



5 ESSENTIAL STEPS TO RESERVED INSTANCE COST OPTIMIZATION

A Quick Guide to Optimizing Your Reserved Instances in 5 Minutes



5

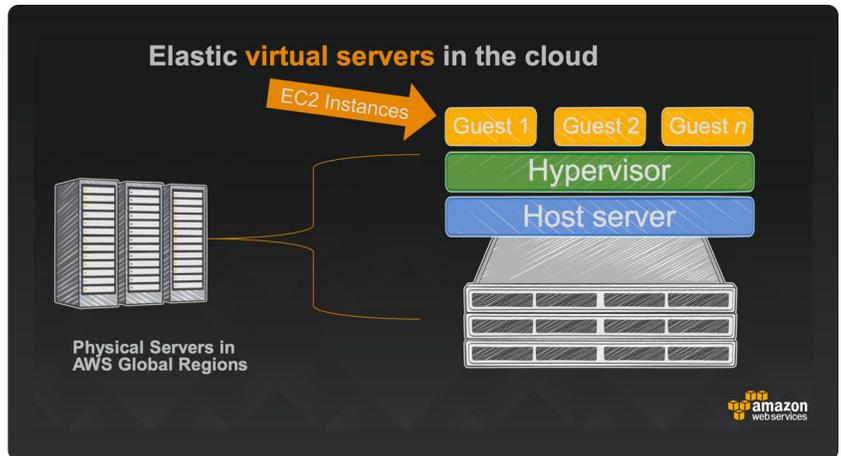
ESSENTIAL STEPS TO RESERVED INSTANCE COST OPTIMIZATION

Introduction

To better understand how to get the best out of your Reserved Instances, it would be useful to have a quick look at what Instances are and what their functions are.

Instances are virtual environments in EC2 (Amazon Elastic Compute Cloud) that help you create a secure, reliable, and scalable computing infrastructure. Virtual servers can be provisioned in different regions and availability zones.

Because it's hard to perfectly determine the scale of your application demands, Instances also come in different configurations depending on your workload.



Types of Instances

To meet varying demand levels, Amazon created a host of instances types you can select from. These instances' configurations differ in pricing, CPU, storage capabilities, network performance, or even the required Availability Zone or Region.

As of today, there are nearly [400 instances across 40+ instance types](#). These instance types are grouped based on their core use-cases.

- ✓ **General Purpose:** Some of the instance types here are Mac and M5 instances like m5.large | m5.xlarge.
- ✓ **Compute Optimized:** Examples are C6g and C5.
- ✓ **Memory Optimized:** Instance types examples under this use-case are X2gd and R5 instances like r5.16xlarge.
- ✓ **Accelerated Computing:** Examples are P4, Inf1, and F1
- ✓ **Storage Optimized:** Examples are: I3, I3n and D3

To find the right instance type, AWS provides the EC2 console or the AWS CLI. For example, you can find different types of instances based on attributes or location with the code snippet below.

```
aws ec2 describe-instance-types --filters "Name=vcpu-info.default-vcpus,Values=48"
aws ec2 describe-instance-type-offerings --location-type "availability-zone" --filters
Name=location,Values=us-east-2a --region us-east-2
```

Finding the Optimal Instance Pricing Option

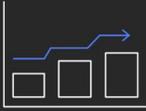
Apart from performance, Cost is another decisive factor application developers use in choosing an instance. Because of this, Amazon has created different pricing options suitable for your organization's needs. There are 7 of these options of which the most commonly used are in the picture below:

In this ebook, we will be focusing on Reserved Instances. They provide significant savings compared to On-Demand instances.

They also provide reliable and stable computing environments. A requirement Spot Instances are not suited for.

More importantly, as we will get to see, when combined with a powerful AWS cost optimization tool, Reserved Instances can provide the most effective computing infrastructure for your needs at a low cost without the attendant long-term commitment you have to cope with.

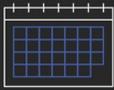
Amazon EC2 purchase options

On-Demand	Reserved Instances	Savings Plan	Spot Instances
Pay for compute capacity by the second with no long-term commitments	Make a 1- or 3-year commitment and receive a significant discount off On-Demand prices	Same great discounts as EC2 RIs with more flexibility	Spare EC2 capacity at savings of up to 90% off On-Demand prices
			
Spiky workloads, to define needs	Committed and steady-state usage	Flexibility to access compute across EC2 and AWS Fargate	Fault-tolerant, flexible, stateless workloads

What are Reserved Instances?

Reserved Instances are cost savings gained from committing to a certain level of EC2 usage. They also come with guaranteed capacity reservations when procured for a specific Availability Zone. As a billing discount on the On-Demand pricing option, organizations can lower their computing spend by up to 72%. There are two-term options offered by Amazon: 1 year or 3-year term with the 3-year term offering a larger discount.

Amazon EC2 Reserved Instances pricing

 Discount up to 75% off the On-Demand price	 Steady state and committed usage	 1- and 3-year terms
 Payment flexibility with 3 upfront payment options (all, partial, none)	 Convertible RIs Change instance family, OS, tenancy, and payment	 Reserve capacity or opt for flexibility across AZs and instance sizes

Types of Reserved Instances

Companies can currently reserve two different types of RIs namely: Standard RIs and Convertible RIs. Scheduled RIs are currently unavailable. The Standard offering class can be modified, bought, and sold. But, it can't be exchanged. On the other hand, the convertible RIs can be modified and exchanged but can't be bought or sold in the RI Marketplace.

Characteristic	Standard	Convertible
Terms (avg. discount off On-Demand)	1yr (40%), 3yr (60%)	1yr (31%), 3yr (54%)
Change Availability Zone, instance size (for Linux OS), networking type	Yes (Using ModifyReservedInstances API and console)	Yes (Using ExchangeReservedInstances API and console)
Change instance families, operating system, tenancy, and payment option		Yes
Benefit from Price Reductions		Yes

Payment Options

AWS also provides three different payment options. Each option provides varying cost savings and flexibility levels. They are All up-front, partial up-front, or no upfront payments options.

Combining the RI types, term options, and payment options give a mix of varying discount proportions. In summary, the more money you commit upfront for a specific Instance type, the larger the savings you get.

What Happens After You Buy Reserved Instances?

After purchasing RIs from AWS or third-party sellers, RIs cost savings are automatically applied when attributes of an instance in use match that of the specified attributes of an RI. The following attributes are to be provided upon launch:

- ✓ Platform
- ✓ Instance type
- ✓ Location i.e Availability Zone
- ✓ Tenancy e.g Dedicated or shared.

You can also select the scope of RIs upon purchase. There's the regional scope that does not reserve capacity but allows size flexibility in applying discounts. For example, if you buy t2.medium RI and you have two running t2.small instances in your account. The discount will be applied automatically as long as they're all in the same region and are of the same instance family type.

There's also the Zonal scope for RIs bought in a particular Availability Zone. You can reserve capacity but don't enjoy size flexibility in discounts application. When you purchase a RI from a specific AZ, any instance in the Availability Zone with matching attributes will have the discount applied on it.

Lastly, you can sell unused RIs at the Reserved Instance Marketplace. AWS charges a service fee of 12% of the upfront RI value though.

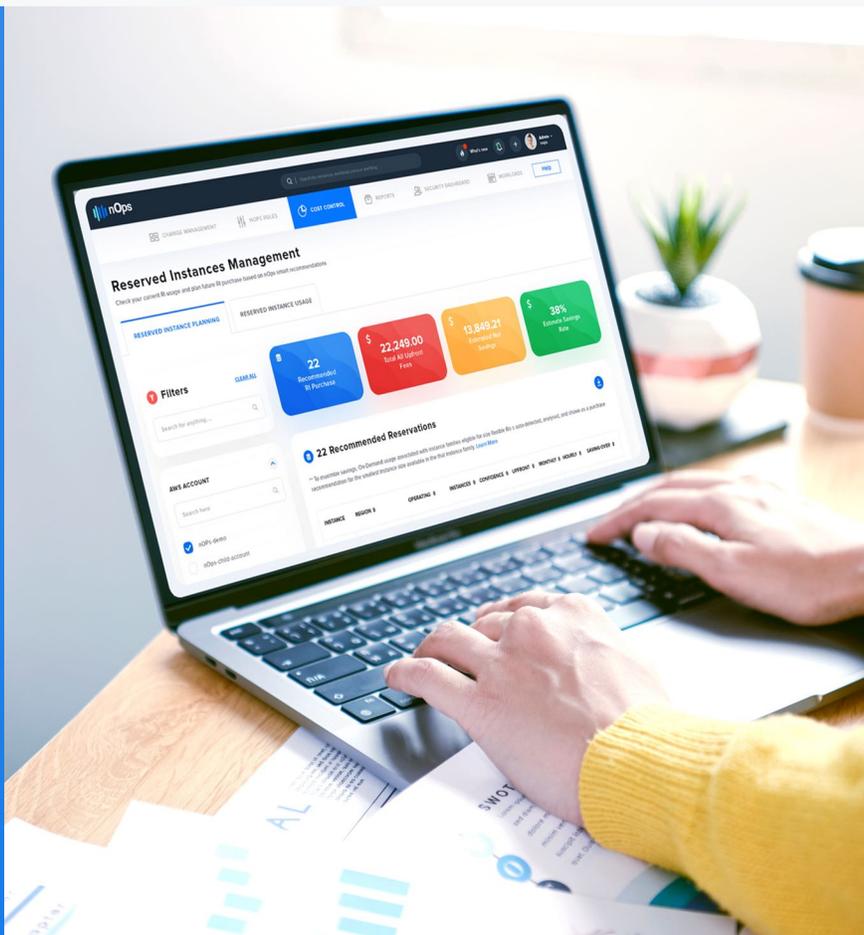
Best Practices For Cost Optimization

Now that we know what Reserved Instances are, what's the best way to optimize our RI cost? In this section, we will highlight some best practices and also mention key steps to Reserved Instance Cost Optimization. We will also show you how to optimize your Reserved Instances in 5 Minutes!

Why is it important?

AWS believes that organizations that focus on cost optimization will be able to achieve their desired goals at the lowest price point. Fully optimized and effectively managed RIs can help you achieve this.

Here are some steps you can put in place to get the best out of your Reserved Instances:



5 Essential Steps to Reserved Instance Cost Optimization

1 Create A Proper Cost Optimization Strategy

The first step to achieving your cost optimization goals is to lay down a proper organizational strategy. It's important to fully understand your application and workload performance and capacity requirements. This will influence your overall compute infrastructure decisions and is critical to your cost management.

Start with setting policies that examine resources usage including how long workloads will be needed. This will probably affect the commitment term options (i.e 1 year or 3 years) you select for example or even the type of EC2 Instances your team uses. Perhaps a combination of Spot Instances and RIs will give the best optimum mix. Which availability zones does your team use most? You also need to set measurable targets you can refine over time. For example, you can set a target that as workload usage increase in the next 12 months by 30%, costs should rise by only 10%. This will likely determine the payment option and RI type you want to choose.

You can start by setting small goals and targets. The important thing is to practice cloud financial management by creating a cost and asset awareness culture. This culture should have processes to analyze your organizational workload components and also identify resource waste. Part of this is to implement reporting cycles that track your cost budget and provide customized reports on changes that can be acted upon quickly.

All this and more will help create a data-driven purchasing strategy.

2 Rightsize Your Existing Infrastructure

AWS considers this step as one of the core pillars of cost optimization. It's also the best way to achieve cost efficiency. This however is a process most organizations ignore before making commitments.

Right-sizing is the process of analyzing your performance and capacity demands while provisioning the exact instances to meet this workload at the lowest possible cost. Some technical components to analyze are your workload's CPU, memory, storage, and networking Instance usage. Analyzing these metrics will help you identify cases of overprovisioning and poor match which you can then correct by downsizing or reconfiguring. It makes no sense to accrue RI commitment discounts that are then wasted on resources you don't need.

Matching demand with the most effective capacity must be an ongoing process as your Instances' needs are constantly changing. Just as you're purchasing RI plans, you need to make sure you select the correct instance type, size to meet the technical requirements of your workload.

A feedback loop that tracks active metrics and makes automated changes is desired. You can use resource monitoring and analysis tools plus cost optimization managers that help you identify opportunities to modify your instances. They also make recommendations that can save you money.

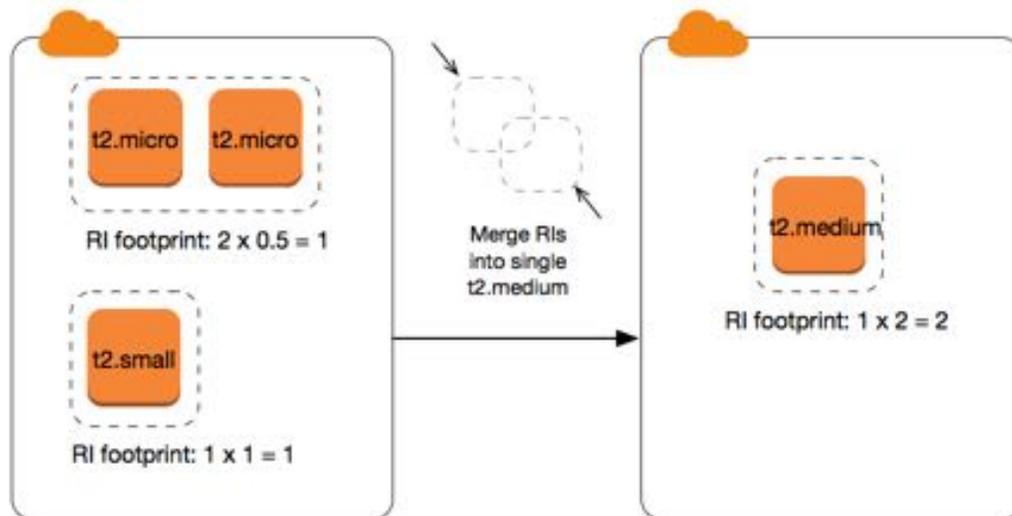
3 Examine and Optimize Reserved Instances

Taking ownership of your resource expense and usage process will help you get the best out of your existing RIs. It will also get the best possible available commitment discounts plans for you. Here are some ways you can optimize your RIs.

- ✓ You can use the AWS Cost Explorer to analyze your usage to determine areas that can be satisfied by RIs. It also identifies EC2 Instances usage that can benefit from Size-Flexible RIs. Size-Flexible RIs were introduced to add additional flexibility in applying your Regional RIs with shared tenancy to all sizes of instances. These sizes are graded on a normalization factor and must be within the same instance family, AWS region, and are for Linux/UNIX RIs. For example, an EC2 RI for a c4.8xlarge instance can be applied to other c4 instances with shared tenancy in the same region as the specified RI. i.e. c4.8xlarge can be substituted 2 c4.4xlarge instances and 16 c4.large instances.

It also means that when you use a larger instance within the same family, you're charged a prorated fee: On-Demand cost + your existing Size-Flexible RI.

- ✓ Optimizing your RIs can also come in the form of modification. You can modify attributes of your Standard or Convertible Reserved Instances such as size, Availability Zone, and scope. To modify an instance size of a Reserved Instance, the new and old instance must have the same size footprint and also be in the same family



You can do this using the AWS Console or using the AWS CLI. Here's a sample AWS CLI Command.



```
modify-reserved-instances --reserved-instances-ids <value> [--client-token <value>]
--target-configurations <value> [--cli-input-json <value>] [--generate-cli-skeleton <value>]
```

Doing all this in different AWS segments can be difficult. It's best to have a Cloud Management tool that will automate all these recommendations.

4 Select the Best Pricing Model

When purchasing AWS RIs Cost Explorer, it's important to choose the best pricing mix. Reserved Instances can be bought with varying levels of payment options and location (i.e region or AZ). Choosing the right mix can give you the highest possible discount which is 72%.

5 Keep Track Of Your RIs Purchases Continuously

This has been partly mentioned in Step 1 and Step 2. Efforts should be made to optimize your RIs over time. Usage can be reviewed periodically before a new purchase is made. Unused resources should also be decommissioned. Keeping track of the latest discount tiers by AWS is also very important.

Optimize Your Reserved Instances in 5 Minutes



Risk-free,
buy-back
guarantee.



Zero
engineering
effort.



Immediate cost
savings.



24/7
continuous cost
optimization.



Shared-savings
pricing model.



5-minute
onboarding
process.

[nOps ShareSave](#) can help you with the 5 aforementioned steps without bothering your developers to continuously monitor hundreds of key metrics.

What is ShareSave? ShareSave service is a real-time, risk-free, automatic life-cycle management of your Amazon EC2 commitments. It adjusts Amazon EC2 commitments in real time based on the environment's capacity requirements and achieves optimal commitment utilization and a dramatic reduction in Amazon EC2 spend.

Simply, only pay for what you use.

With ShareSave, you get automatic execution of the 5 essential steps and you no longer have to worry about:



No Long-term Commitments

Get the flexibility of on-demand instances with the cost savings of 1 year Reserved Instances - Or better! With ShareSave you don't have long-term commitments, no contracts and can cancel at any time. No more concerns about over-provisioning or under-provisioning because you can't predict your Amazon EC2 usage that far in advance. Avoid the financial risk of self-managed RI and SP commitments.



The Complexity of Commitments Lifecycle

Getting the perfect Reserved Instance setup manually can leave you distressed. Apart from the difficult task of accurately predicting the correct resource usage which is made even harder by the 2-3 days delay in getting data from your AWS Cost and Usage Report. There is also the complex choice to be made from hundreds of possible Reserved Instances combinations.

Managing RIs shouldn't be intimidating. ShareSave helps you solve this complexity through 24/7 adjustments of your commitments to maintain a balanced RI life cycle. We achieve this by frequently buying and selling RIs and/or Saving Plans based on your workload's changing capacity requirements.

With ShareSave, optimal commitment consumption becomes a reality.



Your Engineering Team's Declining Productivity

Following all the recommended steps we highlighted is a very time-intensive process. Forcing developers to perform time-consuming activities like forecasting your application's capacity and performance needs, tracking idle resources, and optimizing instances can adversely affect your team's productivity. This is why you need to automate your cost optimization processes.

You can eliminate the time and resources expended on monitoring unused RI commitments with ShareSave. Our AI engine does all the hard work for you by constantly scanning your Amazon CloudWatch and CloudTrail logs and also analyzing your computing resources' data points. We use insights from this analysis to find opportunities to optimize your RI plans.

Free up your developers to create the innovative product features you hired them for. Start using Sharesave today.



Rising Cloud Usage Costs

Self-managing your RI commitments exposes you to the risk of constantly worrying about over-provisioning or under-provisioning your computing infrastructure. Besides that, you're always concerned about your rising usage costs and how to effectively scale your resources without incurring excessive costs.

It's time to gain control of your cloud costs. On average, app owners who use ShareSave enjoy a 45% discount off the AWS On-Demand price. This is possible because we get the most profitable options in the RI marketplace and also take advantage of 3-year no-upfront Savings Plans.

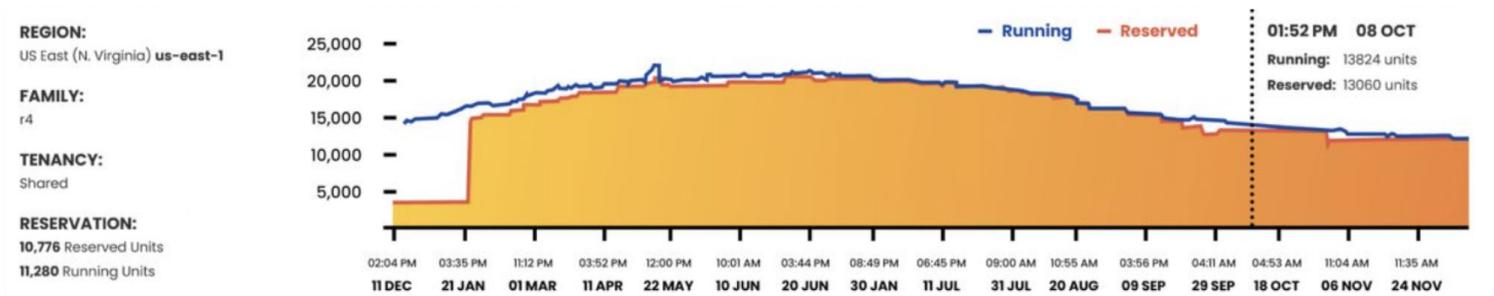
Additionally, we get even more cost savings for you as you scale. How? Our shared-savings pricing model ensures you get a larger percentage of our savings as your usage volume increases.

With our refined process, our customers are able to meet their cost optimization goals by provisioning the best cost-efficient EC2 resources to meet their dynamic workloads.



How Does it Work?

- 1 The ShareSave AI engine collects Amazon CloudWatch and AWS CloudTrail logs and continuously monitors and analyzes infrastructure usage data points.
- 2 ShareSave automatically reacts in real time by purchasing RIs upon an increase in compute usage and selling RIs upon a decrease in compute usage. nOps continuously purchases and sells commitments on an hourly basis, depending on your infrastructure's capacity changes.
- 3 ShareSave grabs the most lucrative discounts in the Amazon EC2 Reserved Instance Marketplace and utilizes 3-year no-upfront commitments when purchasing SPs.



Best of all, ShareSave comes with a **risk-free buy-back guarantee**. In the case of over-provisioning, you would get a full refund for unused commitment dollars. Essentially nOps takes all the risk. No commitments. No long term contracts. Cancel any time.

Onboard in just 5 Minutes

You can enjoy these risk-free cost savings, maximize efficiency and free up your engineers to focus on innovation in 5 minutes. To get started, just follow the steps below:

- 1 Provide nOps only access to your Amazon EC2 usage patterns
- 2 Based on your Amazon EC2 usage patterns, nOps begins to buy RI and/or SP Commitments
- 3 Whether your usage increases or decreases, you pay only for what you use

[Sign up here for a free 30-day trial of nOps](#)

[Visit nOps to learn more about ShareSave](#)