

VMware Cloud Economics

Frequently Asked Questions

Q. What is Cloud Economics?

A. Cloud economics is the study of cloud computing's costs and benefits and the economic principles that underpin them.

Cloud economics involves two primary principles: economies of scale and global reach. Through economies of scale, cloud providers save organizations money because they purchase computing resources in massive quantities at lower costs. When companies utilize these shared resources, they avoid the substantial up-front CAPEX costs of purchasing their own expensive infrastructure. And with a pay-as-you-go pricing model, companies pay only for the resources they actively use, scaling up or down as needed.

The global reach of cloud computing also brings substantial savings. When servers no longer need to be housed on premises—they can be located and accessed from anywhere in the world—companies can dramatically reduce labor costs. Their IT teams no longer need to devote time to deploying and maintaining complex hardware on site.

For more information, see <https://www.vmware.com/topics/glossary/content/cloud-economics.html>

Q. What is the Cloud Economics approach?

A. Our cloud economics approach is to first measure a customer's virtual environment, then project what it would take to move that onto VMware Cloud or Public Cloud, as opposed to leaving it in on-premises datacenter, while including HVAC, electricity usage, labor costs, and more. This can be done on-line at <http://vmctco.vmware.com>, or contact your VMware Representative to get a full Cloud Economics analysis done by one of our Cloud Economists.

Q. Is VMware Cloud more expensive?

A. No, our studies show a 49% savings over on premises traditional environments when moving to the Cloud. Over the last 3 years and across 1000's of studies, the average savings is 49% over traditional infrastructure (Source: VMW Internal Data from TCO Studies; February 9, 2022).

Consider the following:

VMware Cloud on AWS passes through infrastructure savings:

- Scale-out, software-defined, and hyper-converged infrastructure running on x86 servers vs multi-tier, specialized hardware, and disparate infrastructure currently on-premises
- The latest generation of hardware reduces price/performance vs average hardware currently on-premises
- Less standby capacity because it takes minutes to provision a new host, vs. weeks to procure on-premises

VMware Cloud on AWS passes through operations savings:

- Economies of scale – highly skilled VMware employees operate and manage the infrastructure
- Automation – VMware SRE and incident managers work with product developers to implement time-saving automation.

Q. What is VMware Cloud on AWS?

A. VMware Cloud on AWS brings VMware's enterprise class Software-Defined Data Center software to the AWS Cloud and enables customers to run production applications across VMware vSphere®-based private, public and hybrid cloud environments, with optimized access to AWS services. This service is delivered and supported by VMware and its partner community. Find out more online at <http://vmctco.vmware.com>, or contact your VMware Representative to get a full Cloud Economics analysis done by one of our Cloud Economists.

VMware Cloud Economics

Frequently Asked Questions

Q. What technologies are represented in this partnership?

A. VMware Cloud on AWS integrates VMware's flagship compute, storage, and network virtualization products (VMware vSphere®, VMware vSAN™ and VMware NSX®) along with VMware vCenter® management and optimizes it to run on dedicated, elastic, Amazon EC2 bare-metal infrastructure that is fully integrated as part of the AWS Cloud. Customers also get access to 200+ native AWS services over low latency, high bandwidth connection using which customers can enrich the value of their existing applications.

Q. What are the customer use cases for VMware Cloud on AWS?

A. The key use cases include application migration, data center wide evacuation, footprint expansion, virtual desktops infrastructure in the cloud, disaster recovery and application modernization. You can find more detailed information on use cases [here](#).

Q. Have you considered the time, resources, and skills it will take?

A. Most companies underestimate the amount of time and resources it takes to move applications to the public cloud. Even "simple" lift-and-shift requires performance and resiliency testing to ensure the application is meeting user expectations. Most companies also discover they lack the necessary skills, forcing them to rely on expensive consultants and professional services to complete the project.

Q. Have you considered the ROI on re-factoring individual applications?

A. Many applications benefit from re-implementation to native cloud architecture. But equally many don't. Organizations spend money on refactoring just to end up with applications that produce the same results as the originals. It is important to understand the ROI of each refactored application.

Q. Can a native public cloud meet all of your architecture and system requirements?

A. To benefit from native cloud, applications should be re-factored. Organizations must be able to change application source code and configuration. Source code may no longer exist, and configuration changes may break the application. Some applications require features like multicast that is not available in native public clouds (but are supported in VMware Cloud on AWS). It is important to understand all your application requirements.