

# AWS Reserved Instances: Everything You Need to Know

In this guide, we'll tackle the core concepts of AWS Reserved Instances: what they are, how they're billed, where they can be applied, and when to buy them. Feel free to dive right in or jump to a topic that interests you.

[Need to start with EC2 basics? Check out our EC2 Guide.](#)

## Reserved Instances: The Basics

### What is a Reserved Instance?

The term *Reserved Instance* (RI) causes a lot of confusion. This is because a Reserved Instance is not really an instance, but a term agreement for reserved capacity and resource allocation to be used by an EC2 instance fitting specified criteria. By signing a 1 or 3 year commitment, you get a reduced rate on that amount (and type) of usage. Think of it like a coupon or voucher that you've paid for—you still have to apply the reservation correctly to gain its value.

### Why Should I Have Reserved Instances?

There are three main reasons to add Reserved Instances to your portfolio.

#### 1. Savings

You can save up to 65% in comparison to running On-Demand instances. The amount of savings you can achieve depends on the [RI attributes](#) and term of your reservation contract.

#### 2. Avoid Delayed Provisioning

Zonal Reserved Instances come with capacity guarantee. This is especially useful for autoscaling applications that may experience a spike in usage during peak hours when you contend with other AWS users in the same Availability Zone. If you have not reserved capacity in your Availability Zone before the spike occurs, there may be none available exactly at the time that you ask for it resulting in a delay of provisioning by a few minutes.

### 3. Disaster Recovery Strategy

Buying zonal reservations in other Regions can be strategic for securing capacity in the event of a natural disaster or general outage. Let's say the bulk of your infrastructure resides in us-east-1. Purchasing a few reservations in us-west-1 could be smart in the event of a hurricane; if an outage occurs, the demand (and competition) for resources in other regions will spike, putting your workloads at risk of having insufficient resources. Having at least a few zonal reservations in Availability Zones in remote Regions could be part of your disaster recovery strategy.

### What are Reserved Instance Attributes?

There are four attributes that affect AWS Reserved Instance pricing. AWS uses these attributes to match your RI purchase to instance usage:

- **Instance Type:** A specific ratio of CPU/memory/storage/networking capacity optimized for a particular workload type
- **Platform:** Linux/UNIX, SUSE Linux, Red Hat Enterprise Linux, Microsoft Windows Server, and Microsoft SQL Server
- **Scope:** Whether the Reserved Instance applies to a Region or specific Availability Zone.
- **Tenancy:** Default (shared) or dedicated

Let's say you bought a 1 year, no-upfront reservation for an EC2 instance with the following attributes:

- m5.xlarge
- us-west-2a
- Default (shared)
- Linux/Unix

If you were to then go into the EC2 console and launch an m5.xlarge with those same attributes, you'd pay the reserved price (\$0.168/hour) instead of paying the On-Demand hourly rate for that instance (\$0.225/hour). That comes out to a \$500 in savings for 1 year of usage. You'd save \$25,000 per year if you were to run 50 on-demand m5.xlarge, as long as the instances run all the time. Below, we will cover the special reservation considerations that are required when you scale your instance usage up and down during the course of a typical day.

### RI Contracts Options, Types, Costs, & Purchasing Limits

The amount you can save depends on each RI attribute and the contract's details. Longer, all-upfront contracts have the deepest discounts.

### Contract Options:

- **All Upfront Payment:** Pay the whole contract upfront
- **No Upfront Payment:** Make monthly payments
- **Partial Upfront Payment:**
  - 1-year: ~50% now, ~4% monthly
  - 3-year: ~50% now, ~1.5% monthly

### Contract Types:

#### Standard:

- You can modify Availability Zone, scope, platform, and instance size (within the same instance type)
- You can sell these on the AWS Reserved Instance Marketplace

#### Convertible:

- You can exchange one or more Convertible Reserved Instance contracts for another Convertible Reserved Instance with new attributes.
- You cannot sell these on the AWS Reserved Instance Marketplace

Jump to [this topic](#) to read about the pros and cons of Convertible RIs.

Below are *approximate* savings for each reservation plan. Specific savings vary by type and Availability Zone (as of Jan 2019):

- **Convertible 1-year:** ~30%
- **Standard 1-year:** ~40%
- **Convertible 3-year:** ~50%
- **Standard 3-year:** ~60%

### Purchasing Limits for Reserved Instances

There are [a few default limits](#) to keep in mind when planning your reservation purchases:

- **Regional Reservations:** 20 RIs/month
- **Zonal (AZ) Reservations:** 20 RIs/month per zone

These limits are independent of each other. So, for example:.

- In a region with three AZs, you could buy 80 RIs/month:  $(1 \text{ Region} * 20) + (3 \text{ AZs} * 20)$
- In a region with six AZs, you could buy 140 RIs/month

You can exceed these default limits by contacting AWS and requesting these limits be raised for your account. Later in this article, we'll go over the [best purchasing frequency](#) for reservations.

## How are Reserved Instances Applied?

In [a blog post about common RI mistakes](#), our chief data scientist Andrew Paine writes, "Reservations must be thought of in terms of a purchase of concurrent instance counts, not monthly instance hours."

The reason for that distinction comes from a common misunderstanding:

1. AWS bills EC2 usage by the second (**true**)
2. Reservations include a month of usage (**true**)
3. I can share that usage with other instances (**true**)
4. Therefore, I can run 4 m5.xlarge instances 50% of the day and would only need to pay for 2 reservations (**false**)

As long as 2 concurrent m5.xlarge instances are running for the entire month, your 2 reservations will be utilized in full and no on-demand costs will be incurred. If you run 4 concurrent instances during for 12 hours each day (and then turn them off at night), you are wasting 50% of each reservation at night (since you are paying for them whether you use them or not). Furthermore, because each RI provides a "pool" of 3600 seconds every hour that can be split between multiple EC2 instances, during the day, you will be exhausting the pool of seconds after 30 minutes for each RI, and thus will be paying the on-demand price for the last 30 minutes of every hour.

The cycle of your workload needs dictate the type and number of reservation that, you should purchase. In certain scenarios such as concurrent usage concentrated during peak hours of the day, buying reservations would cost you more than running On-Demand instances.

## Regional vs Availability Zonal Reserved Instances

Reserved instances are applied differently depending on whether they are fixed to a Region or Availability Zone.

### Regional Application

- **AZ flexibility:** Regional reservations can be applied across their Availability Zones.
- **Size flexibility:** If reservation is on the Linux/Unix platform, the reservation can be applied to any size of that family instance type (e.g., m5.large - m5.24xlarge). Size flexibility does not apply to Reserved Instances that are purchased for a specific Availability Zone, bare metal instances, Reserved Instances with dedicated tenancy, and

Reserved Instances for Windows, Windows with SQL Standard, Windows with SQL Server Enterprise, Windows with SQL Server Web, RHEL, and SLES.

- **Can be applied to linked accounts:** When accounts are linked to a master billing account and RI discount sharing is enabled, Regional reservations can be borrowed by member accounts when not in use by the purchased account.

## Zonal Application

- **No size flexibility:** Qualifying instances must be an exact match for the tenancy, platform, Availability Zone, instance type, and instance size attributes.
- **Can be applied by linked accounts:** When accounts are linked to a master billing account and RI discount sharing is enabled, Zonal reservations can be borrowed by member accounts when not in use by the purchased account.

It is likely that your accounts will have both Regional and Zonal reservations, so it's important to understand how coverage is prioritized.

Zonal RI coverage application hierarchy:

1. Zonal RI's purchased account
2. Zonal RI's linked account

Regional RI coverage application hierarchy:

1. Regional RI's purchased account
2. Regional RI's linked account

The above shows that reservations can be shared with linked accounts running EC2s that have matching attributes. What does that look like?

Let's say you purchase 2 reservations for m5.xlarge use in the Us-west-2a Availability Zone in account-1. You then link account-1 with account-2 and run 3 m5.xlarge EC2s in account-2. None are running in account-1. Both of your reservations would then be applied to account-2, and the third m5.xlarge would be billed at full price.

## Sharing Reservations

You must use [Consolidated Billing](#) in order to benefit from Reserved Instance Sharing.

Consolidated billing aggregates the costs for a group of AWS accounts in a single AWS account. This account is called the "payer account," while the other accounts are known as "linked accounts." Once you have created a consolidated billing family, Reserved Instance Sharing is turned on by default for the payer account and linked accounts. This can be turned off per-account within the consolidated billing family to ensure reservations apply only to their purchased account if desired..

## Reserved Instance Sharing Highlights:

- Reserved Instances offering capacity reservation cannot be shared with linked accounts; they must be used in the account in which they were purchased.
- Discounts for Reserved Instances purchased on one account are applied to the combined usage for that instance type on the payer account's bill.

## Applying reservation discounts to any size:

Regional Reserved Instances on the [Linux/Unix platform with default tenancy](#) also provide instance size flexibility. This means their discount can be applied to instance usage within that instance type, regardless of size. The discount applied is calculated using a Normalization Factor value.

Instance Size	Normalization Factor
nano	0.25
micro	0.5
small	1
medium	2
large	4
xlarge	8
2xlarge	16
4xlarge	32
8xlarge	64
9xlarge	72
10xlarge	80
12xlarge	96
16xlarge	128
18xlarge	144
24xlarge	192
32xlarge	256

Using the table above, if you have a reservation for an i3.medium instance on the Linux/Unix platform with default tenancy (Normalization Factor of 2) and you run an i3.large instance on the

Linux/Unix platform with default tenancy (4), the running instance receives a 50% discount. The same applies for larger reservations; one large instance reservation can cover 2 medium or 4 small instances of the same family.

Sharing reservations between accounts is a great way to ensure you maximize your reservation investments. Metricly recommends creating a Consolidated Billing Family, linking your accounts, and leaving Reserved Instance Sharing on. Deeper analysis can be performed once your accounts are linked to a payer account. For example, you can [compare spending periods](#) and [monitor dimensional budget spend](#).

## Tips for Planning Reserved Instance Purchasing

To round out your introduction to Reserved Instances, we're going to cover a few simple strategies you can use for your purchase planning. These tips all assume you have completed some [preliminary resource tagging](#).

### How to Pick Areas for Reservations

1. **Organize your spend by instance type.** This helps you identify instance types that have the highest number of running on-demands and where the easier reservation decisions lie..
2. **Focus on architecture least likely to change.** Confer with your engineering management to avoid investing time and money buying reservations a part of your environment that's about to change soon. Instead, pick an area where the architecture is expected to be stable for at least a year. Tagging is the best way to organize this effort.
3. **Review your existing capacity utilization.** Use AWS CloudWatch and optionally an analytics solutions like [Metricly](#) to review your existing capacity utilization. Look for application workloads that exhibit consistent capacity utilization. You will also want to make sure the resource you want to purchase reservations for [is properly sized](#). Otherwise, you will commit to one year for capacity that you are not using.
4. **Separate resources by OS type.** Standard reservations have different stipulations and prices based on the type of OS used. Separating your portfolio this way helps you stay cognizant of any unique key details that should inform your buying decisions. For example, Linux/Unix based Standard reservations are more flexible in that they can apply to any size of the instance family purchased.
5. **Upgrade to the latest generations.** The price on the older generation increase as a way to incentivize users to move to the latest generation. Sometimes this requires you to upgrade your operating system. It's not required but it's better to upgrade before you

make a one year reservation.

These five steps produce a short-list of candidates ready for reservation. Take that list and note where those workloads currently sit both regionally and in what Availability Zone. Then determine whether it makes more sense to purchase the more flexible Regional Reserved Instances, the more dedicated Zonal Reserved Instances, or a combination of both.

## Establish an RI Purchasing Cadence

Knowing when to purchase Reserved Instances is equally important to knowing what resources you'll need. It's likely that your company only purchases new RIs on a yearly or semi-yearly basis. But how often do your infrastructural needs remain unchanged over a whole year--or even six months?

Consider adopting a **monthly or at least a quarterly diversified purchasing strategy**. This benefits you in two ways:

- It limits paying On-Demand prices for added resources
- It prevents creating *expiring reservation cliffs*

## Limit Paying On-Demand Prices for New Resources

Making RI purchases on a yearly basis is not frequent enough to achieve the full benefit of reservations. Infact, even quarterly RI purchasing leaves you open to paying On-Demand prices from additional deployed instances 9 months out of the year.

Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sept	Oct	Nov	Dec
Reserved	On-Demand	On-Demand	Reserved			Reserved			Reserved		
Reserved			Reserved			Reserved			Reserved		
Reserved		On-Demand				Reserved			Reserved		
Reserved						Reserved			Reserved		
Reserved				On-Demand	On-Demand				Reserved		
Reserved				On-Demand	On-Demand				Reserved		
								On-Demand	Reserved		
								On-Demand			
								On-Demand			
								On-Demand			
								On-Demand			
										On-Demand	On-Demand
											On-Demand

## Prevent Expiring Reservation Cliffs

When you purchase reservations only once or twice per year, you create a stressful renewal situation that can get costly. This is because planning your *entire* reservation portfolio for a whole year requires research, talking to stakeholders, and committing to bulk purchasing. If you aren't



ready to act by the time your reservations end, you'll wind up paying On-Demand prices for existing resource usage until new reservations are in place.

## Reserved Instances: Basics Recap

Buying reservations is a smart thing to do. You can use reservations to:

- Save up to 65% of your current spend on capacity and resource utilization
- Prevent Regional outages from impacting service delivery
- Ensure proper reservation coverage as you auto-scale during peak hours

## Intermediate

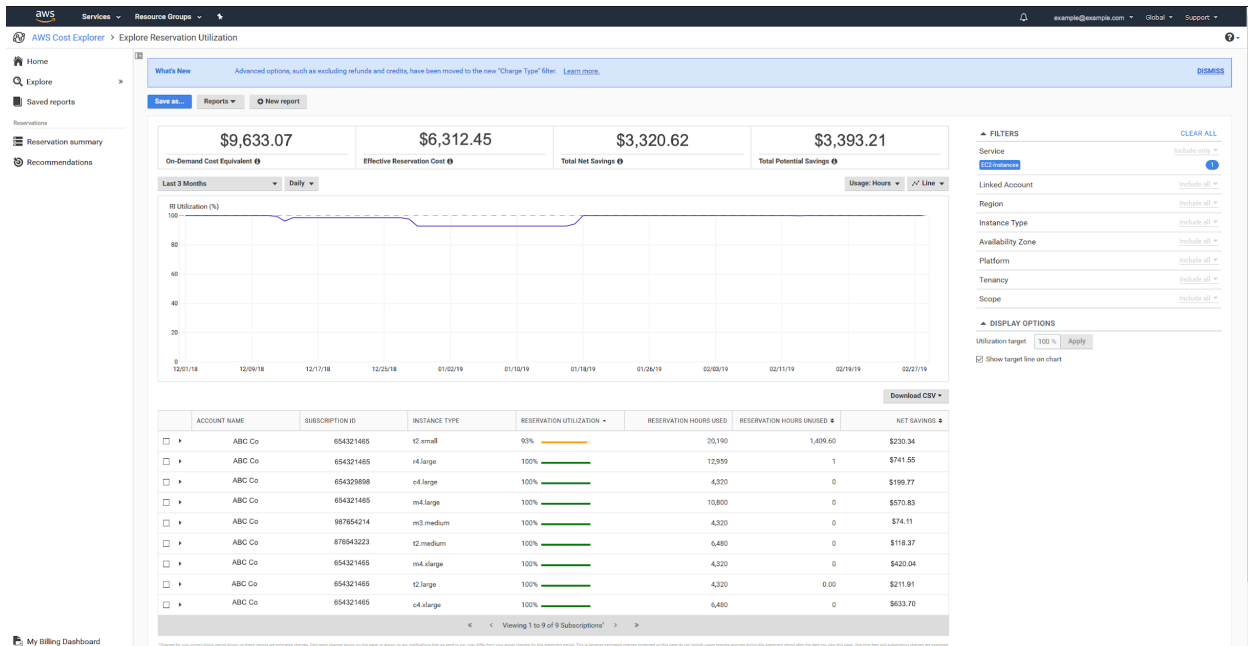
In this section of our guide, we'll tackle some intermediate topics on AWS EC2 Reservations. This section assumes that you have an existing portfolio of reservations and that you have the needed permissions to access your organization's [AWS Billing Console](#).

### 3 Ways to Track Reserved Instance Usage

You can save a lot of money when reservations are applied correctly. Tracking and maximizing Reserved Instance utilization ensures that you get the most out of your commitments. Here are 3 ways to track your Reserved Instance usage:

#### 1. Use Reserved Instance reports in Cost Explorer

AWS Cost Explorer has a few of out-of-the-box tools to get you started with understanding your reservations. [The RI Utilization report](#) provides a visualization that measures the consumption percentage of your purchased RI hours. The image below depicts an account that is close to 100% utilization, however there is a dip between 12/09/2018 and 01/18/2019.



You can drill down into this report using filters for Service, Linked Account, Region, Instance Type, Availability Zone, Platform, Tenancy, and Scope. It's worth noting that this report is updated daily, not hourly. It is not finalized until the end of the month and can take up to 72 hours to become available.

## 2. Create alerts using Reserved Instance Budgets in the AWS Console

Both the severity and duration of a RI utilization dip should be considered when investigating your portfolio's total utilization. Instead of checking the RI Utilization report on a daily basis, [you can set up budget alerts in Cost Explorer](#) and get notifications by email.

The screenshot shows the 'Create a budget' wizard in the AWS Cost Explorer. It is at Step 2, 'Select budget type'. The wizard has four steps in total: Step 1 (Select budget type), Step 2 (Set your budget), Step 3 (Configure alerts), and Step 4 (Confirm budget). Under 'Select which type of budget you would like to create.', there are three radio button options: 'Cost budget' (Monitor your costs against a specified dollar amount and receive alerts when your user-defined thresholds are met.), 'Usage budget' (Monitor your usage of one or more specified usage types or usage type groups and receive alerts when your user-defined thresholds are met.), and 'Reservation budget' (Track the RI Utilization or RI Coverage associated with your reservations. These budgets support Amazon EC2, RDS, Redshift, ElastiCache and Elasticsearch reservation models.). The 'Reservation budget' option is selected. At the bottom right, there are 'Cancel' and 'Set your budget' buttons.

These reservation budget alerts can be scoped to target a specific account's RI usage or usage across all linked accounts. At minimum, the scope has to have a service selected (*Amazon Elastic Compute Cloud - Compute*, for example).

You can also select a data aggregation method to base your budget alerts on by choosing the period of time (daily, monthly, quarterly, annually) considered for your utilization threshold.

The screenshot shows the AWS Budgets 'Create budget' interface. On the left is a navigation sidebar with links to Dashboard, Bills, Cost Explorer, Budgets (highlighted), Reports, Cost Allocation Tags, Payment Methods, Payment History, Consolidated Billing, Preferences, Credits, Tax Settings, and DevPay. The main content area is titled 'Create budget' and includes a sub-header '1 Budget details'. Below this, the 'Budget Type' is set to 'RI Utilization' (selected with a radio button). The 'Name\*' field contains 'EC2 Windows'. The 'Period' dropdown is set to 'Monthly'. Under the '2 Refine your budget' section, there are checkboxes for various filters: Service, Instance Type, Platform, Tenancy, Scope, Region, Availability Zone, and Linked Account, all of which are currently unchecked. A right-hand panel titled 'Create your RI Utilization Budget' provides a three-step guide: 1. Name your budget and select the budget period. 2. Refine your budget criteria by selecting optional filters. 3. Define a utilization threshold and receive alerts via email and/or Amazon Simple Notification Service (SNS).

Multiple notifications with different utilization thresholds can be set to one reservation budget. Notifications that have utilization percentages beneath the set thresholds send an email daily until utilization rises above the threshold.

### 3. Analyze instance-level reservation usage

Gaining perspective on instance-level reservation usage is an important—but challenging—part of understanding and optimizing your AWS bill. This is because per-instance usage data is stored as [detailed billing data](#) (soon to be [AWS Cost & Usage](#)) not currently available in Cost Explorer's reports.

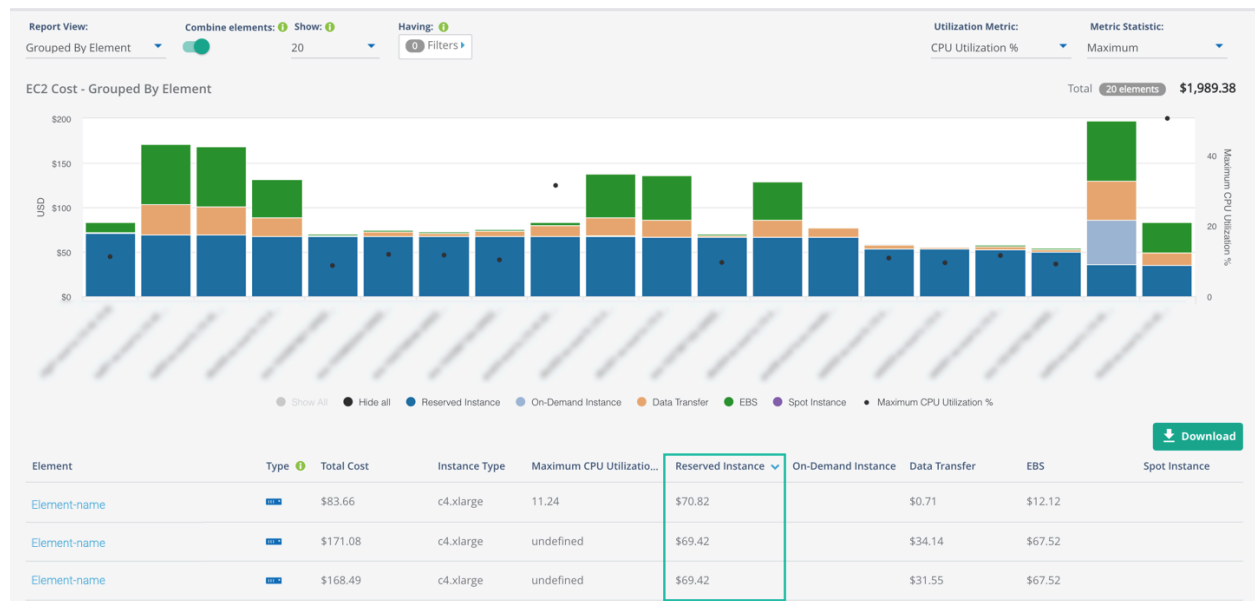
A per-instance dimensional breakdown of RI usage exposes which resources are consuming reservation hours, how many, and calculates a dollar-value for that consumption. This data is listed alongside key attributes such as the instance's region, availability zone, type, family, and other costs. On-demand costs, EBS costs, and Data Transfer costs are also included in detailed billing data, and all three are important to consider when investigating RI consumption.

This data is available via files delivered to an S3 bucket and can be manipulated manually in tools like AWS Quicksight or using your own custom tools. Having [cost allocation tags](#) in place for added context is critical to this stage. [state something here about formulas for amortization rates you'll need and how complicated it is]

After you've prepped your data, it's time to investigate. For this topic's example, we are going to use Metricly's EC2 Cost tool to track reservation utilization at a granular level. Metricly simplifies usage analysis by integrating with your AWS billing accounts and ingesting your detailed billing files—removing your need to download, edit, and upload data into another system.

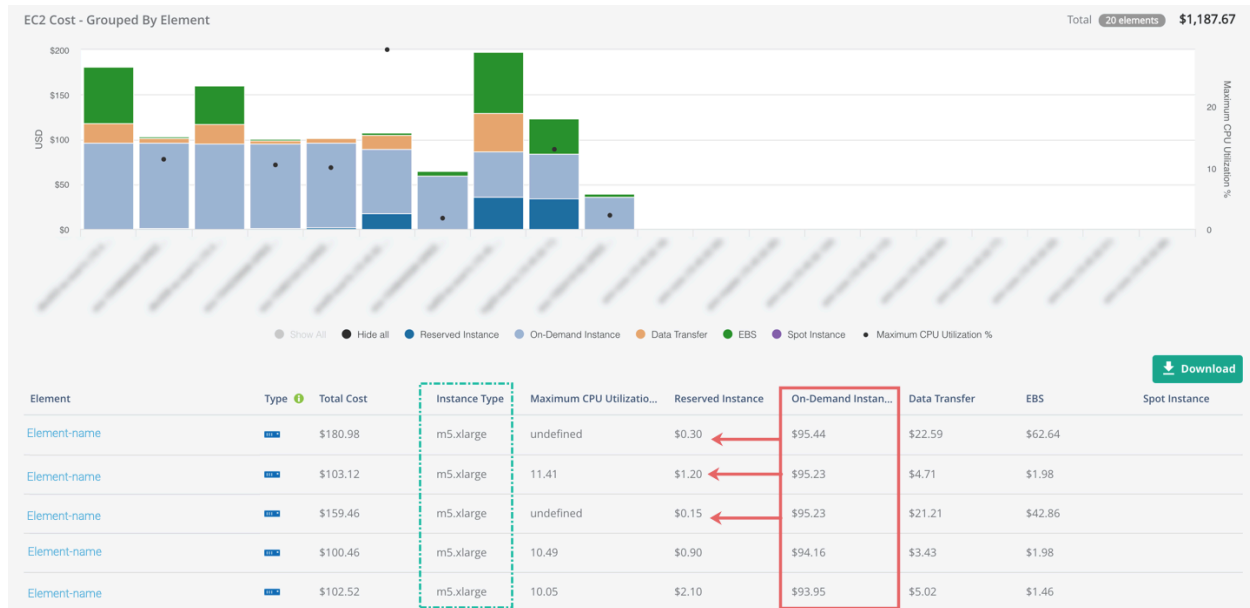
### How to view per-instance RI usage:

1. Open the EC2 Cost tool found in Cost Management
2. Select your time frame (such as daily) to establish the amortization breakdown
3. Group by element to see each resource individually
4. Sort the columns by Reserved Instance or On-Demand Instance



At this stage, you can slice your data however you'd like. For this example, we want to see all of the instances incurring On-Demand costs.

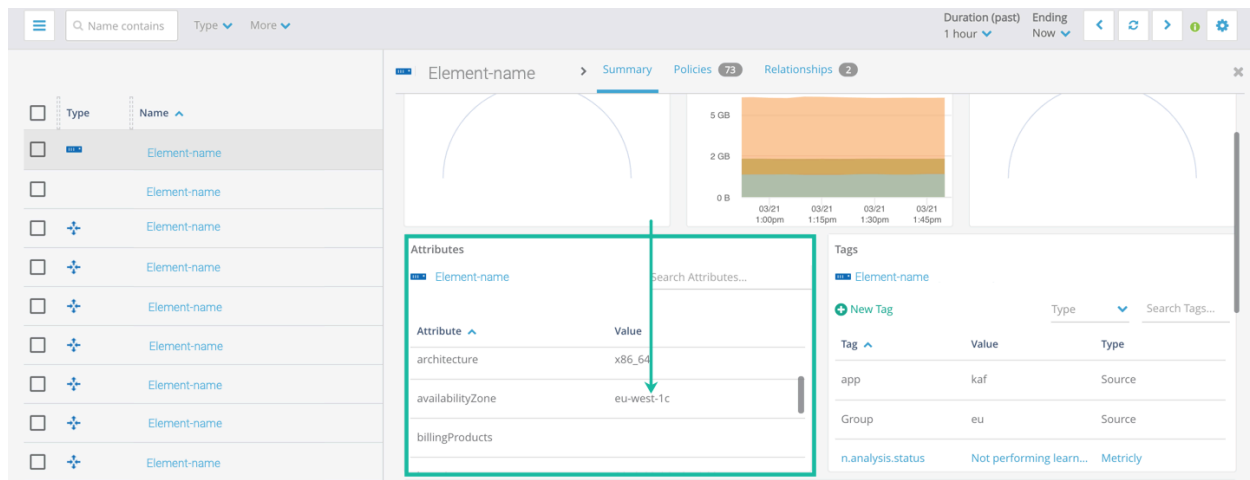
1. Set the filter to **On-Demand > 0**
2. Sort by largest **On-Demand Instance**



Comparing the On-Demand incurred costs of these instances against the costs covered by reservations shows that there could be an insufficient amount of reservations. Each of these instances are m5.xlarge, which is a good starting point for [determining what to purchase](#).

Now we want to know where these under-reserved instances reside.

1. Select the element's name and go to the **Element Detail Panel**
2. Find the Attributes card
3. Scroll to region, availabilityZone, and accountId



These details answer where reservations need to be purchased and allow you to decide:

- If zonal or regional reservations are better

- If you want the reservation on the element's account or a master billing account

You can always go back to the EC2 Cost tool and then filter by those same attributes to determine if there's a particular account or region that needs more reservations. With account-wide RI Usage reports, Reserved Instance Budget Alerts, and Metricly's EC2 Cost tool, you gain full control and visibility into your RI costs and usage.

## Understanding Convertible Reserved Instances

[Earlier in this guide](#) we went over the basic differences between Standard and Convertible Reserved Instances. In this section, we'll tackle Convertible Reserved Instances (CRI) in more depth and show you how to maximize their flexibility to your advantage.

The key selling point of CRIs is that they can be exchanged. The challenge they present, however, is that the exchange process can be confusing.

### 10 things you need to know about exchanging CRIs

1. You must exchange your CRI(s) for one of equal or greater value.
2. You can exchange multiple CRIs for one new CRI.
3. You can exchange one CRI for multiple smaller CRIs.
4. You can exchange Upfront CRIs for Partial Upfront CRIs, and vice versa.
5. You cannot exchange for CRIs in a different Region.
6. You cannot exchange Upfront CRIs for No-Upfront CRIs.
7. Exchanging multiple CRIs of the same term length for one new CRI sets the new expiration date to the furthest in the future of the exchanged CRIs.
8. Exchanging your CRI for one of greater value incurs a true-up cost.
9. Exchanging a No-Upfront CRI to an Upfront CRI gives you better pricing.
10. Merging multiple CRIs with different term lengths sets the new CRI for a 3 year term length.

Let's apply what we know and dive into an example to show the limitations and workarounds you'll need to be familiar with.

**Scenario:** 12 months ago you bought a Full Upfront 3 year r4.large CRI

**Situation:** You no longer need a high-memory instance

**Desired Outcome:** You want a general purpose instance, like an m4.large

**Problem #1:** You can't just convert an r4.large to an m4.large because the m4.large is cheaper

**Problem #2:** You don't want to buy more than you need (2 m4.larges) and incur a true-up charge

**Solution:**

1. First, we are going to use rule number #3 and split-convert our r4.large CRI into its price equivalent in t3.nanos. This gives us the smallest building block of individual CRIs to work with.
2. Now that we have *multiple* t3.nano CRIs of the same term length, we can exchange a subset of our t3.nano CRIs into the m4.large that we originally wanted. This avoids needing to overpay for a direct conversion from r5.large to multiple m4.larges.
3. Finally, we have our desired m4.large and a few remaining t3.nanos that we can use or merge with other t3.nanos in the same region.

All new CRIs in this scenario have 24 months remaining in term length. This is because we did not mix in other contracts with varying durations.

## Submitting an exchange request

We've gone over the rules to keep in mind and laid out a high-level plan for converting an r4.large into an m4.large + t3.nanos, but how do you do it? There are two ways:

- Via the AWS Console
- Via the Command Line Interface

### CRI exchange via AWS Console

1. Navigate to the [EC2 service in your AWS Console](#).
2. Navigate to **Instances > Reserved Instances**.
3. Choose the instance(s) to exchange.
4. Select **Actions > Exchange Reserved Instance**.
5. Choose your desired attributes and select **Find Offering**.
6. Find the **Instance Count** column. Notice it displays the number of new CRIs you receive for the exchange.
7. Select **Exchange** to confirm and commit to the trade.

Your exchange may take a few minutes to be reflected in your account.

### CRI exchange via command line

1. Find and select a target instance using [describe-reserved-instances-offerings](#).
2. Get a quote using [get-reserved-instances-exchange-quote](#).
3. Commit to the quote using [accept-reserved-instances-exchange-quote](#).

## The Pros and Cons of Convertible Reserved Instances

Now that you know the basic rules and exchange process for CRIs, let's recap how they measure up against Standard Reserved Instances.

- CRIs are more flexible, but offer less savings
- CRIs can incur additional fees when exchanging due to the true-up cost of choosing a larger instance
- CRIs cannot be sold on the AWS Reserved Instance Marketplace

CRIs are the better reservation choice for steady-state usage that may anticipate shifts in performance needs. This is because CRIs can switch performance families, operating systems, tenancy, and payment options. We recommend buying the largest CRI that makes sense for you and then, when necessary, split-converting it into smaller components as needs change. This helps you avoid true-up costs of incrementally upgrading your resources. Split-converting from one parent contract also avoids unnecessarily refreshing and complicating your term length expiration dates.

## Selling a Reserved Instance contract

You have the option to sell your contracts before they expire thanks to the [AWS Reserved Instance Marketplace](#). This can be an ideal solution to optimizing your cloud budget when capacity needs change abruptly. Let's go over how the AWS RI Marketplace works.

### Basic requirements of selling a Reserved Instance:

- Reservation must be at least 30 days old
- Reservation must have at least 30 days remaining
- Seller must have a US bank account
- GovCloud Region instances cannot be resold
- Convertible Reserved Instances cannot be resold

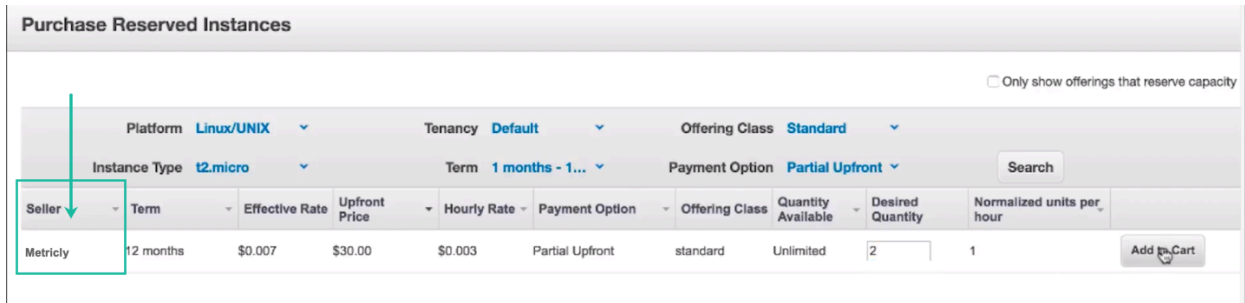
Selling your contracts first requires you to set up a [Seller Registration Page](#) in the AWS Console.

The screenshot shows the 'Account Information' step of the 'Reserved Instance Marketplace' seller registration process. At the top, there are three steps: 'Account Information' (selected), 'Manage Bank Account', and 'Confirmation'. Below this, the title 'Reserved Instance Marketplace' is displayed. Under 'Account Information', a message states: 'Enter your business information below so that Amazon.com may contact you regarding your Selling account. Information about our customers is an important part of our business, and we are not in the business of selling it to others. We share customer information only as described in the [Privacy Notice](#).' A 'Business Information' section contains a 'Business Name' field with a placeholder text: 'Name of the Seller of Record that will appear on any legal documentation Amazon.com sends you.' A 'Continue' button with a right arrow is located at the bottom right of the form.



Once you have set up your Seller Registration Page, provided a bank account, and tax ID, you can then select reservations to put on the marketplace. There is a \$50,000 yearly limit to reselling instances that can be exceeded by contacting AWS and filling out [this form](#).

Buyers can see who is selling the reservation by checking the **Seller** column.



<input type="checkbox"/> Only show offerings that reserve capacity										
Platform <b>Linux/UNIX</b>		Tenancy <b>Default</b>		Offering Class <b>Standard</b>						
Instance Type <b>t2.micro</b>		Term <b>1 months - 1...</b>		Payment Option <b>Partial Upfront</b>		<input type="button" value="Search"/>				
Seller	Term	Effective Rate	Upfront Price	Hourly Rate	Payment Option	Offering Class	Quantity Available	Desired Quantity	Normalized units per hour	
Metricly	12 months	\$0.007	\$30.00	\$0.003	Partial Upfront	standard	Unlimited	<input type="text" value="2"/>	<input type="text" value="1"/>	<input type="button" value="Add to Cart"/>

It can take up to 2 weeks for your first sale to be processed. This is because AWS needs to set up a relationship with your US-based bank to send wire transfers to your account. All other sales typically take between 1-3 days for processing.

Important things to keep in mind:

- A 12% AWS seller fee is applied to the finalized sale based on the total upfront price you set on the reservation
- You retain the reservation and capacity benefits of your on-market reservation until the sale's transaction is complete
- A reservation's remaining duration is rounded down to the nearest month for the buyer
- If you are outside of the US, consider setting up a US bank account before purchasing reservations. **AWS only deposits marketplace sales in US-based accounts.** While alternatives such as Transferwise do exist, you should do a thorough vetting of your options before choosing what works best for you

## Should I sell my Reserved Instance?

Whether or not you sell your Reserved Instance contract should be determined by a two key things:

1. Whether the net return of selling your contract (accounting for the 12% AWS seller fees) is more valuable to you than the reservation's remaining term
2. Whether you foresee needing more capacity of that type in that region for other current or future projects

If your existing reservation can be applied to other resources and still be maximally utilized, it may be best to keep it. Verify across your AWS organization (linked accounts) that you can't use the reservation before selling.

## Creating a future-focused reservation purchase plan

AWS regularly provides new, more powerful instance generations. These newer generations are put at competitive price points to encourage users to adopt future solutions. As time passes, older generations become more expensive. Knowing this is important when deciding what instance generation and term to commit to.

Committing to a 3 year t2 instance contract, for example, is less ideal than committing to the same term length for the newest, t3 generation. Committing to the older generation for such a long time could wind up negating your savings if the price of t2 rises each year. However, committing to a 1 year contract for t2 would mitigate the risk of rising costs if you are not ready to commit to a t3.

Overall, newest generations are the best to commit to because they offer the deepest savings. 3 year commitments are only ideal for newest generations because baseline price rises for older generations and that rise eats into your savings.

## Advanced

### Scheduled Reserved Instances Explained

For bulk workloads that only happen on a routine basis—once per day, per week, or per month—a Standard or Convertible RI contract might go underutilized. This is where Scheduled Reserved Instances shine: you still get reserved capacity and discounts without worrying about overpaying for an underutilized contract.

#### How Scheduled RIs Work

AWS dedicates a pool of EC2 resources within each Availability Zone for Scheduled RI use. These pools support a specific (limited) combination of instance type, operating system, and network. Those limitations should be considered before purchasing a Scheduled Reserved Instance contract.

#### Scheduled RI Limitations

- Only Supports C3, C4, M4, and R3 instance types
- 1 year terms (365 days) required
- Minimum utilization required is 1,200 hours per contract
- You cannot purchase a Scheduled RI contract more than 3 months in advance

- Limited to the following Regions: US East (N. Virginia), US West (Oregon), and Europe (Ireland)
- You cannot cancel, modify, or resell your Scheduled RI contract

## Scheduled RI Availability

Scheduled RIs are bought based on the time slots available in each EC2 pool for a given Availability Zone. Not every EC2 pool has your desired schedule available, so searching through several pools across your desired Availability Zones is important.


These reservations operate around firm timelines and terminate 3 minutes before the end of your purchased timeframe. This is because Scheduled RIs are part of a shared resource pool and must be available for any upcoming reservations in the proceeding timeframe.

Availability is also affected by terminations. Scheduled RIs can be terminated manually, however if they are not relaunched within a few minutes after termination, you will have to wait until the next interval to relaunch.

## Common Reserved Instance Pitfalls to Consider

### Zonal vs Regional Purchasing

It's easy to accidentally purchase a Zonal Reservation when you meant to purchase a Regional Reservation. Selecting the checkbox **Only show offerings that reserve capacity** switches your options from Regional to Zonal. Make sure you uncheck this box when buying Regional reservations.



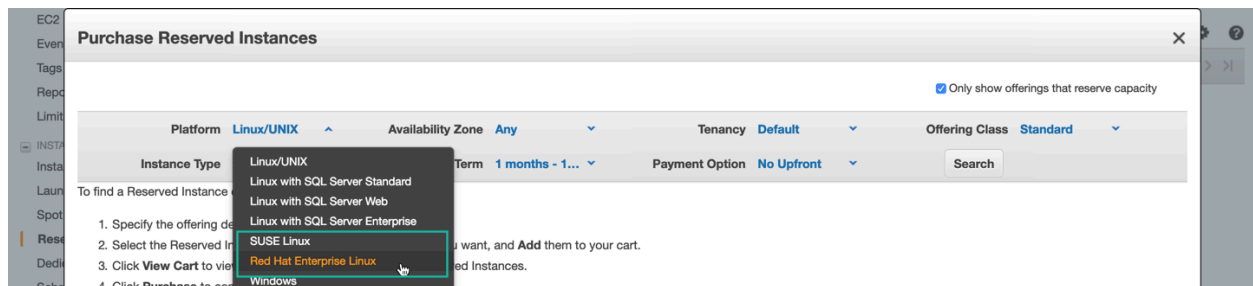
### Reservation Purchasing Cadence

- **Vendor lock-in:** If your company intends to move to another cloud provider, a more frequent purchasing cadence can cause stickiness with AWS that makes it challenging to pinpoint the right time to move. In this instance, you would want to stop purchasing RI Contracts and consider putting your existing ones on the RI Marketplace.
- **All-Upfront purchasing:** If your company prefers to buy bulk all-upfront contracts at certain points in the year, it may not be feasible to change your purchasing cadence.

- **Potential Overhead:** Buying more frequently can capture more savings and prevent large expiring reservation cliffs, but it also requires more regular attention from stakeholders for approval. Make sure your team has the bandwidth for more routine decision making.

## Platform Selections

A common mistake users make is selecting the wrong platform for the type of instance they need. For example, Linux/Unix is the wrong platform choice for running RHEL or SLES. RHEL and SLES have their own options in the Platform dropdown.



## Reserved Instance Trends

What does a typical EC2 portfolio look like for a business hosted by AWS? In this section, we're going to explore some statistics around reservation purchases. Metricly analyzes tens of thousands of Reserved Instance contracts for its customers, and over the years we've noticed a few trends you might find interesting.

### No-Upfront Reserved Instances are the most popular

Over 90% of businesses prefer No-Upfront RIs to any kind of upfront payment. Most customers prefer sacrificing small overall savings to avoid upfront payments altogether. Here's a quick look at purchasing trends from the past 3 years:

#### Overall contract value across payment methods

Year	All Upfront	Partial Upfront	No Upfront
2018	4%	2%	94%
2017	16%	2%	82%
2016	9%	11%	80%

## Convertible Reserved Instances are gaining popularity

Convertible reservations were introduced in late 2016 and have been growing in popularity over the last few years due to the additional flexibility they provide.

### Percentage of total instances across contract types

Year	Standard	Convertible
2018	85%	15%
2017	92%	8%
2016	99%	<1%

Even though their adoption has nearly doubled since 2017, Convertible Reserved Instances still only represent a small subset of reservations. Their more limited savings and benefits may be preventing them from toppling Standard RIs as the go-to contract type. It's also likely that most AWS customers who require this level of flexibility are not quite ready to commit to reservations in general, as their infrastructure needs have yet to become predictable enough.

## Regional Reserved Instances favored over Zonal Reserved Instances

The option to sacrifice actual capacity reservation for flexibility with a regional scope was introduced in late 2016 and quickly become the standard most customers use with 99% opting for this in 2018.

### Percentage of total instances across scope options

Year	Region	Availability Zone
2018	99%	1%
2017	92%	8%
2016	4%	96%

We have found that most AWS customers we work with are not aware that regional reservations do not reserve capacity. Similarly, many AWS customers we work with who have purchased zonal reservations are not aware of the implications of their decision to reserve capacity.

## Regional Standard RIs are the most popular Scope & Offering Class combination

Scope and Offering Class options both affect the flexibility you have in the choice of instances. Take a look at the combination trends across the past 3 years:

### Percentage of total instances across Scope & Offering Class combinations

Year	Region / Convertible	Region / Standard	AZ / Convertible	AZ / Standard
2018	15%	84%	< 1%	1%
2017	8%	84%		8%
2016	< 1%	4%		96%

- Regional Standard Reserved Instance combinations provide flexible scope and (if on Linux/Unix platform) size flexibility. They are also sellable on the RI Marketplace.
- Zonal Convertible Reserved Instance combinations are the least common. In all likelihood, such a combination is due to user error when selecting the Scope of their contract (as described in our pitfalls section). This is because Convertible Instances are designed for flexibility, whereas Zonal scopes are not.

## Top 5 most popular Instance Families

Out of over 150 different Instance types in around 50 different series (including the generation) there is a surprising lack of diversity in the instance types customers reserve. In 2018 over 93% of reservations by instance count fell into just 5 families:

### Percentage of total instances across top 5 instance families

Instance Family 2018	Contribution	Instance Family 2017	Contribution	Instance Family 2016	Contribution
m5	28%	t2	33%	t2	42%

t2	24%	m4	26%	c4	25%
m4	23%	c4	18%	m4	9%
i3	16%	m3	9%	m3	9%
c5	3%	i2	9%	r3	9%
Total:	<b>94%</b>	Total	<b>95%</b>	Total:	<b>94%</b>

- The top 3 instance families have included t2 and m4 instances each year
- The top 3 instance families have contributed over 75% of the total reservations each year

The m5 and c5 families were introduced in November 2017 and are slightly cheaper than the older generation equivalent. We expect those to be gradually replaced over time.

## Beyond

This guide has covered a wide variety of topics on EC2 Reservations, but it doesn't stop here. [Tell us what you'd like to know more about](#) or subscribe to our Cloud Compass Newsletter. Every month we publish new articles that explore ways you can save money and be more efficient in the public cloud.