

Buyer's Guide for **AWS Reserved Instances**

Most organizations can and should benefit from the convenience and cost savings of computing and data storage capabilities offered by the cloud. Many gain these benefits by running their computing workloads on the Amazon EC2 service offered by Amazon Web Services (AWS).

To meet the many different needs and use cases of its customers, AWS offers several ways to purchase compute capacity for its EC2 instances. Reserved Instances (RIs), specifically, present a tremendous opportunity to save significantly on capacity. However, a variety of factors must be considered to accurately budget and forecast cloud costs to fully reap the benefits of AWS Reserved Instances.

In this guide, you will learn:

- How AWS structures, packages, and sells RIs.
- What to consider when purchasing RIs,
- Why you must manage RIs to make good on the cost-savings they offer,
- Effective workflow for managing your RIs, and
- How CloudCheckr provides insights, tools, and automation for optimizing RI investments.

RIs in Context of Amazon's EC2 Instance Offerings

You can purchase Amazon EC2 instances in three ways: On-Demand, Spot, and Reserved. Each differs in cost, payment structure and timing, flexibility in use, and typical use case. RIs offer an extremely attractive option due to their great potential for cost savings.

On-Demand

As the name suggests, you pay a pre-specified hourly rate for the capacity you use. The flexibility to only pay for the capacity you need when you need it comes at a higher cost.

Use On-Demand instances when you have workloads that can't be interrupted, that are short-term, spike in demand, or are unpredictable in capacity.

Spot Instances

Spot allows you to bid on excess On-Demand capacity, and offer even greater potential for savings. The downside? There's no guarantee you'll get it. The upside? The typical, but not guaranteed, Spot rate is substantially lower than On-Demand cost (up to 85% less than On Demand rates).

Use Spot Instances for applications that are not time sensitive, that you couldn't afford to run at the higher On-Demand rate, or for stateless applications that can still function with parts of the application turned off at any given time. You might also use them if you have an immediate need for a lot of extra capacity.

Reserved Instances

You commit to purchasing a specific amount and type of capacity for one or three years, and pay whether you use that capacity or not. Because of the commitment, your hourly rate for capacity can drop the hourly cost by as much as 65 percent from the On-Demand rate.

Use RIs when you have applications that use a fairly steady or predictable amount of compute capacity. In other words, RIs are ideal when you have a pretty good sense of how much capacity you need, so you don't waste money on capacity you won't use.

Why You Need to Know

RIs, when purchased and managed appropriately, can deliver cost savings by dramatically decreasing your AWS bill. The risk of purchasing capacity that goes unused and idle can completely negate any potential cost savings. To manage them properly, you have to take your understanding of RIs a level deeper.

RIs are especially useful to ensure you always have access to reserved capacity in Amazon’s data centers for applications that are critical (unless you are buying regional instances). This guarantees the ability to use or launch that instance, even if capacity issues arise in your availability zone.

Structure of AWS

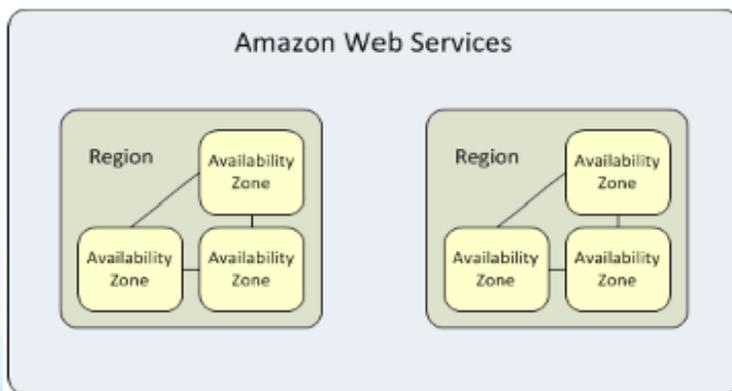
AWS instances are physically grouped in global locations around the world called Regions and Availability Zones (AZs). Each location offers hundreds of different permutations of instances that you can reserve.

AWS Regions and Availability Zones (AZs)

AWS Cloud currently operates data centers around the world and is adding Regions and AZs on a regular basis. AWS offers a variety of availability zones across 15 global regions, including North America, South America, the European Union, Asia Pacific, and China. A separate GovCloud exists to meet iTAR compliance. AWS has announced plans for expanding into additional regions in the near future.

Each AWS Region is composed of one or more Availability Zones (AZs). An AZ is physically housed in one or more separate buildings. For example, the US West Oregon Region has three AZs, which are called US West - Oregon 1a, 1b, and 1c.

The image below provides a quick visualization of the concept of AWS Regions and AZs:



Why You Need to Know

Understanding Regions and AZs can ensure you are making smart RI purchases to reduce potential application latency. For example, you want to launch workloads geographically close to where they’re used, and buy RIs based on that location. In addition, if you want to make a capacity reservation, you must designate a specific AZ for the RI.

Figure 1. AWS Regions and Availability Zones (from AWS)

Permutations of AWS Instances

Regardless of the type of EC2 instance you need (On-Demand, Spot, or Reserved), you have numerous choices for the instance type you use or reserve. These choices are defined by several attributes, including the Region or Availability Zone they're located in, the instance's underlying Operating System, the Tenancy of the hardware, and the Instance Type/Size. The Instance Type specifies the CPU, memory, storage, and networking capacity.

The possible number of instances from which you can choose becomes staggering when you consider how many instance types AWS currently offers (as shown below) multiplied by the number of OSes, Regions/AZs, and tenancy:

Why You Need to Know

Deciding what to purchase takes significant thought, but the real effort begins once you purchase the RIs. At that point you need to constantly realign your workloads with the instances you've reserved to ensure that instances you have purchased at a discount do not go to waste. By understanding the various instance types AWS offers, you better understand the guidelines around shifting workloads to new instances. You'll also need to evaluate and determine if you need to purchase more instances to meet workload needs.

Instance Family	Current Generation Instance Types
General Purpose	t2.nano t2.micro t2.small t2.medium t2.large t2.xlarge t2.2xlarge m4.large m4.xlarge m4.2xlarge m4.4xlarge m4.10xlarge m4.16xlarge m3.medium m3.large m3.xlarge m3.2xlarge
Compute Optimized	c4.large c4.xlarge c4.2xlarge c4.4xlarge c4.8xlarge c3.large c3.xlarge c3.2xlarge c3.4xlarge c3.8xlarge
Memory Optimized	r3.large r3.xlarge r3.2xlarge r3.4xlarge r3.8xlarge r3.16xlarge x1.16xlarge x1.32xlarge
Accelerated Computing	p2.xlarge p2.8xlarge p2.16xlarge g2.2xlarge g2.8xlarge f1.2xlarge f1.16xlarge
Storage Optimized	i2.xlarge i2.2xlarge i2.4xlarge i2.8xlarge
Dense Storage Instances	d2.xlarge d2.2xlarge d2.4xlarge d2.8xlarge

Table 1. Currently available types of Amazon EC2 Instances.

For example, you could run a Linux c3.4xlarge instance (from the Compute Optimized family) in the US West Oregon Region, on a server with Dedicated tenancy.

Considerations When Purchasing RIs

Deciding what type of RIs you want to purchase is only half of the puzzle. Once you know what to purchase, there are still a handful of factors that influence potential cloud cost savings— particularly when you consider purchasing more RIs later on. They also impact the flexibility you have when balancing workloads across your RIs.

Standard versus Convertible RIs

You can purchase RIs in two ways, depending on how much flexibility you need when reallocating them to optimize utilization. You can purchase a Standard RI for the maximum discount possible with less flexibility, or you can purchase Convertible RIs for a smaller discount.

By now, you're likely seeing the common theme with AWS: the more flexibility you need, the more it costs. But of course, the more you can accurately plan and manage for capacity needs, the more you can save.

For both Standard and Convertible RIs, you can use them within any Availability Zone (AZ) in a given AWS region. You can also obtain a capacity reservation if you assign your RIs to a specific AZ.

With Standard RIs, what you reserve is exactly what you get. You can use your RIs within any availability zone in a given AWS region. With Linux OS (except Red Hat and SUSE Linux), you can change the instance size, as well as combine and split RIs with the same start/end hour. You can't trade out for a different Instance Family, Operating System, Tenancy, or Payment Option. Although there is notably less flexibility this way, Standard RIs offer a maximum discount of up to 75 percent of the On-Demand hourly rate.

Convertible RIs let you trade purchased RIs for other RIs, as long as the trade is for RIs with an equal or greater value. You might want to do this if you have capacity available for which you have no use—for example, if you reserved too much capacity, or a project is canceled. The safeguard of added flexibility comes at a price, as the maximum discount you'll get is up to 45 percent of the On-Demand hourly rate. However, you can change Instance Family, Operating System, Tenancy, or Payment Option.

The time period for which you've purchased the convertible RI has an impact on a conversion, too. For example, if after one year of a three-year Convertible RI reservation, you convert it to something else, your new reservation will be for 2 years.

Commitment Period

Standard and Convertible RIs differ in how long you commit to paying for capacity. With Standard RIs, you can purchase with either a one-year or three-year commitment; with Convertible RIs, you can only purchase with a three-year commitment.

Payment Options

Both Standard and Convertible RIs can be purchased in three different ways: by paying All Upfront, Partial Upfront, and No Upfront. With All Upfront, you pay the entire amount for the term to which you've committed. With Partial Upfront, you pay a specific portion of the entire bill upfront, and then monthly payments for the duration of the commitment period. No Upfront simply spreads out the entire payment for the bill in equal payments for the duration of the commitment period. The bigger hourly rate discount comes from paying All Upfront, and decreases as you move toward No Upfront.

Reserved Instance Marketplace

Those who have purchased Standard RIs have the option to sell their unused RI capacity on the AWS Reserved Instance Marketplace. In cases where you have no use for the capacity, this can be a viable option. However, you'll experience some downsides to selling RIs in the Marketplace: resale value is typically less than the original value of your purchase, you'll be charged a seller fee by AWS, AWS will round down what you sell to the nearest month, and it can take a while to sell your RIs.

Master Accounts

Grouping all AWS accounts in your organization under a single master account helps you obtain even greater pricing discounts because you're purchasing more RIs. For example, if you purchase \$4M US worth of RIs, AWS discounts the

Why You Need to Know

When purchasing RIs, you need to have an idea of how much flexibility you'll need for shifting workloads around. For instance, will you need to simply shift an instance in one region to the same instance type in another region, or will you need to switch from a Linux-based system to a Windows-based system? Your choice of Standard versus Convertible RI defines your ability to shift workloads.

In addition, you have to weigh the opportunity cost of commitment and payment options against potential cost savings. If consistent RI capacity is more critical than budget restraints, a three-year term with an All Upfront payment option is ideal to recognize the greatest cost savings. But if your business is like many, the choice is not crystal clear; you'll need to know the cost savings associated with various combinations of time periods and payment options.

purchase by 5 percent. While the bulk purchase discounts can be significant, and it's beneficial to be able to trade and share resources among accounts, the use of master accounts has a drawback that's worth considering: one department within the organization may use the resources that another department reserved, and anticipated using for another purpose.

Managing Your Reserved Instances

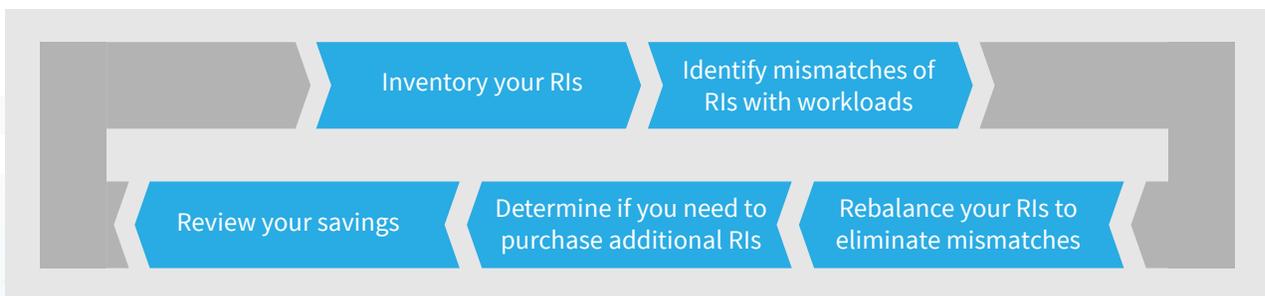
Determining the right AWS RI purchase takes significant time and effort. At this point, many businesses simply step away and let the business use the reserved capacity. But to fully realize the true value that RIs can offer, this is when you need to pay attention the most.

Understandably, the number of moving parts to keep watch over can be overwhelming and even paralyzing, especially if you have thousands of RIs. That's where a reliable workflow for optimizing use of your RIs, coupled with sophisticated algorithms and automation, can help.

In this final part of the guide, you'll learn a repeatable workflow that you can use with CloudCheckr reports, alerts, and automation to make optimally managing your RIs completely possible.

The Workflow

The process you need to follow to manage your RIs boils down to these five steps that are part of continuous RI management process.



Why You Need to Know

You need visibility across all accounts under your AWS master account to have a more comprehensive understanding of who is using what RIs and to have better control over that usage. This visibility avoids situations where resources intended for one project inadvertently get used by another. This also ensures you can allocate costs appropriately across teams or accounts.

How CloudCheckr Helps with Each Step

At each step in the lifecycle management of your RIs, CloudCheckr plays a key role to either do the work for you, or distill out high-level information you need to make the best decision.

Step 1. Inventory your RIs

You can't manage what you don't know about. The CloudCheckr EC2 RI Summary tells you exactly how many RIs you have by Utilization Type (based on payment option and level of usage) and Offering Class (Standard or Convertible).

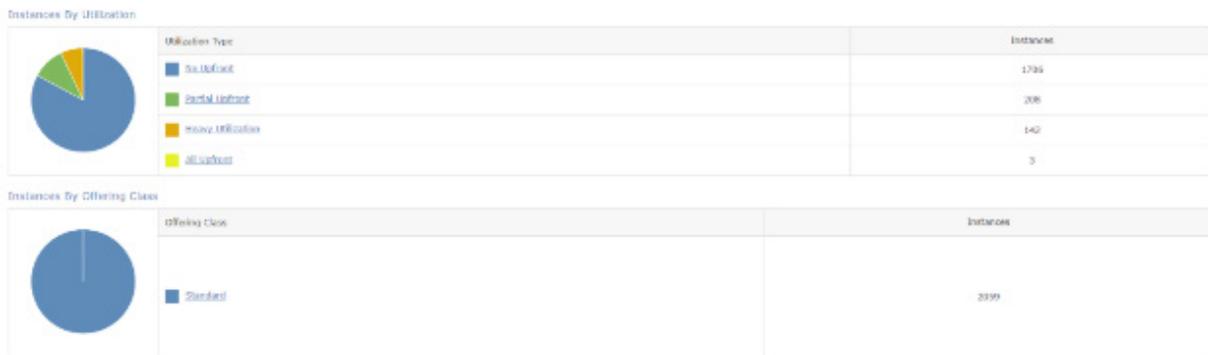


Figure 2. RI Summary Report

The List of Reserved Instances report gives you a complete list of the individual RIs you've purchased, along with details such as the AZ, Region, Hourly Price, Instance Count, and Instance Type, and much more.

	Reserved Instance ID	Availability Zone	Region/Zone	Start Date	End Date	Hourly Price	Instance Count	Instance Type	Offering Type	Offering Class	Scope	On-Demand Hourly Price	Price Protection	Start Time	Term	Time Left	Options Fee	Endgame Fee	Endgame	AMZ Account ID
✓	8-2u3k0-03d-0750-838v4-1288208a2	us-east-2a	US East (Ohio/Virginia)	1 Year	5/12/2017 5/12/2018	\$0.092	1	m4.xlarge	No Utilization	Standard	Availability Zone	\$0.222	Unprotected	11/12/2016 11/02/2016	Default	138 days	\$0.00	0	None	442094636793
⚠	8-7129-2d-1a1t-4m5-4m4-758876142	us-east-2a	US East (Ohio/Virginia)	1 Year	5/12/2017 5/12/2018	\$0.072	1	m4.xlarge	No Utilization	Standard	Region	\$0.201	Unprotected	10/18/2016 5/25/2016	Default	227 days	\$0.00	0	---	218011095617
✓	4-28n2d8-e88d-4081-8721-04708680327	us-east-2a	US East (Ohio/Virginia)	1 Year	5/12/2017 5/12/2018	\$0.092	1	m4.xlarge	Partial Utilization	Standard	Region	\$0.212	Unprotected	10/19/2016 5/26/2016	Default	161 days	\$0.00	0	None	218011095617
✓	5-8P8a-899-4400-918d-4701870718	us-east-2a	US East (Ohio/Virginia)	3 Years	8/22/2015 1/12/2018	\$0.065	1	m4.xlarge	No Utilization	Convertible	Region	\$0.213	Unprotected	8/22/2015 1/12/2018	Default	1,041 days	\$0.00	0	None	218011095617

Figure 3. List of Reserved Instances

Step 2. Identify Mismatches of RIs with Workloads

The List of Reserved Instances gives you a quick visual cue that an RI is not fully utilized—either a green checkmark or a red arrow—with additional details available within the dashboard. For example, you can see who purchased the instance, who’s using it, and more. After all, you’ve paid for the RI—you want to make sure you’re applying it to a workload. If it’s not being applied, knowing why can help you act to remedy it.

Step 3. Rebalance Your RIs to Eliminate Mismatches

Once you’ve inventoried your RIs, you can use CloudCheckr’s RI Rebalancer tool to apply unused RIs to other workloads. For example, if you spun up a workload on an instance in the US West (Oregon) region in the Availability Zone (AZ) us-west-2a, but that workload fails over to a different AZs, for instance us-west-2b, you would benefit by moving the RI from the old AZ (us-west-2a) into the new AZ (us-west-2b).

CloudCheckr matches up any unused RIs with workloads that need capacity to align parameters, rules, and utilization needs to provide actionable insights. For example, the app integrates OS, Region, AZ, Instance Type, and the amount of flexibility for redirecting based on whether you purchased a Standard or Convertible RI.

CloudCheckr does all this deep analysis under the hood, then enables you to act on the best recommended action to rebalance RIs directly in your dashboard. With the click of a button, you can rebalance your RI and the workload. Optionally click a checkbox and count on CloudCheckr to automatically rebalance your workload whenever it discovers a mismatch.

Reserved Instance ID	Product Description	Instance	Instance IDs matching before	Change	Instance Type	Region	Availability Zone	Instance Count	Instance IDs To Match	AWB Account	
✓ 71127512-2419-46d1-b046-716825216112	Linux	Default	None	From To	m3.xlarge m3.xlarge	US East (Northern Virginia)	us-east-1a	1 1	i-784887ad, i-0a3b4076, i-22879511, i-01c19165e, i-1714f564, i-079180e0, i-0a3b4076, i-22879511, i-01c19165e, i-1714f564, i-079180e0	218511780027 (AmazonChime)	Rebalance

Figure 4. CloudCheckr Rebalance tool

Step 4. Determine if You Need to Purchase Additional RIs

Once you’ve optimized workloads by rebalancing your RIs, you may discover that you still need additional RIs to meet capacity needs. Based on those needs, CloudCheckr provides a comprehensive Purchase by Instance Report to show recommended RIs to purchase, along with a list of RIs available for purchase.

Additionally, the report includes critical details that allow you to make the purchase decision to best fit your budget and desired level of cost savings. It even shows you how long it will take for you reach the break even point of your purchase to start realizing cost savings.

Recommended Reserved Instance Purchase

Region	Instance Type	Platform	Tenancy	Number Recommended	On-Demand Hourly Cost	No Upfront	Partial Upfront	All Upfront
US East (Northern Virginia)	m3.large	WindowsStandard	Default	1	\$0.7540	RI Hourly Cost: \$0.4920 RI On-Demand Hourly Cost: \$0.4420 (\$0.4420 hourly cost * 8,760 hours in year) = \$3,850.80 upfront / 8,760 hours in year) Current Hourly Cost: \$0.5111 Optimal Hourly Cost: \$0.5050 Monthly Savings: \$22.98 Break Even: 6/10/16 Returns On Investment (vs On-Demand): 104% Returns On Investment (vs No Upfront): 111%	RI Hourly Cost: \$0.5195 RI On-Demand Hourly Cost: \$0.3790 (\$0.3790 hourly cost * 8,760 hours in year) = \$3,328.00 upfront / 8,760 hours in year) Current Hourly Cost: \$411.14 Optimal Hourly Cost: \$100.71 (including amortized RI cost) Monthly Savings: \$240.42 (\$22.20 including amortized RI cost) Break Even: 8/10/16 Total Savings: \$1,815.11 Returns On Investment (vs On-Demand): 117.61% Returns On Investment (vs No Upfront): 12.21%	RI Hourly Cost: \$2.8500 RI On-Demand Hourly Cost: \$0.5007 (\$0.5007 hourly cost * 8,760 hours in year) = \$4,378.00 upfront / 8,760 hours in year) Current Hourly Cost: \$411.14 Optimal Hourly Cost: \$0.08 (\$0.08 hourly cost including amortized RI cost) Monthly Savings: \$410.14 (\$410.22 including amortized RI cost) Break Even: 0/10/16 Total Savings: \$4,691.63 Returns On Investment (vs On-Demand): 93.51% Returns On Investment (vs No Upfront): 18.24%

EC2 instances available for RI purchase

Instance	Time Remaining	On-Demand Hourly Cost	On-Demand Monthly Cost
i3.2xlarge (New/Existing)	\$0.0076 (including usage of 176.00 hours out of 720 hours in a month)	\$0.7040	\$411.24

RI Purchasing Options

Term	On-Demand Hourly Cost	RI Hourly Cost	Monthly Savings	Upfront Cost	Break Even	Total Savings (including upfront cost)	Returns On Investment (vs On-Demand)	Returns On Investment (vs No Upfront)	Total Commitment
1 Year All Upfront	\$411.14	\$0.00	\$411.14	\$3,228.00	8 months	\$1,689.62	22.20%	18.24%	\$3,228.00
1 Year Partial Upfront	\$411.14	\$180.71	\$240.42	\$1,210.00	3 months	\$1,621.11	110.62%	42.26%	\$3,211.50
1 Year No Upfront	\$411.14	\$302.66	\$108.48	\$0.00	9 months	\$1,061.71	N/A	N/A	\$3,971.62

Figure 5. Purchase by Instance Recommendation Report

Step 5. Review Your Savings

Sooner or later, your company should verify that the investment you made in RIs paid off—and by how much—to justify continued investment. The Historical Savings Report summarizes how much savings you realized by using RIs versus comparable On-Demand resources.

Total Savings	Amortized Reserved Upfront Cost	Net Savings
\$193,721.94	\$24,179.97	\$169,541.97

Month	Size	Type	Tenancy	Region	Hours	Savings	Amortized Cost
December 2016	m3.large	WindowsSqlStandardVpc	Default	US East (Northern Virginia)	180	\$47.16	\$0.00
December 2016	m3.medium	LinuxVpc	Default	US East (Northern Virginia)	61,069	\$1,399.19	\$198.43
December 2016	c4.large	LinuxVpc	Default	US East (Northern Virginia)	3,912	\$256.69	\$126.03
December 2016	m4.large	LinuxVpc	Default	US East (Northern Virginia)	12,972	\$734.87	\$312.08
December 2016	t2.medium	LinuxVpc	Default	US East (Northern Virginia)	36,555	\$587.12	\$121.11
December 2016	t2.micro	LinuxVpc	Default	US East (Northern Virginia)	34,792	\$150.20	\$34.82
December 2016	t2.small	LinuxVpc	Default	US West (Oregon)	7,286	\$47.27	\$0.00
December 2016	m3.medium	LinuxVpc	Default	US West (Oregon)	47,592	\$988.66	\$193.92
December 2016	t1.micro	LinuxVpc	Default	US East (Northern Virginia)	3,391	\$31.37	\$30.56
December 2016	t2.medium	LinuxVpc	Default	US West (Oregon)	2,440	\$43.43	\$0.00
December 2016	c3.xlarge	LinuxVpc	Default	US East (Northern Virginia)	2,086	\$297.46	\$0.00
December 2016	m3.large	LinuxVpc	Default	US East (Northern Virginia)	15,388	\$889.14	\$287.96
December 2016	c4.xlarge	LinuxVpc	Default	US West (Oregon)	6,317	\$986.59	\$1,198.29

Figure 6. Historical Savings Report

Actionable Insights for Keeping RIs in Check

So how often should you check to see if you have underutilized RIs or need to purchase additional RIs to meet workload needs? How do you know when your purchase for an RI failed, is pending, or is about to expire so that you can determine if you need to renew it? How do you know when AWS has dropped its price on an instance type so that you can re-negotiate for a better price?

Ideally, you'd immediately know when any of those situations occurred so that you can take the appropriate action. With the RI Lifecycle Management screen, this view lets you set up email notifications that get triggered when these situations occur.

The screenshot shows the 'EC2 Reserved Instances' configuration page. It lists several events with corresponding email notification fields, all containing the placeholder 'user@domain.com':

- EC2 RI is purchased**: Email to notify: user@domain.com
- EC2 RI purchase is in pending mode**: Email to notify: user@domain.com
- EC2 RI purchase failed**: Email to notify: user@domain.com
- EC2 RI is not fully utilized**: Email to notify: user@domain.com
- 30 days before an EC2 RI expires**: Email to notify: user@domain.com
- 7 days before an EC2 RI expires**: Email to notify: user@domain.com
- EC2 RI price is dropped within 30 days of purchase**: Email to notify: user@domain.com

Figure 7. RI Lifecycle Management Email Notifications

Using CloudCheckr to Make Good on the Cost Saving Potential of RIs

If your enterprise-level organization has purchased hundreds or thousands of RIs, you need to carefully manage them to realize the greatest amount of cost savings. With the many variables you must take into consideration to optimally manage your RIs, manually doing so becomes increasingly complicated.

CloudCheckr, a unified cloud governance solution, is specifically designed for AWS to provide the help you need to meet the challenge of managing your RIs. Following our five step workflow, while leveraging CloudCheckr's easy-to-use tools, automated processes, and actionable reports gives you all that you need to get the most from your AWS investment.