloud cost logs, track your costs through tagging, set up reports/threshold rules, and create/deploy reports through email. Some examples of Cloud applications are CloudAware, CloudCheckr, Armor, Turbot, etc.

Monitoring/Administration: Organizations should dedicate an employee to monitor Cloud costs using the Cloud cost application. They should take ownership of costs, monitoring, and alerting users when unusual costs are

incurred, higher than normal costs are noticed. Monitoring costs on a regular basis will help to keep Cloud costs manageable and keep organizations ahead of forecasted costs.

I hope these simple practices help you with your organization’s challenges when implementing Cloud infrastructure. For follow-up questions, comments, or concerns – [please reach out to Trexin Consulting.](#)

REFERENCES:

- <https://www.udemy.com/course/aws-certified-cloud-practitioner-new/>
- <https://aws.amazon.com/blogs/aws-cloud-financial-management/five-things-you-should-do-to-create-an-accurate-on-premises-vs-cloud-comparison-model/>
- <https://www.columbusglobal.com/en-us/blog/on-premises-vs-cloud-based-solutions>

A QUICK GLANCE AT ON-PREMISES SERVER COSTS VERSUS CLOUD COMPUTING COSTS:

On-Premises

9%
software licenses

- Customization & implementation
- Hardware
- IT personnel
- Maintenance
- Training




Ongoing costs

- Apply filters, patches, upgrade
- Downtime
- Performance tuning
- Rewrite integrations
- Upgrade dependent applications
- Ongoing burden on IT
- Maintain/upgrade hardware
- Maintain/upgrade network
- Maintain/upgrade security
- Maintain/upgrade database

Cloud Computing

68%
subscription fee

- Implementation, Customization & training



Ongoing costs

- Subscription fee

Source: <https://intellias.com/Cloud-computing-vs-on-premises-comparison-guide/>

CLOUD VS. ON-PREMISE AT A GLANCE

	Cloud	On-premise
Time to Implementation	Minutes	Varies - could take multiple days to order, ship, install and configure
Net present value	Not applicable	Capital expense depreciates over time
Tax Implications	Operating expenses - deducted as costs	Capital expenses
Flexibility	Scales with needs	Scaling up requires additional equipment. Cannot reduce costs if scaling down is needed.
Ongoing costs	Pay as you go. If servers are not needed, there is no cost.	Cost of maintenance, upgrades, energy usage and hardware replacement after typical 3 to 5 year lifecycle
Reliability	Google guarantees 99.999% uptime	No guarantee
Personnel Cost	Zero	\$250 per server per month
Security	Google is manned & monitored 24/7 by a dedicated security, privacy, and internal audit teams.	Typically provided by a single or few employees who are not 100% dedicated to security and possibly not trained in security

Source: <https://www.mypickCloud.com/2018/11/13/what-is-the-true-cost-of-being-on-premise/>



This TIP was written by Michael Litwin. Michael welcomes comments and discussion on this topic and can be reached at michael.litwin@trexin.com.