Cloud Economics Journey

Introductory Guide

This infographic is designed to help new customers realise the full potential of AWS Cloud, by showing you ways to maximise business value throughout your cloud journey.

What if you could:

1 Common reasons for using AWS

Millions of customers are using AWS at scale for reasons such as:

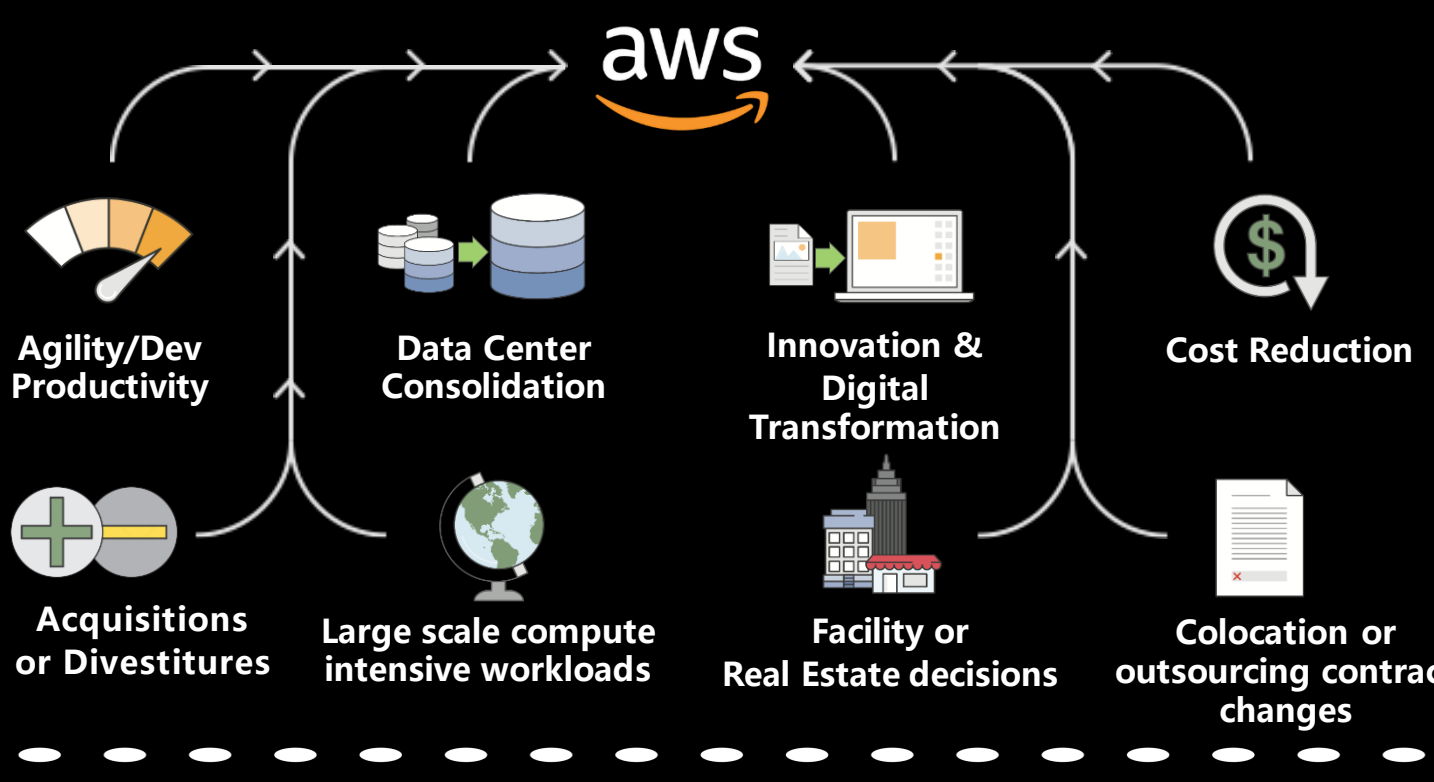

2 The AWS Cloud Value Framework

Customers have realised business value beyond just cost savings

Cost savings (TCO)

What is it?
Infrastructure cost savings/ avoidance from moving to the cloud

Example
30% reduction in TCO (Globe)

Staff productivity

What is it?
Efficiency improvement by function on a task-by-task basis

Example
Over 500 hours per year of server configuration time saved (Sage)

Operational resilience

What is it?
Benefit of improving SLAs and reducing unplanned outage

Example
Critical workloads run in multiple AZs and Regions for robust DR (Expedia)

Business agility

What is it?
Deploying new features/ applications faster and reducing errors

Example
Migrated 30 applications in 50 days (National Australia Bank)

Cost impact → Value impact

3 On-Premises costs to include in any comparison with AWS

The cost of AWS includes cost drivers that are often hidden in an on-premises deployment. To create an like for like comparison, keep in mind the following components of on-premises and colocation environments.

1 Server Costs	Hardware – Server, Rack Chassis PDUs, ToR Switches (+Maintenance)	Software - OS, Virtualization Licenses (+Maintenance)	Facilities Cost: Space, Power, Cooling
2 Storage Costs	Hardware – Storage Disks, SAN/FC Switches	Software - Backup	Facilities Cost: Space, Power, Cooling
3 Network Costs	Network Hardware – LAN Switches, Load Balancer Bandwidth costs	Software – Network Monitoring	Facilities Cost: Space, Power, Cooling
4 IT Labor Costs	Server Admin, Virtualization Admin, Storage Admin, Network Admin, Support Team		

Setting the right expectations for AWS spend

Most organisations operate to a budget and want to understand how much AWS will cost. In order to generate the right estimate, keep in mind the following suggestions.

1 Estimate your cost based on the best practices you plan to apply

Consider the following cost efficiency levers when forming your price estimate.

i) Match supply with demand

With AWS you can align your cost to demand for IT, avoiding waste and meeting peak business demand.

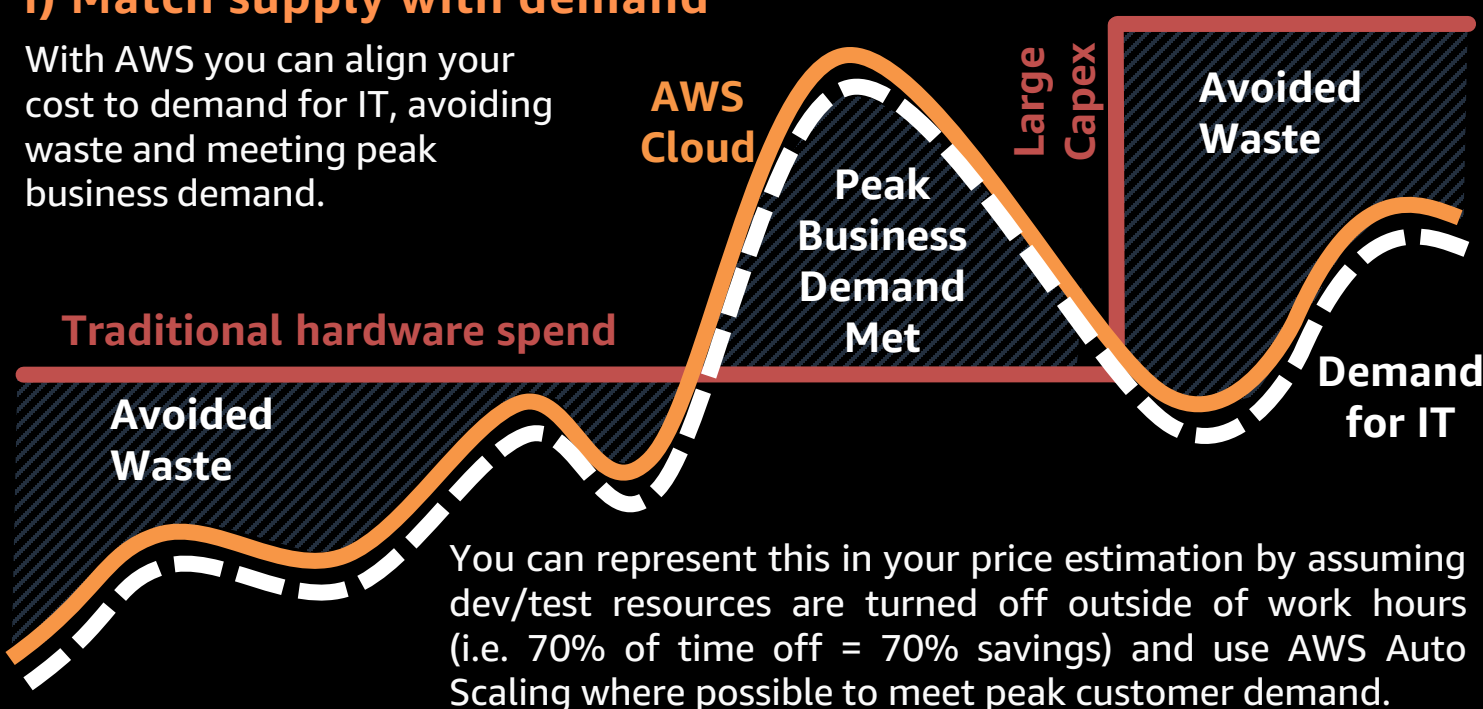

You can represent this in your price estimation by assuming dev/test resources are turned off outside of work hours (i.e. 70% of time off = 70% savings) and use AWS Auto Scaling where possible to meet peak customer demand.

ii) Pick the right pricing model

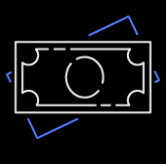

Pick from one of the three pricing models (on-demand, Reserved Instances, Amazon EC2 Spot) in your cost estimation. Learn more on pages 5 and 6 of this document.

Learn more about Pricing Options

Click the icon to learn more

iii) Fit storage to your needs

Storage type can have a big impact on pricing and cost. For example, long-term archival storage (Amazon Glacier) can be 20x cheaper than persistent local storage (Amazon EBS-GP2).

Learn more about Storage Options

Click the icon to learn more

Combining optimisation levers has a significant cost impact

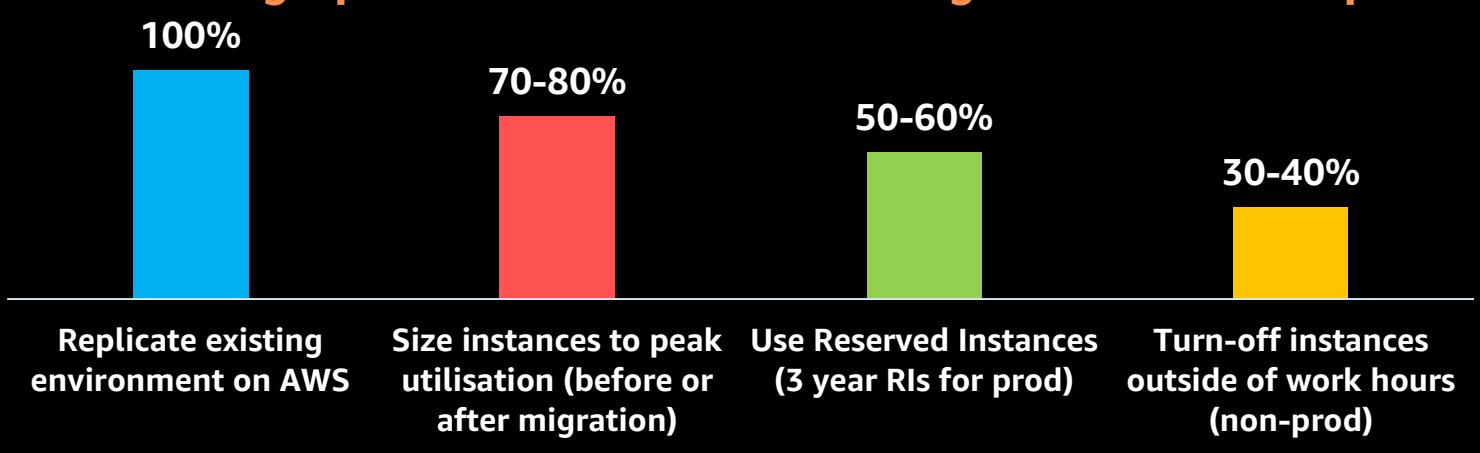

2 Preparing cost estimates

New workloads

Net new workloads should consider the best practices above when using any tool suggested below.

Migrating Workloads

For migrations you should determine your peak resource utilisation (i.e. peak CPU and RAM), and the expected usage pattern (e.g. % of time unused) before creating your price estimate. Tools like your hypervisor resource utilisation report or TSO Logic will be able to provide the data points.

The following tools help you estimate your predicted spend on AWS Cloud:

Self-serve

Simple Monthly Calculator

AWS pricing calculator

Both calculators are useful for pricing, the newer AWS pricing calculator will supersede the Simple Monthly Calculator.

Supported options

Migration Acceleration Program (MAP)

For larger migrations or experiments, you may qualify for support such as detailed cost modelling and/or programs such as the Migration Acceleration Program.

Contact AWS Sales

Reach out to your Account Manager or AWS Sales if you'd like to learn more about the MAP or TSO Logic.

Click the icons to learn more

Setup for Cost Visibility & Optimisation

AWS Cloud provides much greater transparency into your IT infrastructure spending.
You can now see spend down to departments, teams or even an individual level.
AWS also provides powerful tools allowing you to predict, manage and optimise your spend.
The steps below will help you get started.

1 Gain insights into your costs

AWS provides transparency into where your spend is being incurred.

We encourage customers to supplement their monthly spend review process with a more frequent (e.g. weekly) approach using AWS Cost Explorer to catch unexpected spend at the time it occurs. AWS Budgets can send warnings based on user defined thresholds.

Monthly AWS invoice

Move towards tools with greater speed to insight

1. AWS Cost Explorer

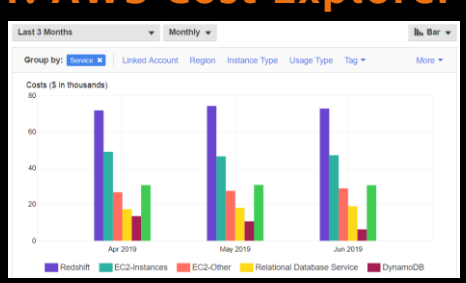

2. AWS budgets

Learn more about AWS cost management tools at: https://aws.amazon.com/aws-cost-management/

2 What can AWS Cost Explorer do?

Cost Explorer is available to all customers, free of charge in the AWS Console.

It visualises your costs ($) and usage (e.g. GBs, Hours), and allows you to drill down via grouping and filtering functionality. Both engineers and budget owners should use AWS Cost Explorer as part of a weekly cost review to avoid any spend surprises.

Questions that Cost Explorer can answer

- What is my spend by AWS Product (incl. AWS Marketplace)?
- How have costs changed over time by AWS Account?
- Which team has cost optimisation opportunity?

Setting up AWS Cost Explorer

Learn more about AWS Cost Explorer	Use Cost Explorer to Analyse spend & usage	Go to the AWS Cost Explorer Console
		

Click the icons to learn more

3 AWS Budgets for warnings

AWS Budgets can provide warnings (via email or SNS notifications) when user specified cost, usage, or reservation thresholds are reached.

AWS Budgets improves awareness of your AWS spend, enabling you to act quickly when actual values deviate from expectations.

Scenarios where AWS Budgets can help

- I want to spend at most $100 in my training account, inform me when it reaches 50%
- Let me know when my account is forecast to reach 110% of my monthly budget
- Warn me when my discounts are not applying as expected (e.g. RI Utilisation drops below 95%)

Setting up AWS Budgets

Learn more about AWS Budgets	How to create an AWS Budget	Go to the AWS Budgets dashboard
		

Click the icons to learn more

Setup for Cost Visibility & Optimisation

Once customers start taking advantage of the tools that provide cost transparency, many also want to know: what is this spend for and who created these resources? The steps below will help you answer these questions.

4 Improve your cost allocation granularity

On AWS you can have multiple accounts grouped under a payer account. We encourage you to think about using a multi-account structure to categorise your spend, improving your ability to know where spend is being incurred.

Beyond accounts, resource tagging enables even more granular insight. Tagging can be used to let you know who created/owns a resource which is useful for informing those who have idle resources.

Benefits

Visibility + granular allocation enables:

- Showback
- Chargeback
- Ownership
- Responsible & efficient behavior across larger orgs.

Multi-account structure (AWS Organizations)

A multi-account structure via AWS Organizations will separate your resources and spend by business unit, team, application, and environment.

When getting started, define and agree a multi-account structure and tagging policy with stakeholders such as Finance, Engineering, and Business teams.

Grouping spend and resources by accounts will save you time in the long-run as there is less need to implement tagging enforcement policies and tools.

The AWS Landing Zone solution helps customers more quickly set up a secure, multi-account AWS environment based on AWS best practices.
An example multi-account structure is shown below.

All accounts connected to: **Payer Account**

Digital Transformation Team's Accounts

- Web App
  - Dev acct.
  - Test acct.
  - Prod acct.
- Sandbox
  - Training acct.

Platforms Team's Accounts

- Shared Services
  - Dev acct.
  - Test acct.
  - Prod acct.

Getting started with multi-account structures

Learn more about AWS Organizations

Launch faster using AWS landing zone

Click the icons to learn more

Tagging resources and tagging enforcement

Resource tags can be used like accounts to allocate spend. Advantages of tags include: 1) supporting optimisation automation (e.g. resource turn-off can be based on tags), 2) providing greater detail than account-level categorisation. Disadvantages of tags are that they require enforcement (e.g. correct for misspelling / missing tags).

Similar to accounts, stakeholders from across the business should be engaged to define valid tag names (e.g. Cost Centre), valid values (e.g. Cost Centre A), and to agree the method of tagging enforcement.

Getting started with resource tagging

Tagging best practices whitepaper

AWS tagging strategies

Cloud Custodian: Enterprise grade tagging enforcement

Click the icons to learn more

Cost Optimisation Levers

The following methods are commonly used by customers to improve cost efficiency on the cloud

1 Select the right instance size for your workloads

On-premises environments are often over-sized because they need to be provisioned for peak expected demand 3-6 months into the future (e.g. sized for end-of-year sales peak in July). AWS allows you can select the cheapest instance for what you actually need and up-size resources when required.

Right Sizing

Pre-migration, use your hypervisor resource utilisation report or a discovery tool like TSO Logic to fit your AWS environment to your actual IT need. Sizing down pre-migration reduces total effort as re-sizing resources in Production is more difficult. The example below shows basic steps to right size resources that are already running.

1. Use a tool to find underutilised resources and total savings potential

2. Agree when to resize, how many to resize, and understand any other constraints

3. Size to what's needed (m4.4xlarge -> m4.xlarge saves 87%)

4. Review application Performance

5. Celebrate the savings win

Getting started with Right Sizing

Blog: EC2 Resource optimization tool

Learn more about Cost Explorer: Rightsizing Recommendations

Click the icons to learn more

2 Use Reserved Instances

Reserved Instances (RIs) are a commitment in exchange for discount. They provide up to 75% discount off vs. on-demand instances. RIs are not physical instances, but a billing discount applied onto On-Demand usage. Each hours, RIs are checked against running on-demand instances, if a match is found, a discount is applied, otherwise the discount is unused for the hour.

Considerations before buying RIs

- If you have spare IT budget, upfront RIs become more attractive as they provide greater savings
- If you need to change instance type, size, or operating system within the next year, consider Convertible RIs
- Consider filtering out non-prod AWS accounts when running an RI recommendation as some non-prod resources can be turned off outside of work hours
- If you use a tool like AWS CloudWatch or AWS Trusted Advisor, consider removing low-utilisation instances from your RI recommendation as you may want to right size the resource first

Services offering RIs

- Amazon EC2 & EC2 Hosts
- Amazon RDS
- Amazon Redshift
- Amazon ElastiCache
- Amazon Elasticsearch
- Amazon DynamoDB*
- Amazon CloudFront*

* Reserved Capacity, but not an RI

RI parameters can be adjusted to cater to your needs

The following RI parameters affect the amount of discount.

	Less discount		Greater discount
Payment option	No-upfront	Partial upfront	All upfront
Duration	1 year		3 year
Operating System	Others (typically)		Open source
Instance type / size	Older generations (typically)		Newer generation
Class	Convertible		Standard

Getting started with Reserved Instances

Learn about EC2 Reserved Instances

How to purchase Reserved Instances

AWS Console RI Tools

- RI Coverage
- RI Recommendations
- RI Utilisation
- RI Budget Warnings

Click the icons to learn more

Cost Optimisation Levers

The following methods are commonly used by customers to improve cost efficiency

3 Schedule on-off your non-production workloads

One simple method to reduce costs is to stop resources that are not in use, and then start those resources again when their capacity is needed.

Benefits of Scheduling

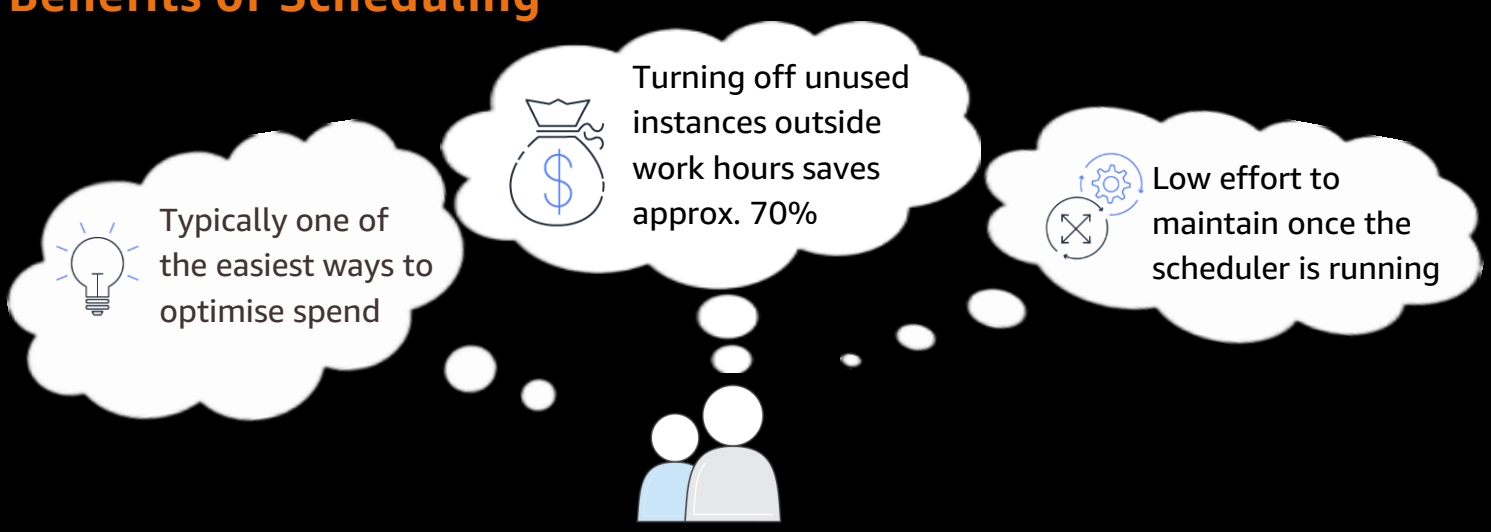

Shut down unused instances with AWS Instance Scheduler

AWS Instance Scheduler enables customers to configure custom start and stop schedules for their EC2 and RDS Instances.

Setting up AWS Instance Scheduler

Instance Scheduler landing page

View the implementation guide

Click the icons to learn more

4 Use Amazon EC2 Spot Instances

Amazon EC2 Spot Instances let you take advantage of unused EC2 capacity in the AWS Cloud. Spot Instances are available at up to a 90% discount compared to On-Demand prices.

How best to use Spot Instances

As Spot is made up of the spare capacity in AWS data centers, you have the option to hibernate, stop or terminate your Spot Instances when EC2 reclaims the capacity back with two-minutes of notice.

With the right architecture, customers can use Spot even in production environments. For example, Spot can be architected in a way (via EC2 Fleet) that can switch between on-demand, RI, and Spot based on the cheapest available instance without interruption to your application.

SPOT IS IDEAL FOR:

- ☑ Fault-tolerant
- ☑ Flexible
- ☑ Loosely coupled
- ☑ Stateless workloads

Customer Spot Case Studies

Find out about how other customers leveraged Spot

Click the icon to learn more

Workloads suitable for Spot

Big Data

CI/CD

Containerised Workloads

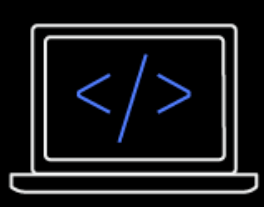

Web App/Services

HPC

Getting Started with Spot Instances

New Spot pricing model

Getting started guide

Introduction to EC2 Fleet

Click the icons to learn more

5 Ongoing training and learning

Recommended resources to continue learning about AWS Cost include:

AWS Cost Labs | AWS Well Architected | AWS Cost Management Blog