

MASTERING THE FUNDAMENTALS OF AWS COST EFFICIENCY

Giving IT, Operations and Finance teams the means to
build a strong cloud cost management practice



CONTENTS

| | |
|--|----|
| INTRODUCTION | 3 |
| STAGE 1 GAINING BASIC COST VISIBILITY Putting an end to surprises in the AWS bill | 4 |
| STAGE 2 MANAGING COST ALLOCATION AND CHARGEBACK Everyone sees what they've spent, and where they spent it..... | 8 |
| STAGE 3 OPTIMIZING USAGE EFFICIENCY Stopping waste, improving savings..... | 13 |
| STAGE 4 DEVELOPING A RESERVED INSTANCE PURCHASING STRATEGY Purchasing AWS Reserved Instances with a purpose | 18 |
| STAGE 5 UNDERSTANDING THE BUSINESS VALUE OF INCREASING CLOUD SPENDING Aligning cloud costs to doing business | 25 |
| IT'S A CYCLE OF CONTINUOUS CLOUD COST MANAGEMENT IMPROVEMENT A never-ending story of savings..... | 30 |

INTRODUCTION

THE COSTS OF RUNNING your AWS infrastructure can add up if left unchecked. While companies can create killer apps to serve their users well, they can also leave a lot of savings on the table without the right tools to optimize their spending.

Having worked with some of the biggest AWS spenders in the world, Cloudability has identified notable trends in the way the most successful companies monitor, manage and mitigate their AWS costs. Following these guidelines will allow you to scale your own cloud infrastructure efficiently.

We collected all the winning trends that we know of and boiled them down into what we call **The Five Stages of AWS Cost Efficiency**.

- STAGE 1** Gaining Basic Cost Visibility
- STAGE 2** Managing Cost Allocation and Chargeback
- STAGE 3** Optimizing Usage Efficiency
- STAGE 4** Developing a Reserved Instance Purchasing Strategy
- STAGE 5** Understanding the Business Value of Increasing Cloud Spending

Use this e-book as a guide to learn and understand each of these stages, and apply the key tips, visualizations and reports to drive better cloud cost management for your own organization.

STAGE 1

GAINING BASIC COST VISIBILITY

Putting an end to surprises in the AWS bill

THE FIRST STEP IN EFFICIENT CLOUD USE for any size of company is to be able to see and understand what it's spending and whether that spending is going up or down. This goes beyond the monthly statement and into a much more granular view.

First of all, anyone trying to wrangle their company's AWS accounts needs to make sure that ALL of the company accounts are known and included in this process. Secret projects kick off, new resources spin up to take care of overflow work—this is just the way business gets done in the world of tech. But getting every known AWS resource and its billing into one place is the first winning step toward efficient cloud cost management.

So to reiterate, even within an individual company there can be multiple master payer AWS accounts. To understand the total cloud costs of your business, these accounts must be identified and their data aggregated into the same cost management tool.

Daily monitoring of month-to-date costs, rolling averages and estimates of end-of-month spending delivers accountability and confidence to operational teams and management. This visibility into cloud cost and usage data are foundational pieces of information that teams can use to optimize their cloud spending.



“Breaking the bill down into a daily view is a great first line of defense against unwanted cloud spending anomalies”

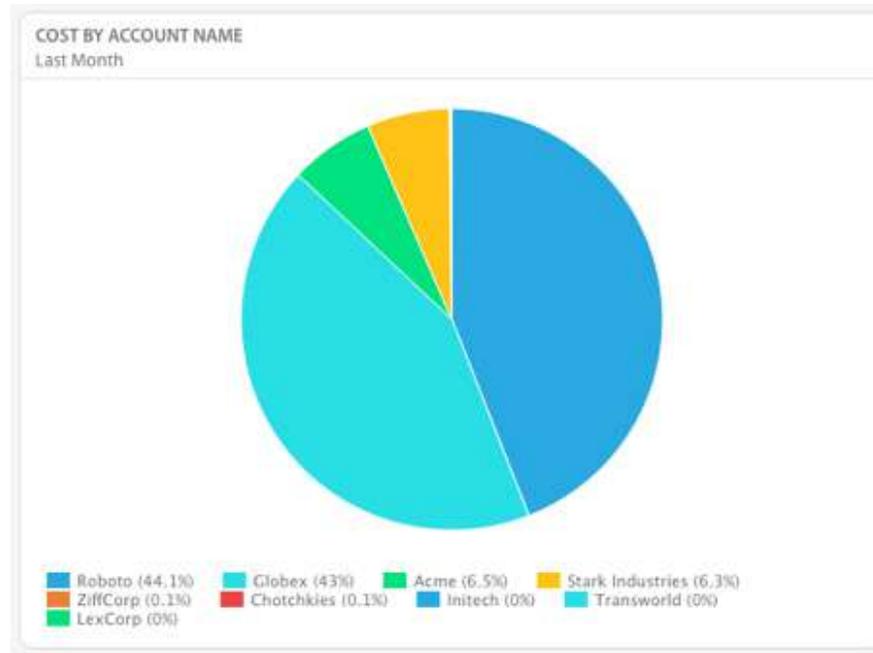
TRANSFORM THE MONTHLY VIEW INTO A DAILY ONE

Even at the monthly view, poring over an AWS bill is an intense process. Breaking down those numbers across all the products and services into a daily view adds a few layers of complexity.

Breaking the bill down into a daily view is a great first line of defense against unwanted cloud spending anomalies. For instance, finance and operational teams can learn of spiked usage within days or weeks instead of finding out a whole month later. This allows teams to adjust their cloud spending and provision operational teams to accommodate for these spikes in usage.

| EXTREME SPREADSHEET | | | | | | | | | | | | | | |
|---------------------|------|-------|------|-------|------|-------|------|-------|------|------|------|------|------|-----|
| 1200 | 1245 | 1236 | 2034 | 3456 | 6.09 | 3.12 | 5678 | Δ | 7809 | 1430 | 5634 | 565 | | |
| Σ | 3789 | 4500 | 2.45 | 1430 | 450 | 542 | 9845 | 2.45 | 565 | 3456 | 1236 | 120 | | |
| 5632 | π | 4583 | 1023 | 6709 | 700 | 127 | 4567 | 120 | 6157 | 12 | 6149 | 324 | | |
| Σ | 3456 | 34.12 | 2.45 | 1430 | 6.09 | 34.12 | 3012 | 2.45 | 565 | 3456 | 5637 | 234 | | |
| 1200 | 1245 | 1236 | 2034 | 3456 | 1200 | 900 | 1278 | Ω | 7809 | 1201 | 9907 | 565 | | |
| Σ | 3456 | 34.12 | 2.45 | 1430 | 6.09 | 890 | 5678 | 2.45 | 565 | 3456 | 1236 | 120 | | |
| | | 110 | 3456 | 34.12 | 2.45 | 1430 | 6.09 | 12.56 | 5678 | 2.45 | 565 | 3456 | 1236 | 120 |
| | | 1200 | 1245 | 1236 | 2034 | 3456 | 8.99 | 5699 | 5678 | Δ | 7809 | 1430 | 5634 | 565 |
| | | Σ | 3456 | 34.12 | 2.45 | 1430 | 5.66 | 34.12 | 5678 | 2.45 | 565 | 3456 | 1236 | 120 |
| | | 213 | 1236 | 6709 | 542 | 900 | 7809 | 1200 | 8900 | 34.1 | 230 | 3126 | 1090 | 550 |
| | | 120 | 1245 | 1236 | 2034 | 3456 | 8.99 | 5699 | 5678 | Δ | 7809 | 1430 | 5634 | 565 |
| | | 110 | 3456 | 34.12 | 2.45 | 1430 | 6.09 | 12.56 | 5678 | 2.45 | 565 | 3456 | 1236 | 120 |

The AWS Detailed Billing Report can be quite long and detailed. Parsing specific billing information is a big task when done by hand.



LEFT Instead of one large bill, breaking usage down by account provides a greater degree of detail and visibility into spending.

BELOW Taking this billing view down even further into environments of accounts adds more granularity.

| | ACCOUNT NAME | ENVIRONMENT | COST (TOTAL) |
|--------------------------|-------------------|-------------|--------------|
| <input type="checkbox"/> | Wayne Enterprises | TEST | \$82,239.65 |
| <input type="checkbox"/> | Wayne Enterprises | development | \$45,144.20 |
| <input type="checkbox"/> | Wayne Enterprises | production | \$42,174.97 |
| <input type="checkbox"/> | Wayne Enterprises | qa | \$38,482.40 |

SEPARATE CLOUD SPENDING BY ACCOUNTS

Along with a granular, daily view of their cloud spending, businesses can organize cost and usage data by accounts or teams. This means organizing and breaking down the monthly AWS bill into specific teams and their projects. This financial and operational accountability gives teams the confidence to know where to scale operations up, where to tone things down and to truly understand where cloud investments are going.

In addition to identifying spikes in cloud cost in usage by the day or hour, businesses now have the means of detecting and classifying anomalies and spikes in usage to specific operational teams. This requires proper tagging of cloud resources to get the most mileage, which we'll cover later in the guide.

Separating cloud spending by accounts also provides a way to deliver consistent reporting and accountability to team leads and relevant stakeholders. As the business scales its operations, and likely increases its cloud spending, no one is left with a surprising bill if proper cloud cost management reporting and visibility are in place.

WHAT'S NEXT?

Transitioning from a default monthly bill into a daily view by accounts is the first step toward building a culture of cloud cost management. Once this foundation is in place, IT and operations teams can start to identify which parts of their AWS spending can improve and grow.

Financial teams can now have visibility into what is typically a complex billing process. Reducing surprises and being able to add context to spikes in the AWS bill is music to the ears of any financial department.

STAGE 2

MANAGING COST ALLOCATION AND CHARGEBACK

Everyone sees what they've spent, and where they spent it.

NOW IT'S TIME TO VISUALIZE and understand where the dollars are going. Proper tagging and classification of AWS resources are not default features. But businesses willing to do the tagging work upfront yield massive visibility and tracking benefits later on.

Getting management and teammates around a common cost allocation strategy, developing a programmatic and repeatable tagging policy and delivering specific reporting views for specific budget holders are key elements toward building a sustainable cost allocation and chargeback plan.

BUILD A COST ALLOCATION STRATEGY AND TAXONOMY THAT SUITS EVERYONE

Before creating and sending reports to budget owners and stakeholders, AWS users should bring operations, IT, and finance together to create a cloud cost taxonomy that everyone agrees with. This taxonomy becomes the common ground and language to deliver cloud cost management reports and insights.

This taxonomy should assist teams in visualizing cloud cost and usage data and answering questions like, "How much are we spending on staging, dev and test environments?"

While finance might not typically be involved in those areas of the business, a common



cost-allocation strategy and plan should include a working knowledge of terms and definitions so all stakeholders can understand why operations spends what it needs to to keep building, developing and testing new apps and products.

This also empowers finance to ask questions about how teams and projects can improve the return on their spending, such as identifying which teams or projects are most efficient, and which may need some help getting there. The ability to break a master AWS account into smaller, understandable chunks enables clear, transparent communication so all teams can work toward improving the return on their cloud investment.

IMPLEMENT A TAGGING POLICY TO ADD SPECIFIC BILLING DIMENSIONS

Linked accounts alone can't deliver the level of granular cloud cost reporting that businesses need to understand how to save and where efficiencies exist. By implementing a tagging policy, AWS billing details gain new depth and meaning.

Tagging doesn't necessarily mean adding complexity. Taking the time to tag an AWS account upfront gives cloud cost management and reporting tools the means to dig deep into the cost and usage data. This allows users to attribute all kinds of cloud product usage, costs and details to the right teams.



Also, understanding these tagging challenges can help teams properly utilize AWS tags and get the most out of their billing data.

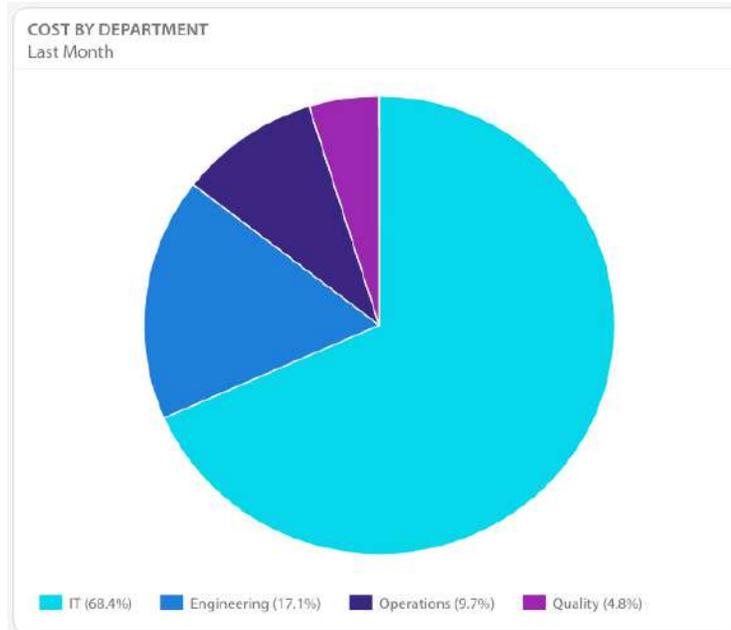
Learn all seven secrets of AWS tags by reading this article!



- » **MAKING THE MOST OF THE 10 TAGS:** AWS allows up to 10 tags per resource. Users who put all of these tags to use gain the most detail and granularity in their reporting!
- » **TAGS ARE CASE SENSITIVE:** This means the tags 'hotdog' and 'Hotdog' will show up differently within the account. Be precise with your tags and diligent about ensuring their consistency.
- » **TAGGING ISN'T RETROACTIVE:** Tags begin to organize and track data on the day that the user applies them, and moving forward until the user removes the tags. This means data from before the user added tags is not collected. This is the same with any programmatic tagging solution. In short, always tag and start as soon as possible!
- » **AWS TAGS REPRESENT A SLICE IN TIME:** Users can build reports on tagged data, but they only represent costs and usage of data within that time frame. Changing tagging values the day after a report will not change the previous tagging data.

Programmatically tagging AWS accounts is the most consistent way to ensure various resources are identified and can yield detailed reporting. Utilizing a programmatic tagging solution, via script or API, reduces the chances for errors and inconsistency.

There are also services (like Puppet, Chef and Cloud Formation), that can help tag and identify untagged resources.



LEFT Finance teams might appreciate a cloud spending breakdown by department to aid with chargebacks.

BELOW Operations management teams might care more about the actual usage and cost of their various AWS products, like EC2 instances.

| | APPLICATION ↕ | INSTANCE TYPE ↕ | USAGE HOURS ↕ | COST (TOTAL BLENDED) ↕ |
|--------------------------|---------------|-----------------|---------------|------------------------|
| <input type="checkbox"/> | frontend | c1.xlarge | 282,762 | \$77,499.03 |
| <input type="checkbox"/> | frontend | c1.medium | 201,978 | \$5,492.61 |
| <input type="checkbox"/> | app | c1.xlarge | 146,700 | \$12,988.95 |
| <input type="checkbox"/> | web-service | c1.xlarge | 131,379 | \$7,880.06 |
| <input type="checkbox"/> | backend | c1.xlarge | 130,063 | \$18,913.07 |
| <input type="checkbox"/> | frontend | m1.small | 111,078 | \$1,319.10 |
| <input type="checkbox"/> | frontend | m1.medium | 103,427 | \$5,206.12 |
| <input type="checkbox"/> | backend | c1.medium | 80,160 | \$1,237.52 |

“Successful companies turn the flood of AWS cloud cost and usage data into simple, actionable charts that can save teams up to 30 to 50%”

WHAT'S
NEXT?

SPECIFIC VIEWS FOR SPECIFIC BUDGET HOLDERS

Once an agreeable cost allocation taxonomy and tagging strategy is in place, AWS users can feel much more confident about their cloud cost and usage reports. This means delivering the right data to the right people so that they can take full control over their cloud spending.

Successful companies can turn the flood of AWS cloud cost and usage data into simple, actionable charts for finance, operations and engineering teams to quickly understand what they're spending on a daily basis, while flagging opportunities to optimize or repair their cloud cost and usage to align with a healthy business bottom line.

Once users achieve a more granular visibility into the monthly AWS bill and can build a cost allocation taxonomy and tagging strategy that all teams agree on, it's time to dig into the data and find ways to save. The next stage encourages the ongoing maintenance and iterative improvement of cloud usage to discover opportunities to streamline systems and save businesses anywhere from 30 to 50% of their cloud costs.

STAGE 3

OPTIMIZING USAGE EFFICIENCY

Stopping waste, improving savings.

JUST AS HOMEOWNERS TRY to turn off water and electricity when not in use to save money, IT and operations managers should use a similar tactic for achieving cloud savings within their existing infrastructure.



Do all of the cloud resources need to be on during nights and weekends? This is an opportunity to save if planned properly.

There are 168 hours in a week, and 108 of them are nights and weekends. Companies often over-provision resources during this downtime, leaving everything on, all of the time. They treat their cloud resources like a datacenter instead of harnessing the pay-as-you-use model of AWS and the cloud.

With hard evidence (uncovered through valiant tagging and cost allocation reporting) that AWS resources aren't being used during specific times of the month, IT and operations managers can find ways to scale and save.

Discovering these opportunities to cut back, optimize and save using insights from AWS cost and usage data is what “Optimizing Usage Efficiency” is all about. The goal is to identify which resources can be turned off, sized down or auto-scaled (using programming) back during non-peak hours. During this phase it’s also important to determine which instances users need to buy reservations (AWS Reserved Instances) for, because this is a major method of saving on AWS costs.

PUT THE TAGGING SCHEMAS TO WORK

Proper AWS account tagging is critical to getting the most out of any cloud cost management tool. Once that hard work is done, usage and billing reports gain a whole new level of granularity and detail. From these details arise ways to fine-tune and augment AWS resources to gain cloud savings over time.

At a bare minimum, the “name” tag, when used properly, can yield a lot of data about why the particular AWS resource exists. Users typically stuff as much detail into these as possible, creating name tags that might say, for example:

```
name="<environment>-<team>-  
<customer>-<project>-<cost-  
center>-<so-on-so-forth>..."
```

As messy as it might seem, these name tags alone can yield a lot of information. It’s not the most sustainable method, however.

Role tags are another basic tag that can tell a lot about an AWS resource. The role tag can identify which layer of the technical stack requires a given resource (e.g. database, platform, web app, etc.). Using a role tag is a simple way to visualize those layers and associate

Pro-tip: Create “tag or terminate” rules. Basically, if an AWS resource has insufficient tagging to identify its role and need, it is terminated.



usage to each to identify where fine-tuning and savings may reside.

Ideally, using as many tags as contextually possible is the best way for AWS users to track, understand, and gain insights as to where efficiency and savings opportunities might exist.

These are just a few ways proper tagging can reveal key cloud usage and cost insights so that IT and operations can confidently make adjustments to yield savings.

LOOK FOR UNDERUTILIZED INSTANCES

With access to the right data, AWS users can gain some quick efficiency wins by identifying

UNDERUTILIZED EC2 INSTANCES: <10% AVG CPU & <1G BANDWIDTH (RUNNING AT LEAST 72 HOURS)
Last 7 days

| INSTANCE ID | INSTANCE SIZE | ACCOUNT NAME | AVG CPU UTILIZATION | DISK I/O | TOTAL BANDWIDTH |
|-------------|---------------|--------------|---------------------|----------|-----------------|
| i-a3fc2fab | c1.medium | Acme | 0.77% | 18,108 | 359.76MB |
| i-56826005 | c1.medium | Acme | 3.90% | 0 | 287.61MB |
| i-10f35330 | m1.small | Acme | 0.87% | 20,012 | 241.87MB |
| i-6b53852a | c1.xlarge | Acme | 0.01% | 2,123 | 63.47MB |
| i-2b2df555 | m1.small | Acme | 0.19% | 4,484 | 63.08MB |
| i-7b3ac269 | m1.small | Acme | 0.14% | 4,229 | 55.53MB |
| i-7528e473 | m1.small | Acme | 0.15% | 12,927 | 49.80MB |
| i-acdbc8ba | m3.medium | Acme | 0.09% | 0 | 30.34MB |

Filtering AWS resources down by low utilization reveals which instances aren't being used much. Look for instances with low CPU and Disk I/O performance — these are key indicators of underutilization.

“When AWS users can pinpoint their ‘off-hours’ with their utilization data, they can confidently spin down their resources during those times.”

resources that are underutilized. Users can seek various combinations of low CPU, bandwidth and disk I/O utilization to identify which resources may be over-provisioned.

AWS users should be wary not to just turn off everything that seems underutilized without concrete reason. This can cause resourcing strains on engineering teams and cause production downtime. As mentioned earlier, a proper tagging strategy and taxonomy can ensure that everyone’s on the same page about what qualifies as an underutilized AWS resource, and what is business-critical in its existing state.

SPIN DOWN UNUSED RESOURCES DURING OFF-HOURS

Of a 24-hour day, if only 10 specific hours of that day are known to have employees working on research and development, why run those resources during the 14 or so “off-hours?” Using AWS cost and usage data to generate a report on when things are actually being used can yield the confidence to turn those resources down during off-hours.

Usually, production resources for a modern app or service are likely to always be running at full capacity. Developmental and research-related resources usually have a bit more leeway in terms of identifying downtime.

By using a filtered report that collects specific hours of the day and the usage of non-production resources, users can pinpoint when development resources hit peak hours and when they don’t see much work at all.

By using a filtered report that collects specific hours of the day and the usage of non-production resources, users can pinpoint when development resources hit peak hours and when they don’t see much work at all.

Once again, the accuracy of these recommendations are only as good as the account tagging efforts (as mentioned earlier in the e-book!) employed by your team.



WHAT'S
NEXT?

The goal of Stage III is to break down all AWS cost and usage data to know if resources are being used to their potential. If the resources are over-provisioned, maybe it's time to talk to engineering to come up with ways to shift resources and their sizes and types around to fit the workload better.

Because why pay to park for an entire day *if you're only going to be there for an hour?*

Once this kind of usage optimization happens, and operations teams are confident that they're running the most efficient scenario possible (because they have the right data to do so), it's time to explore ways to use AWS Reserved Instances to lock in a better rate on finely-tuned cloud environments. This is where some serious savings can occur.

STAGE 4

DEVELOPING A RESERVED INSTANCE PURCHASING STRATEGY

Purchasing AWS Reserved Instances with a purpose.

RESERVED INSTANCES (RI) are an AWS billing construct that allow users to reserve AWS resources and capacity for a lower rate. Each of these reservations are made for a particular class of instance in a particular Availability Zone within a specific region. While users take advantage of significantly lower rates, they must commit to paying for all hours of a given term (either one or three years).

The strategic purchase and application of RIs can yield significant savings for any size of IT and operations team. To get there, teams must make the most of what they can learn from their monthly AWS cost and usage data.

SAVINGS FROM A LOWER RATE ... AT A COST

RIs are basically coupons for AWS resources. These coupons require an upfront fee and a dedication to using the coupon for a one to three-year term. There's an opportunity to save as much as 65% on a given AWS resource, assuming the AWS resources are optimized for the business and the right RIs are purchased — that's quite the discount for users willing to make the commitment.

As users learned in Stage III, cloud operations teams need to figure out the most efficient way to run AWS resources before even considering RIs. If this pre-planning and strategy isn't in place, teams run the risk of not only committing to RIs, but buying the wrong RIs while their unoptimized resources continue to run with the on-demand rate. *Not good!*



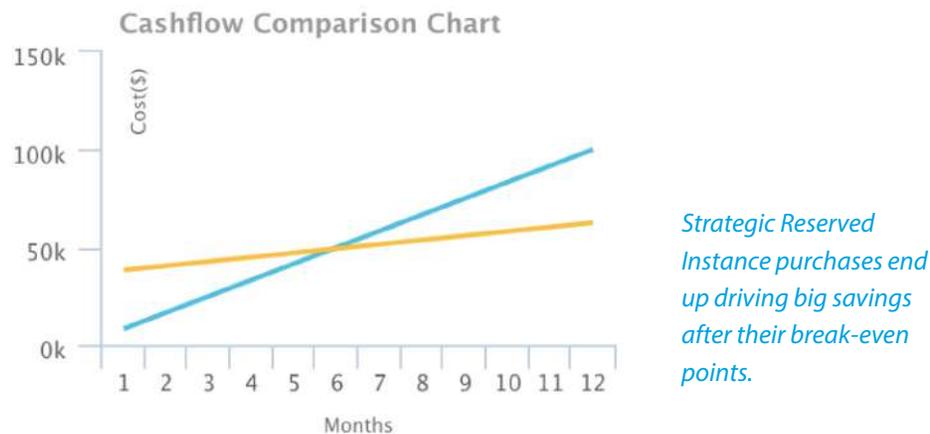
In short, Stage III helps users find their right-sized AWS setup. This lets them make the most of Stage IV, where they can gain the confidence to identify and commit to RIs that lead to the most savings. *Way better.*

But RIs require an upfront cost to purchase, and users are locked into an RI in one or three-year terms. They receive a lower rate on hourly usage for every single hour of the year or three-year term.

Specific planning using detailed AWS cost and usage reports can help users determine which type and term of RI to purchase for the best fit. The ultimate goal is to invest in and apply an RI to reach its break-even point (where savings regains upfront and RI fees) and beyond (all savings, all good).

The break-even point might take a few months of usage, but users can know that beyond that point, IT and operations are coming out on top savings-wise.

RESERVING REGION-SPECIFIC CAPACITY



AWS users with region-specific apps, products and services may want to invest in RIs that give them a lower rate locked into a specific service region.

Instances on AWS are divided across several geographical areas, called regions. The west coast of the U.S., for example, has two regions: US-West-1, in Northern California, and US-West-2, in Oregon.

Each region contains several Availability Zones (AZs), which are distinct locations within that region. AZs within the same region are connected through low-latency links, but each has its own power and cooling systems, and operates independently from other AZs.

Having capacity reserved in a specific Availability Zone can be particularly useful if an infrastructure uses auto-scaling and frequently experiences spikes in usage.

Here are examples of AWS Availability Zones.

| Code | Name |
|----------------|--------------------------------------|
| ap-northeast-1 | Asia Pacific (Tokyo) Region |
| ap-southeast-1 | Asia Pacific (Singapore) Region |
| ap-southeast-2 | Asia Pacific (Sydney) Region |
| eu-west-1 | EU (Ireland) Region |
| sa-east-1 | South America (Sao Paulo) Region |
| us-east-1 | US East (Northern Virginia) Region |
| us-west-1 | US West (Northern California) Region |
| us-west-2 | US West (Oregon) Region |

“Businesses not only reserve capacity for savings, but to protect against potential downtime or outages.”

Here’s an example: Imagine a social networking app with most of its infrastructure in US-East 1A. If the app is successful, there will hopefully be massive spikes in usage. Generally, apps like this would autoscale to suit these new needs. However, this won’t work if there’s no more available capacity in that Availability Zone. Purchasing Reserved Instances in that AZ is one way to protect from that usage spike.

Keep in mind that Amazon does not guarantee the capacity — it does, however, place those reservations first in line.

PROTECTING AGAINST OUTAGES

The third reason for purchasing a Reserved Instance is a newer trend; businesses are beginning to reserve capacity in other regions as a means to protect against potential downtime or outages.

In the case of a natural disaster, IT and operations teams that run safety-critical apps and services may require moving from one AZ to another. A reservation in a backup AZ would save users “a spot in the front of the line,” ensuring service uptime.

The cheapest way to use reservations this way is to have some All Upfront reservations sitting in a particular AZ. This type of RI planning is more costly upfront, but acts as an uptime “insurance policy” depending on the critical requirement of a given app or service (e.g. public safety, medical, other apps that require near-100% uptime).

UNDERSTAND RESERVED INSTANCE PRICING AND TYPES

Unlike on-demand rates, which charge users by varying hours of use, RIs apply a lower rate to every hour of a chosen term (one or three years). AWS users must pay a usage fee for RIs, which is a key figure to take into account when planning a long-term investment.

Examples of RI fees include:

- » **ALL UPFRONT:** Pay for the entire reservation term in one upfront payment. This payment option offers the highest savings rate.
- » **PARTIAL UPFRONT:** Pay for part of the reservation term in an upfront payment, then pay the remainder in monthly installments. This option costs more than All Upfront, but less than No Upfront.
- » **NO UPFRONT:** Pay for the reservation in monthly installments throughout the term's duration. This payment option offers the lowest savings rate.

The more a company can pay upfront, the more long-term savings are available. Every business has a different cost of capital, so there's no definitive answer when it comes to which RI type to invest in.

By using AWS cost and usage data to identify where the best purchases can be made, IT and operations can feel confident in making one to three-year investments in RIs and yield the most savings over time. Using actual cloud cost and usage reporting is a foundational piece of a sustainable RI purchasing strategy.

CREATE A SUSTAINABLE RI PURCHASING STRATEGY

Here are some key components of RI purchasing strategies that Cloudability has seen work well.

- » **BASE NEW PURCHASES ON COST AND USAGE DATA.** Proper account tagging and visualization can generate key reports to base strategic RI purchasing upon.
- » **MODIFY EXISTING RIs.** If users have already purchased RIs, it's important to optimize and modify those before committing to new RI purchases.
- » **FIND THE SMALLEST RI PURCHASE TO DO THE MOST GOOD.** This is similar to a proof of concept of an RI strategy. Users can use cost and usage data and reporting to identify a small set of AWS resources that can benefit from a long-term RI buy.
- » **START WITH A QUICK PURCHASE, ITERATE.** Identifying these small purchases (mentioned above) avoids any large engineering and finance discussions while allowing teams to start trickling savings in. Once the proof of concept becomes a tangible tactic, continue in small iterative batches to begin optimizing cloud spending while visualizing the savings from the RI strategy.
- » **INSTITUTIONALIZE MONTHLY PURCHASES.** Once the small, iterative purchases continue to generate results, they can evolve into bigger, monthly adjustments.
- » **PUT SOMEONE AT THE HELM.** By placing RI planning, purchasing and governance into the hands of a technical and financially savvy person, businesses can now have a go-to specialist. This person ensures that RI buying and modifications adjust to the changing nature of any business, while promoting a healthy culture of cloud cost management.

**WHAT'S
NEXT?**

The first four stages outline how to visualize the monthly bill into a daily view to better understand the nuances of cloud costs, how to tag and allocate costs to properly chargeback cloud spending to different parts of the business, how to identify areas of underutilization through usage data and how to achieve savings by purchasing AWS RIs using a repeatable, iterative strategy.

The fifth fundamental step focuses on communicating cloud cost management and connecting returns on investment with other stakeholders. Now is the time for AWS users to begin discussing how cloud environment growth and its savings impacts business unit costs, and learn the means to inform other teams about the ways cloud cost management is helping the bottom line and isn't just a growing, unmonitored figure.

STAGE 5

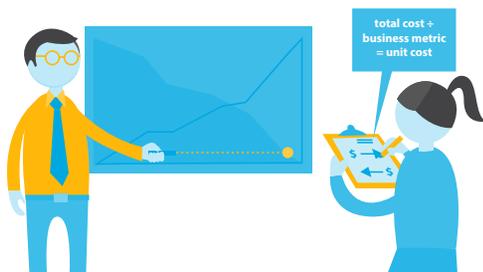
UNDERSTANDING THE BUSINESS VALUE OF INCREASING CLOUD SPENDING

Aligning cloud costs to doing business.

EVEN WITH A CLOUD COST MANAGEMENT PRACTICE in place, a lot of companies' cloud bills still increase. But, that isn't necessarily a bad thing. Assuming other business objectives, like user growth and revenue, are in place, a growing cloud budget is a sign of a healthy business.

Building a repeatable AWS cost management strategy might cut costs down, but without aligning those lower cloud costs to a unit of business, it's hard to show the true impact of cloud savings to the executive team. AWS users who put cloud cost management into play will now need to identify how to add business unit costs into the cloud metric mix.

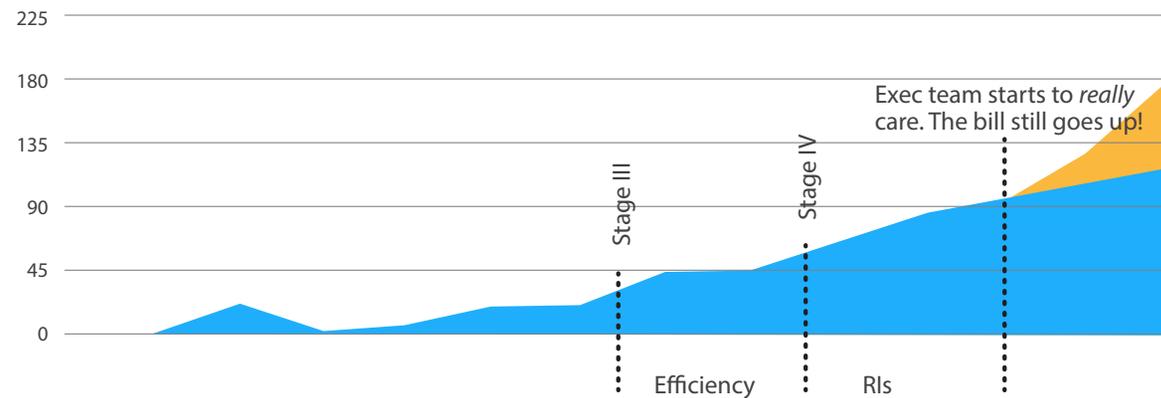
Positioning cloud savings and business unit cost together is the key here. Show cloud cost management benefits as a business asset, and not a growing liability.



KNOW WHEN IT REALLY STARTS TO MATTER

AWS users who find success with cloud cost management will often see their cloud costs continue to inch higher and higher. They might be explaining charts like this one:

An increasing bill isn't necessarily a bad thing

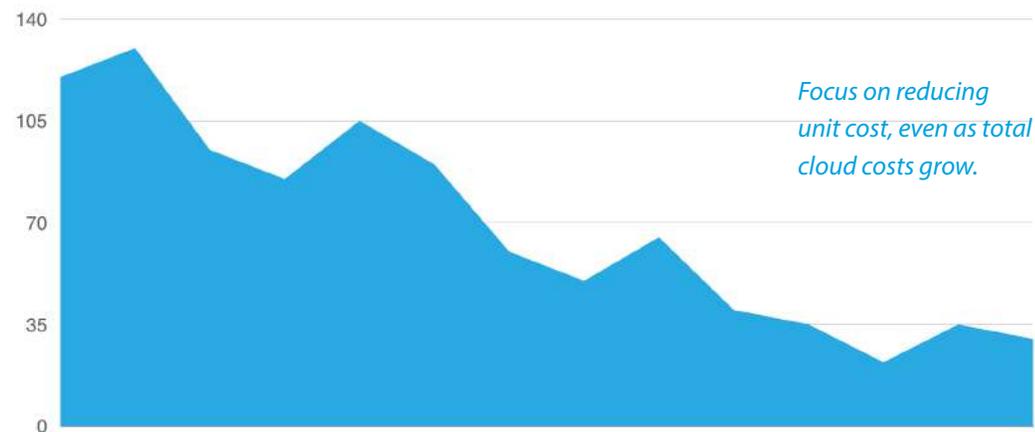


Executive teams tend to start caring when the cloud use goes from experimental to exponential costs. Whatever the size of your business, leadership always starts to care when the costs rise without a good reason why. For some companies, cloud costs climb into the millions.

A natural response to this trend is for executives to begin putting pressure on IT and operations to justify these costs. With a cloud cost management strategy in practice, usage

optimization and waste reduction policies should already be in place to help relieve some of this executive pressure.

What helps align executives with cloud cost management and the general growth of cloud environments is to visualize the cost of operating in the cloud with key business unit costs.



Aligning a stable, lean cloud cost with a lowering unit cost means the business is getting the most bang for its buck. For instance, every dollar spent on the cloud might have yielded 10 units of value. Showing that cloud costs help lower business unit costs (because it's easier to launch the app or send content and data to serve users) also shows a positive return on investment. From this perspective, even though overall cloud spend is increasing, it shows a definite business benefit that can put non-technical executives at ease.

Export cloud cost and usage data into Google Sheets to start the business unit cost conversation. Do this easily with the Cloudability API



IDENTIFY UNIT COST

This unit cost can be difficult to define. What does the business value at the highest level? At the operational level? For instance, one department might consider the unit cost being the cost to generate an app user's subscription. Another department might consider it the cost to deliver data and content to the user. Depending on the business metric, this formula helps IT and operations align costs of the cloud to key business metrics.

```
Total cost / business metric = unit cost
```

Obviously, it's not always that simple. But, if AWS users lay out a proper tagging taxonomy, cost allocation and chargeback foundation, users can split out total cloud costs among various teams or applications. Now the department that favors generating subscriptions can see how cloud spending impacts their work. The same goes for the team delivering data and content from our example above.

Linking accounts within a consolidated billing structure gives business teams the granularity to understand where cloud spending is going and to what return on investment it generates. Adding separate linked accounts for each discrete application and environment adds another layer of depth and understanding to inform various stakeholders within the organization.

Putting environments into separate accounts makes usage optimization easier and allows quick separation of specific business costs like COGS (e.g. production) from R&D (e.g. development) spending.

By following these steps, a cloud cost-related unit cost can now become a key metric within the executive dashboard. This is a much more elegant, informative figure than simply referencing the monthly AWS bill.

MANAGE THOSE UNIT COSTS

Besides generating this useful figure, AWS users should use cloud cost and usage reporting to assist in managing the unit costs of various teams and stakeholders. It's the simplest way to drill down into the value of costs as they're accrued. It's also a strong vehicle in encouraging transparent, honest conversations around cloud usage and spending.

This can be particularly interesting for engineering or operations groups who frequently spin up resources, allowing these teams to see a direct link between their infrastructure choices and the growth and cost margins of the business. Similarly, using unit cost to measure the effectiveness of efficiency efforts and Reserved Instance purchases can be a great way to visualize its impact on cloud cost management.

So as long as revenue continues to scale with the unit being measured, unit cost will mark one of the most important numbers that a business tracks. Aligning cloud environment efforts to lower these business unit costs is something all teams should strive for.



WHAT'S NEXT?

That's it! That's all there is to it. *Just kidding.*

These five fundamentals of AWS cost management outline an ongoing cycle that IT, operations and finance must balance and continue in order to maintain a healthy cloud cost management culture. At the core of this cycle is having the means to collect, control, visualize and act intelligently upon cloud cost and usage insights.

IT'S A CYCLE OF CONTINUOUS CLOUD COST MANAGEMENT IMPROVEMENT

A never-ending story of savings

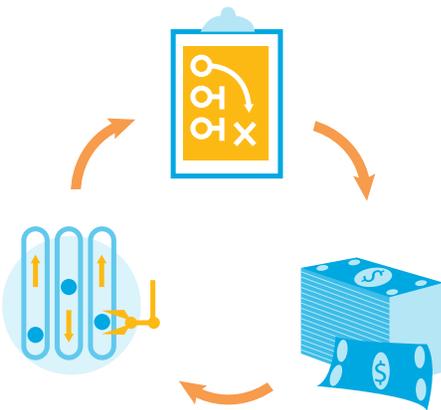
ITERATE AND REPEAT TO SUCCEED

Building a sustainable cloud cost management strategy and process isn't a one-and-done solution. IT, operations and finance teams must continually work together, using the right cloud cost and usage data, to optimize this process over time.

Business infrastructures and policies evolve, and the overall footprint will grow.

Organizations that are nimble, transparent and willing to grow their cloud cost governance practices will ensure continuous efficiency.

At Cloudability, we've seen companies have the most success when they work this model as an iterative cycle into the way they manage their infrastructure. They rinse and repeat, so to speak, working toward growing a scalable cloud environment while putting their costs to good use, while infusing confidence into cloud management and executive teams.



ABOUT CLOUDABILITY

YOUR AWS COSTS are too complex to manage with spreadsheets, and too important to leave to chance. Let Cloudability collect, store and analyze all of your AWS cost and usage data. Then use our powerful Analytics tools and Reserved Instance Planner to create and share reports with anyone in your organization.

Thousands of AWS users across more than 100 countries, including General Electric, Pitney Bowes, Atlassian, OpenDNS/Cisco and Metamarkets, have used Cloudability to manage more than \$4B in cloud spending, and those numbers are going up every day.

Want to learn more? **Check out cloudability.com and start a free, 14-day trial today.**

“Cloudability was the only solution on the market with the technology and expertise needed to keep up with our cloud growth.”

—Mike Driscoll, CEO
Metamarkets

“We have much more granular visibility than anyone has ever been used to.”

—Chris Drumgoole, *COO of IT*

GE