
Amazon Cost and Usage Reports

User Guide

亚马逊云科技



Amazon Cost and Usage Reports: User Guide

Table of Contents

What are Amazon Cost and Usage Reports?	1
How Cost and Usage Reports work	1
Report timeline	1
Report files	2
Report columns	2
Using your report	2
Creating reports	3
Setting up an Amazon S3 bucket for Cost and Usage Reports	3
Creating Cost and Usage Reports	4
Managing reports	6
Viewing your report	6
Viewing the latest report version	6
Viewing your finalized report	7
Understanding your report versions	7
Amazon CUR delivery timeline	8
Creating new Cost and Usage Report versions	8
Overwriting previous Cost and Usage Reports	8
Cost and Usage Reports manifest files	9
Editing reports	11
Using Cost and Usage Reports for Amazon Organizations	12
Managing Cost and Usage Reports as a member account	12
Managing Cost and Usage Reports as a management account	12
Querying reports using Athena	13
Setting up Athena with CloudFormation	13
Setting up Athena manually	14
Creating an Athena table	15
Creating a report status table	15
Uploading your report partitions	16
Running Athena queries	16
Column names	17
Other resources	17
Loading report data to Amazon QuickSight	17
Loading report data to Amazon Redshift	17
Configuring Amazon CUR using Billing Conductor	20
Differences between Billing Conductor Amazon CUR and standard Amazon CUR	20
Create pro forma Cost and Usage Reports for a billing group	20
Data dictionary	22
Identity details	22
identity/LinItemId	22
identity/TimeInterval	22
Billing details	23
B	23
I	23
P	24
Line item details	24
A	24
B	24
C	24
L	25
N	26
O	27
P	27
R	27
T	28

U	28
Reservation details	29
A	29
E	30
M	30
N	31
R	32
S	33
T	33
U	34
Pricing details	36
L	36
P	36
R	37
T	37
U	37
Product details	37
A	37
C	38
D	38
E	40
F	41
G	42
I	43
L	45
M	46
N	47
O	47
P	48
R	50
S	51
T	53
U	55
V	55
W	56
Resource tags details	56
resourceTags/user:creator	57
resourceTags/user:name	57
resourceTags/user:owner	57
resourceTags/user:purpose	57
Savings Plans details	57
A	57
E	57
I	58
N	58
O	58
P	59
R	59
S	60
T	61
U	61
Cost Categories details	61
costCategory/Project	61
costCategory/Team	62
costCategory/Environment	62
Discount details	62
discount/BundledDiscount	62

discount/TotalDiscount	62
Use cases	63
Understanding Savings Plans	63
Savings Plans line items	63
Understanding reservations	65
Understanding your reservation line items	66
Understanding your amortized reservation data	69
Monitoring your size flexible reservations for Amazon EC2	71
Monitoring your On-Demand capacity reservations	72
Understanding your data transfer costs	74
All data transfer costs	74
Amazon EC2 data transfer costs	74
VPC peering data transfer costs	75
Legacy reports	76
Detailed Billing Reports	76
Migrating From DBR to Amazon CUR	76
Comparing benefits of the Cost and Usage Reports (Amazon CUR)	77
Key differences between the Detailed Billing Report and the Cost and Usage Reports	78
Reporting on advanced charge types	80
Understanding unused reservation costs	81
Scenario 1: RI usage is 100%	81
Scenario 2: Partial RI usage	82
Scenario 3: Capacity reservation	84
Scenario 4: Size flexible reservations	85
Monthly Report	86
Monthly cost allocation	87
Amazon Usage Report	87
Troubleshooting	88
There are no report files in the Amazon S3 bucket	88
One of my report data partitions is empty	88
My Cost and Usage Report data doesn't match the data in other Billing and Cost Management features	89
I want to backfill data because I changed the settings of my report	89
My report file folder in Amazon S3 is in an unnamed folder	89
I can't select the option to include resource IDs on my report	90
My Cost and Usage Report queries for Amazon Athena don't work on Amazon Redshift, or my Amazon Redshift queries don't work on Amazon Athena	90
The columns included in my report have changed from a previous month	90
The queries or tables based on my report don't work because the columns in my report have changed	91
I need help querying my report	91
I can't find the billing data for my Amazon EC2 Dedicated Host	91
I don't understand the billing data for my Amazon EC2 Elastic IP addresses	91
I use consolidated billing and I don't understand the difference between unblended and blended rates or costs	92
Some line items in my report have a blended rate or blended cost of 0	92
I don't understand how All Upfront Reserved Instances are amortized in my report	92
Contacting customer support	93
Security and permissions	94
Quotas and restrictions	95
Cost and Usage Reports	95
Document history	96
Amazon glossary	97

What are Amazon Cost and Usage Reports?

The Amazon Cost and Usage Reports (Amazon CUR) contains the most comprehensive set of cost and usage data available. You can use Cost and Usage Reports to publish your Amazon billing reports to an Amazon Simple Storage Service (Amazon S3) bucket that you own. You can receive reports that break down your costs by the hour, day, or month, by product or product resource, or by tags that you define yourself. Amazon updates the report in your bucket once a day in comma-separated value (CSV) format. You can view the reports using spreadsheet software such as Microsoft Excel or Apache OpenOffice Calc, or access them from an application using the Amazon S3 API.

Amazon Cost and Usage Reports tracks your Amazon usage and provides estimated charges associated with your account. Each report contains line items for each unique combination of Amazon products, usage type, and operation that you use in your Amazon account. You can customize the Amazon Cost and Usage Reports to aggregate the information either by the hour, day, or month.

Amazon Cost and Usage Reports can do the following:

- Deliver report files to your Amazon S3 bucket
- Update the report up to three times a day
- Create, retrieve, and delete your reports using the Amazon CUR API Reference

How Cost and Usage Reports work

After you create a Cost and Usage Report, Amazon sends your report to the Amazon S3 bucket that you specify. Amazon updates your report at least once a day until your charges are finalized.

Your report files consist of a .csv file or a collection of .csv files and a manifest file. You can choose to configure your report data for integration with Amazon Athena, Amazon Redshift, or Amazon QuickSight.

Report timeline

After you create your report, it can take up to 24 hours for Amazon to deliver the first report to your Amazon S3 bucket.

After delivery starts, Amazon updates the report files at least once a day. Each report update in a given month is cumulative, so each version of the report includes all of the billing data for the month to date. The report updates that you receive throughout the month are estimates. The charges are subject to change as you continue to use your Amazon services.

Note

Different Amazon services provide your usage-based billing information at different times, so you might notice updates to a certain hour or day come in at different times.

Amazon builds on previous reports until the end of the billing period. Amazon finalizes your report's usage charges after issuing an invoice at the end of the month. After the end of the report billing period, Amazon generates a new report for the next month with none of the information from the previous report.

After your report is finalized, Amazon might update the report if Amazon applies refunds, credits, or Amazon Support fees to your usage for the month. Because Developer, Business, and Enterprise Support are calculated based on final usage charges, those are reflected on the sixth or seventh of the month for the prior month's Cost and Usage Report. Amazon applies credits or refunds based on the terms of your agreement or contract with Amazon.

Report files

Your report is a .csv file or a collection of .csv files stored in an Amazon S3 bucket. The number of files that your report generates depends on your selection for report versioning and the size of your report.

When you create a report, you can choose to create new report versions or overwrite the existing report version with every update. If you choose to create new report versions, then your report generates more files with every update.

The size of an individual report can grow to more than a gigabyte and might exceed the capacity of desktop spreadsheet applications to display every line. If a report is larger than most applications can handle (around 1 million rows), then Amazon splits the report into multiple files that are stored in the same folder in the Amazon S3 bucket.

Amazon also generates refunds into separate files. Amazon issues refunds after the closing of a monthly bill.

For more information on report files, file-naming conventions, and versioning, see [Understanding your report versions \(p. 7\)](#).

Report columns

Each report includes several columns with details about your Amazon costs and usage. The columns that Amazon includes in your report depend on your usage during the month.

Every report includes columns with the **identity/**, **bill/**, and **lineltem/** prefixes. All other columns are included only if your monthly Amazon usage generates data to populate those columns. For example, your report includes **savingsPlan/** columns only if you used Savings Plans during that month.

To learn more about the columns in your report, see the [Data dictionary \(p. 22\)](#).

Using your report

You can download your report from the Amazon S3 console, query the report using Amazon Athena, or upload the report into Amazon Redshift or Amazon QuickSight.

- For more information about creating an Amazon S3 bucket and using Athena to query your data, see [Querying Cost and Usage Reports using Amazon Athena \(p. 13\)](#).
- For more information about uploading to Amazon Redshift, see [Loading report data to Amazon Redshift \(p. 17\)](#).
- For more information about uploading to Amazon QuickSight, see [Loading report data to Amazon QuickSight \(p. 17\)](#).

Creating Cost and Usage Reports

You can use the **Cost & Usage Reports** page in the Billing and Cost Management console to create Cost and Usage Reports. In the following section, you'll find information on how to get started with Cost and Usage Reports.

Topics

- [Setting up an Amazon S3 bucket for Cost and Usage Reports \(p. 3\)](#)
- [Creating Cost and Usage Reports \(p. 4\)](#)

Setting up an Amazon S3 bucket for Cost and Usage Reports

To receive billing reports, you must have an Amazon S3 bucket in your Amazon account to receive and store your reports. When setting up a Cost and Usage Report in the billing console you can select an existing Amazon S3 bucket you own or setup up a new one. In either case, you'll be asked to review and confirm the application of the following default bucket policy. Editing this policy in the Amazon S3 console or changing the bucket owner after you've setup a Cost and Usage Report will prevent Amazon from being able to deliver your reports. Storing the billing reports data in your Amazon S3 bucket is billed at standard Amazon S3 rates. For more information, see [Quotas and restrictions \(p. 95\)](#).

The following policy is applied to every bucket during the setup of a Cost and Usage Report:

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "billingreports.amazonaws.com"
      },
      "Action": [
        "s3:GetBucketAcl",
        "s3:GetBucketPolicy"
      ],
      "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
      "Condition": {
        "StringEquals": {
          "aws:SourceArn": "arn:aws:cur:us-east-1:${AccountId}:definition/*",
          "aws:SourceAccount": "${AccountId}"
        }
      }
    },
    {
      "Sid": "Stmt1335892526596",
      "Effect": "Allow",
      "Principal": {
        "Service": "billingreports.amazonaws.com"
      },
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*",
      "Condition": {
        "StringEquals": {
          "aws:SourceArn": "arn:aws:cur:us-east-1:${AccountId}:definition/*",
          "aws:SourceAccount": "${AccountId}"
        }
      }
    }
  ]
}
```



```
}  
  }  
]  
}
```

This default policy helps ensure that the Cost and Usage Report data can be read by the bucket owner and confirms the bucket is owned by the account which set-up the Cost and Usage Report. Specifically:

- Every time a Cost and Usage Report is delivered, Amazon first confirms whether the bucket is still owned by the account which setup the report. If the bucket ownership has changed, the report will not be delivered. This helps to ensure the security of the account's billing data. This bucket policy allows Amazon ("Effect": "Allow") to check which account owns the bucket ("Action": ["s3:GetBucketAcl", "s3:GetBucketPolicy"]).
- To deliver reports to your bucket, Amazon needs write permissions for that bucket. To do this, the bucket policy grants ("Effect": "Allow") the Amazon Cost and Usage Reports service ("Service": "billingreports.amazonaws.com") permission to deliver ("Action": "s3:PutObject") reports to the bucket you own ("Resource": "arn:aws:s3::DOC-EXAMPLE-BUCKET/*").

This bucket policy does not give Amazon permissions to read or delete any objects in your bucket, including the Cost and Usage Reports after they've been delivered.

- To ensure the bucket owner can access these reports, Amazon further applies a `BucketOwnerFullControl` ACL to the reports when delivering them. By default, Amazon S3 objects such as these reports can only be read by the user or service principal who wrote them. To provide you or the bucket owner with permission to read the reports Amazon must apply the `BucketOwnerFullControl` ACL. The ACL grants the bucket owner `Permission.FullControl` for these reports.

If you see an **Invalid bucket** error in your billing console for Cost and Usage Report, you should verify that this policy and bucket ownership haven't changed after report setup.

Creating Cost and Usage Reports

You can use the **Cost & Usage Reports** page of the Billing and Cost Management console to create Cost and Usage Reports. You can create up to 10 reports for an individual Amazon account.

Note

It can take up to 24 hours for Amazon to start delivering reports to your Amazon S3 bucket. After delivery starts, Amazon updates the Amazon Cost and Usage Reports files at least once a day.

To create Cost and Usage Reports

1. Sign in to the Billing and Cost Management console at <https://console.amazonaws.cn/billing/home#/>
2. On the navigation pane, choose **Cost & Usage Reports**.
3. Choose **Create report**.
4. For **Report name**, enter a name for your report.
5. For **Additional report details**, select **Include resource IDs** to include the IDs of each individual resource in the report.

Note

Including resource IDs will create individual line items for each of your resources. This can increase the size of your Cost and Usage Reports files significantly, based on your Amazon usage.

6. For **Data refresh settings**, select whether you want the Amazon Cost and Usage Reports to refresh if Amazon applies refunds, credits, or support fees to your account after finalizing your bill. When a report refreshes, a new report is uploaded to Amazon S3.
7. Choose **Next**.
8. For **S3 bucket**, choose **Configure**.
9. In the **Configure S3 Bucket** dialog box, do one of the following:
 - Select an existing bucket from the drop down list and choose **Next**.
 - Enter a bucket name and the Region where you want to create a new bucket and choose **Next**.
10. Review the bucket policy, and select **I have confirmed that this policy is correct** and choose **Save**.
11. For **Report path prefix**, enter the report path prefix that you want prepended to the name of your report.
12. For **Time granularity**, choose one of the following:
 - **Hourly** if you want the line items in the report to be aggregated by the hour.
 - **Daily** if you want the line items in the report to be aggregated by the day.
 - **Monthly** if you want the line items in the report to be aggregated by month.
13. For **Report versioning**, choose whether you want each version of the report to overwrite the previous version of the report or to be delivered in addition to the previous versions.

Overwriting reports can save on Amazon S3 storage costs. Delivering new report versions can improve auditability of billing data over time.
14. For **Enable report data integration for**, select whether you want to enable your Cost and Usage Reports to integrate with Amazon Athena, Amazon Redshift, or Amazon QuickSight. The report is compressed in the following formats:
 - **Athena**: parquet format
 - **Amazon Redshift or Amazon QuickSight**: .gz compression
15. Choose **Next**.
16. After you have reviewed the settings for your report, choose **Review and Complete**.

You can always return to the Cost & Usage Reports section of your Billing and Cost Management console to see when your reports were last updated.

Managing your Cost and Usage Reports

In the following section, you'll find information about managing your Cost and Usage Reports (Amazon CUR).

Topics

- [Viewing your report \(p. 6\)](#)
- [Understanding your report versions \(p. 7\)](#)
- [Editing your Cost and Usage Reports configuration \(p. 11\)](#)
- [Using Cost and Usage Reports for Amazon Organizations \(p. 12\)](#)

Viewing your report

To view information about your Cost and Usage Report, use the Billing and Cost Management console. To view your report's files, you can use the Amazon S3 console.

Use the following procedures to find your report and report files.

To view your report details and files

1. Sign in to the Billing and Cost Management console at <https://console.amazonaws.cn/billing/home#/>
2. On the navigation pane, choose **Cost & Usage Reports**.
3. From your list of reports, choose the name of the report that you want to view.
4. On the **Report Details** page, you can view the report's settings.
5. To view the report's files, note the **Report path prefix** on the **Report Details** page.
6. Choose the bucket name listed under **Amazon S3 bucket**. The link opens this bucket in the Amazon S3 console.
7. From the list of objects in the bucket, choose the folder named with the first part of the **Report path prefix** that you noted in step 5. For example, if your **Report path prefix** is **example-report-prefix/example-report-name**, then choose the folder named **example-report-prefix**.
8. From the list of objects in the folder, choose the folder named with the second part of the **Report path prefix** that you noted in step 5. For example, if your **Report path prefix** is **example-report-prefix/example-report-name**, then choose the folder named **example-report-name**. This folder contains your report files.

Viewing the latest report version

Amazon updates your Cost and Usage Report at least once a day until your charges are finalized. When you create a report, you can choose to create new report versions or overwrite the existing report version with every update.

If you configured your report to create new report versions with every update, then use the **assemblyId** in the manifest file to find the latest report files.

To view your latest report files when you have multiple report versions

1. Sign in to the Billing and Cost Management console at <https://console.amazonaws.cn/billing/home#/>
2. On the navigation pane, choose **Cost & Usage Reports**.
3. From your list of reports, choose the name of the report that you want to view.
4. On the **Report Details** page, note the **Report path prefix**.
5. Choose the bucket name listed under Amazon S3 bucket. The link opens this bucket in the Amazon S3 console.
6. From the list of objects in the bucket, choose the folder named with the first part of the **Report path prefix** that you noted in step 4. For example, if your **Report path prefix** is **example-report-prefix/example-report-name**, then choose the folder named **example-report-prefix**.
7. From the list of objects in the folder, choose the folder named with the second part of the **Report path prefix** that you noted in step 4. For example, if your **Report path prefix** is **example-report-prefix/example-report-name**, then choose the folder named **example-report-name**.
8. Open the folder named with the latest billing period (in the YYYYMMDD-YYYYMMDD format).
9. Open the **example-report-name-Manifest.json** file.
10. At the top of the manifest file, note the **assemblyId**. The **assemblyId** value corresponds to the name of the folder with the latest report files.
11. Return to the Amazon S3 console page where you're viewing the folder named with the latest billing period.
12. Open the folder named with the **assemblyId** value that you noted in step 10. For example, if the **assemblyId** value is **20210129T123456Z**, then open the folder named **20210129T123456Z/**. This folder contains your latest report files.

Viewing your finalized report

After issuing your invoice at the end of the month, Amazon finalizes your report's usage charges. To determine if a line item on your report is final, review the **bill/InvoiceId** column. If the line item is final, then the **bill/InvoiceId** column is populated with an Amazon invoice ID. If the line item is not yet final, then the **bill/InvoiceId** column is populated with the word "Estimated".

To determine if your entire report is finalized, review the **bill/BillType** column. If the report is final, then the **bill/BillType** column is populated with the word "Anniversary". If the report is not yet final, then the **bill/BillType** column is populated with the word "Estimated".

Note

After your report is finalized, Amazon might update the report if Amazon applies refunds, credits, or support fees to your usage for the month. Because Developer, Business, and Enterprise Support are calculated based on final usage charges, those are reflected on the sixth or seventh of the month for the prior month's report. Amazon applies credits or refunds based on the terms of your agreement or contract with Amazon.

Understanding your report versions

Amazon updates your Cost and Usage Report at least once a day until your charges are finalized. When you create a report, you can choose to create new report versions or overwrite the existing report version with every update.

Your report files include a .csv file or a collection of .csv files and the manifest file. Your report can also include any additional files that support your data's integration with Amazon Athena, Amazon Redshift, or Amazon QuickSight.

The following sections describe the file organization and naming conventions based on the report versioning that you choose.

Cost and Usage Reports delivery timeline

During the report period, Amazon delivers a new report and a new manifest file each time Amazon updates the report. Amazon builds on previous reports until the end of the billing period. After the end of the report billing period, Amazon generates a new report with none of the information from the previous report.

Creating new Cost and Usage Report versions

When you choose to keep your previous Cost and Usage Reports, your Amazon CUR uses the following Amazon S3 organization and naming conventions.

```
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<assemblyId>/<example-report-name>-<file-number>.csv.<zip|gz>
```

- `report-prefix` = The prefix that you assign to the report.
- `report-name` = The name that you assign to the report.
- `yyyymmdd-yyyymmdd` = The range of dates that the report covers. Reports are finalized at the end of the date range.
- `assemblyId` = An ID that Amazon creates each time that the report is updated.
- `file-number` = If the update includes a large file, Amazon might split it into multiple files. The `file-number` tracks the different files in an update.
- `csv` = The format of the report files.
- `zip` or `gz` = The type of compression applied to the report files.

For example, your report could be delivered as a collection of the following files.

```
<example-report-prefix>/<example-report-name>/20160101-20160131/<123456789>/<example-report-name>-<1>.csv.<zip>  
<example-report-prefix>/<example-report-name>/20160101-20160131/<123456789>/<example-report-name>-<2>.csv.<zip>  
<example-report-prefix>/<example-report-name>/20160101-20160131/<123456789>/<example-report-name>-<3>.csv.<zip>  
<example-report-prefix>/<example-report-name>/20160101-20160131/<123456789>/<example-report-name>-Manifest.json  
<example-report-prefix>/<example-report-name>/20160101-20160131/<example-report-name>-Manifest.json
```

Amazon delivers all reports in a report date range to the same `report-prefix/report-name/yyyymmdd-yyyymmdd` folder. Amazon gives each report a unique ID and delivers it to the `assemblyId` subfolder in the date range folder. If the report is too large for a single file, the report is split into multiple files and delivered to the same `assemblyId` folder.

For more information on manifesting files when you keep a previous report, see [Cost and Usage Reports manifest files \(p. 9\)](#)

Overwriting previous Cost and Usage Reports

When you choose to overwrite your previous Cost and Usage Reports, your Amazon CUR uses the following Amazon S3 organization and naming conventions.

```
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>-<file-number>.csv.<zip|gz>
```

- `report-prefix` = The prefix that you assign to the report.
- `report-name` = The name that you assign to the report.
- `yyyymmdd-yyyymmdd` = The range of dates that the report covers. Amazon finalizes reports at the end of the date range.
-
- `file-number` = If the update includes a large file, Amazon might split it into multiple files. The `file-number` tracks the different files in an update.
- `csv` = The format of the report files.
- `zip` or `gz` = The type of compression applied to the report files.

For example, your report could be delivered as a collection of the following files.

```
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>-<1>.csv.<zip>  
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>-<2>.csv.<zip><example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>-<3>.csv.<zip>  
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>-Manifest.json
```

Athena specifications

If you chose Athena support when you created your Amazon CUR, the file naming conventions are the same as when you choose to overwrite your Amazon CUR except for the format and compression. Athena Amazon CUR files use `.parquet` instead. For example, your report could be delivered as a collection of the following files.

```
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>.parquet  
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/  
<cost_and_usage_data_status>  
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>-Manifest.json  
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>-create-table.sql  
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/crawler-cfn.yml
```

Amazon CloudFormation specifications

In addition to the Amazon CUR files, Amazon also delivers an Amazon CloudFormation template that you can use to set up an Amazon CloudFormation stack that enables you to query Amazon S3 data using Athena. If you don't want to use the Amazon CloudFormation template, you can use the provided SQL to create your own Athena tables. For more information, see [Querying Cost and Usage Reports using Amazon Athena](#) (p. 13).

Cost and Usage Reports manifest files

When Amazon updates Amazon CUR, Amazon also creates and delivers manifest files that you can use for Amazon Athena, Amazon Redshift, or Amazon QuickSight.

Manifest files use the naming conventions, and lists the following:

- All of the detail columns that are included in the report to date
- A list of report files if the report was split into multiple files
- The time period covered by the report, and other information.

```
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>-Manifest.json  
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<assemblyId>/<example-report-name>-Manifest.json  
<example-report-prefix>/<example-report-name>/<example-report-name>/year=2018/month=12/<example-report-name>-Manifest.json
```

Creating new Cost and Usage Report versions

When you keep the previous Cost and Usage Reports, the manifest file is delivered to both the date range folder and the `assemblyId` folder. Each time Amazon creates a new Amazon CUR for a date range, it overwrites the manifest file stored in the date range folder with an updated manifest file. Amazon delivers the same updated manifest file to the `assemblyId` folder along with the files for that update. Manifest files in the `assemblyId` folder aren't overwritten.

Overwriting the previous Cost and Usage Reports

When you overwrite the previous Amazon CUR, the manifest file is delivered to the `month=mm` folder. The manifest file is overwritten along with the report files.

Amazon Redshift specifications

If you chose the option for Amazon Redshift support in your Amazon CUR, Amazon also creates and delivers a file with the SQL commands that you need to upload your report into Amazon Redshift. You can open the SQL file with a regular text editor. The SQL file uses the following naming convention.

```
<example-report-prefix>/<example-report-name>/yyyymmdd-yyyymmdd/<assemblyId>/<example-report-name>-RedshiftCommands.sql
```

If you use the commands in the `RedshiftCommands` file, you don't need to open the `RedshiftManifest` file.

Important

The manifest file determines which report files the `copy` command in the `RedshiftCommands` file uploads. Deleting or removing the manifest file breaks the `copy` command in the `RedshiftCommands` file.

Amazon Athena specifications

If you chose the option for Amazon Athena support in your Amazon CUR, Amazon also creates and delivers multiple files to help set up all of the resources that you need. Amazon delivers a Amazon CloudFormation template, a SQL file with the SQL to create your Athena table manually, and a file with the SQL to check your Amazon CUR refresh status. These files use the following naming conventions.

```
<example-report-prefix>/<example-report-name>/<example-report-name>/yyyymmdd-yyyymmdd/crawler-cfn.yml  
<example-report-prefix>/<example-report-name>/<example-report-name>/yyyymmdd-yyyymmdd/<example-report-name>-create-table.sql
```

```
<example-report-prefix>/<example-report-name>/<example-report-name>/yyyymmdd-yyyymmdd/  
<cost_and_usage_data_status>
```

Editing your Cost and Usage Reports configuration

You can use the **Cost & Usage Reports** page in the Billing and Cost Management console to edit Cost and Usage Reports.

Note

Report names can't be edited. If you chose **Overwrite** for **Report versioning**, you're unable to edit the report name, whether the report includes resource IDs, time granularity, or the report versioning. If you delete a report set to **Overwrite** and create a new report with the same name, Amazon S3 bucket, and path prefix, your data could corrupt and become inaccurate.

To edit Cost and Usage Reports

1. Sign in to the Billing and Cost Management console at <https://console.amazonaws.cn/billing/home#/>
2. On the navigation pane, choose **Cost & Usage Reports**.
3. Select the report that you want to edit and choose **Edit report**.
4. (Versioned reports only) For **Additional report details**, select **include resource IDs** to add IDs for individual resources.
5. Select the **Data refresh settings** if you want Amazon Cost and Usage Reports to refresh when Amazon applies refunds, credits, or support fees to your account after your bill is finalized.

When a report refreshes, a new report is uploaded to Amazon S3.

6. Choose **Next**.
7. For **S3 bucket**, enter the name of the Amazon S3 bucket where you want the reports delivered.
8. Choose **Verify**.

Note

The bucket must have appropriate permissions to be valid. For more information on adding permissions to the bucket, see [Setting Bucket and Object Access Permissions](#) in the *Amazon Simple Storage Service User Guide*.

9. For **Report path prefix**, enter the report path prefix that you want prepended to the name of your report.
10. (Versioned reports only) For **Time granularity**, choose one of the following:
 - **Hourly**: If you want the line items in the report to be aggregated by the hour.
 - **Daily**: If you want the line items in the report to be aggregated by the day.
 - **Monthly** if you want the line items in the report to be aggregated by month.
11. (Versioned reports only) For **Report versioning**, choose whether you want each version of the report to overwrite the previous version of the report, or to be delivered in addition to the previous versions.
12. For **Enable report data integration for**, select whether you want to enable your Amazon CUR to integrate with Amazon Athena, Amazon Redshift, or Amazon QuickSight. The report is compressed in the following formats:
 - **Athena**: Parquet format
 - **Amazon Redshift or Amazon QuickSight**: .gz compression
13. Choose **Save**.

Using Cost and Usage Reports for Amazon Organizations

In Amazon Organizations, both management accounts and member accounts can create Cost and Usage Reports. The IAM policies that allow or restrict the ability to create a report are the same for both types of accounts.

Note

The account that creates the Cost and Usage Report must also own the Amazon S3 bucket that Amazon sends the reports to. Avoid configuring a Cost and Usage Report with a bucket owned by another account. For more information the Amazon S3 bucket setup requirements, see [Setting up an Amazon S3 bucket for Cost and Usage Reports \(p. 3\)](#).

Managing Cost and Usage Reports as a member account

If you have permissions to create a Cost and Usage Report for a member account within an organization, you can create a report for only the member account's cost and usage data. The member account receives reports for its cost and usage during the time that the account has been a member of its current organization.

For example, say a member account leaves organization A and joins organization B on the 15th of the month. Then, the member account creates a report. Because the member account created a report after joining organization B, the member account's report for the month includes billing data for only the time that the account has been a member of organization B.

After a member account joins a new organization, the member account's cost and usage are recorded in the new organization's reports. This is the same outcome for a management account that converts to a member account and joins a new organization.

When a member account leaves an organization or converts into a standalone account, the member account can still access the previous reports as long as they have permissions to the Amazon S3 bucket where the previous reports are stored.

Managing Cost and Usage Reports as a management account

If you're an administrator of an Amazon Organizations management account and you don't want member accounts to create a report, you can apply a service control policy (SCP) that prevents member accounts from creating reports. The SCP can prevent member accounts from creating new reports, but it doesn't delete reports created previously.

Note

SCPs apply only to member accounts. To prevent a management account from creating a report, modify the IAM policies assigned to the IAM users in the management account.

For more information on consolidated billing, see [Consolidated billing for Amazon Organizations](#) in the *Amazon Billing User Guide*.

Querying Cost and Usage Reports using Amazon Athena

Amazon Athena is a serverless query service that you can use to analyze the data from your Amazon Cost and Usage Reports (Amazon CUR) in Amazon Simple Storage Service (Amazon S3) using standard SQL. This helps you avoid having to create your own data warehouse solutions to query Amazon CUR data.

We strongly recommend that you create both a new Amazon S3 bucket and a new Amazon CUR report to use with Athena. Amazon CUR supports only the Apache Parquet compression format for Athena and automatically overwrites previous reports that are stored in your S3 bucket.

This section outlines how to use Athena with Cost and Usage Reports. For a full description of the Athena service, see the [Amazon Athena User Guide](#).

Topics

- [Setting up Athena using Amazon CloudFormation templates \(p. 13\)](#)
- [Setting up Athena manually \(p. 14\)](#)
- [Running Amazon Athena queries \(p. 16\)](#)
- [Loading report data to other resources \(p. 17\)](#)

For a demonstration of querying reports using Athena, see the following video.

[Analyze Cost and Usage Reports using Amazon Athena](#)

Setting up Athena using Amazon CloudFormation templates

Important

Amazon CloudFormation doesn't support cross-Region resources. If you plan to use an Amazon CloudFormation template, you must create all resources in the same Amazon Region. The Region must support the following services:

- Amazon Lambda
- Amazon Simple Storage Service (Amazon S3)
- Amazon Glue
- Amazon Athena

To streamline and automate integration of your Cost and Usage Reports with Athena, Amazon provides an Amazon CloudFormation template with several key resources along with the reports that you set up for Athena integration. The Amazon CloudFormation template includes an Amazon Glue crawler, an Amazon Glue database, and an Amazon Lambda event.

The Athena integration setup process using Amazon CloudFormation removes any Amazon S3 events that your bucket might already have. This can negatively affect any existing event-based processes that you have for an existing Amazon CUR report. We strongly recommend that you create both a new Amazon S3 bucket and a new Amazon CUR report to use with Athena.

Before you can use a CloudFormation template to automate Athena integration, make sure that you do the following:

- Create a new Amazon S3 bucket for your reports. For more information, see [Creating a bucket](#) in the *Amazon S3 User Guide*.
- [Create a new report \(p. 4\)](#) to use with Athena. During the setup process, for **Enable report data integration for**, choose **Athena**.
- Wait for the first report to be delivered to your Amazon S3 bucket. It can take up to 24 hours for Amazon to deliver your first report.

To use the Athena Amazon CloudFormation template

1. Open the Amazon S3 console at <https://console.amazonaws.cn/s3/>.
2. From the list of buckets, choose the bucket where you chose to receive your Amazon CUR report.
3. Choose your report path prefix (*your-report-path-prefix/*). Then, choose your report name (*your-report-name/*).
4. Choose the `.yaml` template file.
5. Choose **Object actions**, and then choose **Download**.
6. Open the Amazon CloudFormation console at <https://console.amazonaws.cn/cloudformation>.
7. If you have never used Amazon CloudFormation before, choose **Create New Stack**. Otherwise, choose **Create Stack**.
8. Under **Prepare template**, choose **Template is ready**.
9. Under **Template source**, choose **Upload a template file**.
10. Choose **Choose file**.
11. Choose the downloaded `.yaml` template, and then choose **Open**.
12. Choose **Next**.
13. For **Stack name**, enter a name for your template and choose **Next**.
14. Choose **Next**.
15. At the bottom of the page, select **I acknowledge that Amazon CloudFormation might create IAM resources**.

This template creates the following resources:

- Three IAM roles
 - An Amazon Glue database
 - An Amazon Glue crawler
 - Two Lambda functions
 - An Amazon S3 notification
16. Choose **Create stack**.

Setting up Athena manually

We strongly recommend that you use the Amazon CloudFormation template to create your table instead of creating it yourself. The provided SQL query creates a table that covers only a single month of data, but the Amazon CloudFormation template creates a table that can include multiple months and that updates automatically. For more information on how to set up the Amazon CloudFormation template, see [the section called "Setting up Athena with CloudFormation" \(p. 13\)](#).

If you choose not to use the Amazon CloudFormation template to set up your Athena table, manually follow the steps below. You need to create a table before you can run SQL queries on your Amazon

CUR data. You will need to do this step at least once a month and the table only includes data from the current Amazon CUR.

As part of the table creation process, Amazon transforms the Amazon CUR column names. For more information about the transformation process, see [the section called "Column names" \(p. 17\)](#).

- [Creating an Athena table \(p. 15\)](#)
- [Creating a Cost and Usage Reports status table \(p. 15\)](#)
- [Uploading your report partitions \(p. 16\)](#)

Creating an Athena table

Amazon includes the SQL that you need to run to create this table in your Amazon CUR bucket.

To create your Athena table

1. Sign in to the Amazon Web Services Management Console and open the Amazon S3 console at <https://console.amazonaws.cn/s3/>.
2. From the list of buckets, choose the bucket where you chose to receive your Cost and Usage Reports.
3. Navigate the path *your-report-prefix-your-report-name-path-to-report*.

The exact path depends on whether your Amazon CUR is set to overwrite previous versions. For more information, see [Cost and Usage Reports delivery timeline \(p. 8\)](#).

4. Open the file *my-report-name-create-table.sql*.
5. Copy the SQL from the file, starting with `CREATE` and ending with `LOCATION 's3://your-report-prefix/your-report-name/the-rest-of-the=path'`. Take note of the first line, as you need the database name and table to create the Athena database.
6. Open the Athena console at <https://console.amazonaws.cn/athena/>.
7. In the **New query 1** query pane, paste the following SQL. For *<database name>*.*<table name>*, use the database and table name from the first line of the SQL that you copied.

```
CREATE DATABASE <database name>
```

8. Choose **Run query**.
9. In the dropdown menu, choose the database that you just created.
10. In the **New query 1** query pane, paste the rest of the SQL from the SQL file.
11. Choose **Run query**.

After you create your table, you need to load your partitions before you can run a query. For more information, see [Uploading your report partitions \(p. 16\)](#).

Creating a Cost and Usage Reports status table

Amazon refreshes your Amazon CUR multiple times a day. There isn't a way for Athena to tell when Amazon is in the process of refreshing your report, which can lead to query results with a combination of old and new data. To mitigate this, create a table to track whether Amazon is refreshing your Cost and Usage Reports and query that table to see if Amazon is refreshing your data. You only need to create this table once. After that, Amazon keeps the table up to date.

To create your refresh table

1. Open the Athena console at <https://console.amazonaws.cn/athena/>.

2. In the **New query 1** query pane, paste the following SQL.

```
CREATE EXTERNAL TABLE IF NOT EXISTS cost_and_usage_data_status(  
  status STRING)  
ROW FORMAT SERDE  
  'org.apache.hadoop.hive ql.io.parquet.serde.ParquetHiveSerDe'  
WITH SERDEPROPERTIES (  
  'serialization.format' = '1'  
)  
LOCATION 's3://{S3_Bucket_Name}/{Report_Key}/cost_and_usage_data_status/'
```

3. Choose **Run query**.

To check whether Amazon is refreshing your data, use the Athena console to run the following SQL query.

```
select status from cost_and_usage_data_status
```

Uploading your report partitions

To query your Cost and Usage Reports data, you need to upload the data into your Athena table. You must do this for each new Amazon CUR report that Amazon delivers to you.

To upload your latest partitions

1. Open the Athena console at <https://console.amazonaws.cn/athena/>.
2. Choose the ... next to your table.
3. Choose **Load Partitions**.

If you don't upload your partitions, Athena returns either no results or an error message that indicates missing data.

Running Amazon Athena queries

To run Athena queries on your data, first use the Athena console to check whether Amazon is refreshing your data and then run your query on the Athena console. When you run your SQL, make sure that the correct database is selected from the dropdown list. You can use the following SQL to check the status.

```
select status from cost_and_usage_data_status
```

The two possible results are `READY` and `UPDATING`. If the status is `READY`, then you can query your Athena database. If the status is `UPDATING`, then Athena might return incomplete results.

After you've confirmed that Amazon is refreshing your data, you can run your own queries. For example, the following query shows year-to-date costs by service for each month in the example database called `mycostandusage_parquet`.

```
SELECT line_item_product_code,  
sum(line_item_blended_cost) AS cost, month  
FROM mycostandusage_parquet  
WHERE year='2018'  
GROUP BY line_item_product_code, month
```

```
HAVING sum(line_item_blended_cost) > 0  
ORDER BY line_item_product_code;
```

Column names

Athena column name restrictions are different from the Cost and Usage Reports column name restrictions. This means that when your Amazon CUR data is uploaded into an Athena table, the column names change. Amazon makes the following changes:

- An underscore is added in front of uppercase letters
- Uppercase letters are replaced with lowercase letters
- Any non-alphanumeric characters are replaced with an underscore
- Duplicate underscores are removed
- Any leading and trailing underscores are removed
- If the column name is longer than the allowed length of column names, underscores are removed from left to right

Note

After applying these rules, some of the resource tag columns will have duplicate names. Amazon merges columns when there are more than one columns with the same name.

For example, the column name `ExampleColumnName : Example Column Name Continued` becomes `example_column_name_example_column_name_continued`.

Loading report data to other resources

You can upload Cost and Usage Reports to Amazon Redshift and Amazon QuickSight to analyze your Amazon cost and usage.

Topics

- [Loading report data to Amazon QuickSight \(p. 17\)](#)
- [Loading report data to Amazon Redshift \(p. 17\)](#)

Loading report data to Amazon QuickSight

You can upload your Cost and Usage Reports into Amazon QuickSight.

For more information about uploading to Amazon QuickSight, see [Creating a Data Set Using Amazon S3 Files](#) in the *Amazon QuickSight User Guide*.

Loading report data to Amazon Redshift

This section shows how you can upload Amazon CUR to Amazon Redshift to analyze your Amazon costs and usage.

Important

Amazon Redshift columns aren't case sensitive and has stricter character limitations than user-defined tags. To prevent conflicts between Amazon Redshift and user-defined tags, Amazon replaces your tags with the tags `userTag0`, `userTag1`, `userTag2`, etc. After you create an Amazon Redshift table and upload your report into it, you can create an Amazon Redshift table

that maps the Amazon-defined tags to your user-defined tags. The tag table allows you to look up your original tags.

For example, if you have the tags `OWNER` and `Owner`, Amazon Redshift doesn't allow you to create a table with two columns named "owner". Instead, you create a report table with the columns `userTag0` and `userTag1` instead of `OWNER` and `Owner`, and then create a table with the columns `remappedUserTag` and `userTag`. The `remappedUserTag` column stores the Amazon-defined tags `userTag0` and `userTag1`, and the `userTag` column stores your original tags, `OWNER` and `Owner`.

Amazon provides the commands to create your Amazon Redshift table, upload your report, create your tag table, and insert all of the tag rows into your tag table. The commands are provided to you in the `RedshiftCommands.sql` file that is stored alongside your manifest file in S3, and in the **Redshift file Helper file** in the Billing and Cost Management console. Amazon also provides a `RedshiftManifest` file, which controls which report the commands in the `RedshiftCommand` file uploads. Deleting or removing the `RedshiftManifest` file breaks the copy command in the `RedshiftCommands` file.

To find the `RedshiftCommands.sql` file in the Billing and Cost Management console

1. Sign in to the Billing and Cost Management console at <https://console.amazonaws.cn/billing/home#/>
2. In the navigation pane, choose **Cost & Usage Reports**.
3. Choose the report that you want to upload to Amazon Redshift.
4. Next to **You have enabled viewing reports in the following service(s)**, choose **Amazon Redshift**.
5. Copy the commands from the dialog box and paste them into your SQL client.

The following procedure assumes familiarity with databases and Amazon Redshift.

To upload an Cost and Usage Reports to Amazon Redshift

1. Create an Amazon Redshift cluster.

For more information, see [Creating a Cluster](#) in the *Amazon Redshift Cluster Management Guide*.
2. Sign in to the Amazon Web Services Management Console and open the Amazon S3 console at <https://console.amazonaws.cn/s3/>.
3. Navigate to the Amazon S3 location where you store your Amazon CUR.
4. Open the `RedshiftCommands.sql` file.

The file contains customized commands to create an Amazon Redshift table, upload the Amazon CUR from Amazon S3, and create a tag table that allows user-defined tags to be imported into Amazon Redshift.

5. In the copy command, replace `<AWS_ROLE>` with the ARN of an IAM role that has permissions to access the Amazon S3 bucket where you store your Amazon CUR.
6. Replace `<S3_BUCKET_REGION>` with the Region your Amazon S3 bucket is in. For example, `us-east-1`.
7. Use a SQL client to connect to the cluster.

For more information, see [Accessing Amazon Redshift Clusters and Databases](#) in the *Amazon Redshift Cluster Management Guide*.

8. Copy the SQL commands from the `RedshiftCommands.sql` file to your SQL client in the following order:
 - create table - This command creates an Amazon Redshift table with a schema customized to match your report.
 - copy - This command uses the provided IAM role to upload the Amazon CUR files from S3 to Amazon Redshift.

- create tag table - This command creates a table that allows you to map Amazon-defined tags to your user-defined tags.
 - insert - These commands insert the user-defined tags into the tag table.
9. After you have copied all of the data from your Amazon CUR into Amazon Redshift, you can query the data using SQL. For more information about querying data in Amazon Redshift, see [Amazon Redshift SQL](#) in the *Amazon Redshift Database Developer Guide*.

Note

The number of columns in Cost and Usage Reports can change from month to month, such as when a new cost allocation tag is created or a service adds a new product attribute. We recommend that you copy the data from your Amazon CUR into a new table every month, and then copy the columns that interest you into a separate month-by-month table.

Configuring Cost and Usage Reports using Billing Conductor pro forma billing data

You can create pro forma Amazon Cost and Usage Reports (Amazon CUR) for each billing group that you create in Billing Conductor. The pro forma Amazon CUR has the same file format, granularity, and columns as the standard Amazon CUR. Pro forma contains the most comprehensive set of cost and usage data available for a given period of time. For more information about Billing Conductor, see the [Billing Conductor User Guide](#).

Topics

- [Understanding the differences between Billing Conductor Amazon CUR and standard Amazon CUR \(p. 20\)](#)
- [Create pro forma Cost and Usage Reports for a billing group \(p. 20\)](#)

Understanding the differences between Billing Conductor Amazon CUR and standard Amazon CUR

There are a few differences between the standard Cost and Usage Reports and pro forma Amazon CUR created using the Billing Conductor configuration.

- The standard Amazon CUR computes the cost and usage for each account in your consolidated billing family. A pro forma Amazon CUR per billing group only includes the accounts in the billing group at the time of computation.
- The standard Amazon CUR populates the invoice column once and invoice is generated by Amazon. A pro forma Amazon CUR doesn't populate the invoice column. Currently, no invoice is generated, or issued by Amazon based on pro forma billing data.

Create pro forma Cost and Usage Reports for a billing group

Use the following steps to generate a pro forma Amazon CUR for a billing group.

To create pro forma Cost and Usage Reports for a billing group

1. Sign in to the Billing and Cost Management console at <https://console.amazonaws.cn/billing/home#/>
2. On the navigation pane, choose **Cost & Usage Reports**
3. On the top right of the **Report table**, choose **Settings**.

4. Turn on the **Pro forma** data view.
5. Choose **Enable**.
6. Choose **Create report**.
7. For **Report name**, enter a name for your report.
8. For **Data view**, choose **pro forma**.
9. Choose a billing group.
10. For **Additional report details**, choose **Include resource IDs** to include the IDs of each individual resources in the report.
11. For **Data refresh settings**, choose if you want Cost and Usage Reports to refresh if Amazon applies refunds, credits, or support fees to your account after finalizing your bill. When a report refreshes, a new report is uploaded to Amazon S3.
12. Choose **Next**.
13. For **S3 bucket**, choose **Configure**.
14. In the **Configure S3 Bucket** dialog box, do one of the following:
 - Choose an existing bucket from the drop down list and choose **Next**.
 - Enter a bucket name and the Amazon Region where you want to create a new bucket and choose **Next**.
15. Review the bucket policy, select **I have confirmed that this policy is correct**, and choose **Save**.
16. For **Report path prefix**, enter the report path prefix that you want prepended to the name of your report.

This step is optional for Amazon Redshift or Amazon QuickSight, but required for Amazon Athena. If you don't specify a prefix, the default prefix is the name that you specified for the report in step 7, and the date range for the report in the following format: `/report-name/date-range/`

17. For **Time granularity**, choose one of the following:
 - **Hourly** if you want the line items in the report to be aggregated by the hour.
 - **Daily** if you want the line items in the report to be aggregated by the day.
18. For **Report versioning**, choose whether you want each version of the report to overwrite the previous version of the report, or to be delivered in addition to the previous versions.

Overwriting reports can save on Amazon S3 storage costs. Delivering new report versions can improve auditability of billing data over time.
19. For **Enable report data integration for**, choose whether you want to upload your Cost and Usage Reports to Amazon Athena, Amazon Redshift, or Amazon QuickSight. The report is compressed in the following formats:
 - **Athena**: parquet format
 - **Amazon Redshift or Amazon QuickSight**: .gz compression
20. Choose **Next**.
21. After you have reviewed the settings for your report, choose **Review and Complete**.

Data dictionary

Cost and Usage Reports contain details about your usage. The following sections list and describe a subset of columns that you see in your report, and the corresponding definitions.

To download the full list of the columns that can appear in Amazon Cost and Usage Reports (Amazon CUR) and the services that the columns apply to, download [Column_Attribute_Service.zip](#). This comma-separated values (CSV) list includes **Identity**, **Bill**, **LineItem**, **Reservation**, **Pricing**, and **Product** columns.

Every Cost and Usage Report includes **Identity**, **Bill**, and **LineItem** columns. All other columns are included in your report only if your monthly Amazon usage generates data to populate those columns.

Topics

- [Identity details \(p. 22\)](#)
- [Billing details \(p. 23\)](#)
- [Line item details \(p. 24\)](#)
- [Reservation details \(p. 29\)](#)
- [Pricing details \(p. 36\)](#)
- [Product details \(p. 37\)](#)
- [Resource tags details \(p. 56\)](#)
- [Savings Plans details \(p. 57\)](#)
- [Cost Categories details \(p. 61\)](#)
- [Discount details \(p. 62\)](#)

Identity details

Columns under the **identity** header in Amazon Cost and Usage Reports are static fields that appear in all Cost and Usage Reports.

You can use the identity line items in the report to find specific line items that have been split across multiple Amazon CUR files. These include the following columns:

identity/LineItemId

- **Description:** An ID that identifies every line item in a single given version of the Amazon CUR. The line item ID isn't consistent between different Cost and Usage Reports and can't be used to identify the same line item across different reports.
- **Example:** A report created for November 29 can be large enough to require multiple files. The **LineItemId** is consistent between the November 29 Amazon CUR files, but doesn't match the **LineItemId** for the same resource in the November 30 report.

identity/TimeInterval

- **Description:** The time interval that this line item applies to, in the following format: `YYYY-MM-DDTHH:mm:ssZ/YYYY-MM-DDTHH:mm:ssZ`. The time interval is in UTC and can be either daily or hourly, depending on the granularity of the report.
- **Example:** The **TimeInterval** `2017-11-01T00:00:00Z/2017-12-01T00:00:00Z` includes the entire month of November 2017.

Billing details

Columns under the **bill** header in Amazon Cost and Usage Reports are static fields that appear in all Cost and Usage Reports. You can use the billing line items in the report to find details about the specific bill covered by the report, such as the charge type and the beginning and end of the billing period. This includes the following columns:

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | VWXYZ

B

bill/BillingEntity

The Amazon seller that your account is with. Possible values are the following:

- **China (Beijing)** – Amazon Web Services China (Beijing) Region operated by Sinnet
- **China (Ningxia)** – Amazon Web Services China (Ningxia) Region operated by NWCD

bill/BillingPeriodEndDate

The end date of the billing period that is covered by this report, in UTC. The format is `YYYY-MM-DDTHH:mm:ssZ`.

bill/BillingPeriodStartDate

The start date of the billing period that is covered by this report, in UTC. The format is `YYYY-MM-DDTHH:mm:ssZ`.

bill/BillType

The type of bill that this report covers. There are three bill types:

- **Anniversary** – Line items for services that you used during the month
- **Purchase** – Line items for upfront service fees
- **Refund** – Line items for refunds

I

bill/InvoiceId

The ID associated with a specific line item. Until the report is final, the **InvoiceId** is blank.

bill/InvoicingEntity

The Amazon entity that issues the invoice. Possible values include:

- **Amazon Web Services** - The entity that issues invoices to customer globally, where applicable.
- **Amazon Web Services EMEA SARL** - The entity that issues invoices to customers based in Europe, the Middle East, and Africa (excluding South Africa).
- **Amazon Internet Services Pvt. Ltd** - The entity that issues invoices to customers based in India.

P

bill/PayerAccountId

The account ID of the paying account. For an organization in Amazon Organizations, this is the account ID of the management account.

Line item details

Columns under the **lineitem** header in Amazon Cost and Usage Reports are static fields that appear in all Cost and Usage Reports. They cover all of the cost and usage information for your usage. This includes the following columns:

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | VWXYZ

A

lineitem/AvailabilityZone

The Availability Zone that hosts this line item. For example, `us-east-1a` or `us-east-1b`.

B

lineitem/BlendedCost

The `BlendedRate` multiplied by the `UsageAmount`.

Note

BlendedCost is blank for line items that have a **LineItemType** of **Discount**. Discounts are calculated using only the unblended cost of a member account, aggregated by member account and SKU. As a result, **BlendedCost** is not available for discounts.

lineitem/BlendedRate

The `BlendedRate` is the average cost incurred for each SKU across an organization.

For example, the Amazon S3 blended rates are the total cost of storage divided by the amount of data stored per month. For accounts with RIs, the blended rates are calculated as the average costs of the RIs and the On-Demand Instances.

Blended rates are calculated at the management account level, and used to allocate costs to each member account. For more information, see [Blended Rates and Costs](#) in the *Amazon Billing User Guide*.

C

lineitem/CurrencyCode

The currency that this line item is shown in. All Amazon customers are billed in US dollars by default. To change your billing currency, see [Changing which currency you use to pay your bill](#) in the *Amazon Billing User Guide*.

L

lineItem/LegalEntity

The provider of your Amazon services. Possible values are the following:

- **Amazon Web Services** – The entity that sells Amazon services.
- **Amazon Internet Services Pvt. Ltd** – The local Indian entity that acts as a reseller for Amazon services in India.

lineItem/LineItemDescription

The description of the line item type. For example, the description of a usage line item summarizes what type of usage you incurred during a specific time period.

For size-flexible RIs, the description corresponds to the RI the benefit was applied to. For example, if a line item corresponds to a `t2.micro` and a `t2.small` RI was applied to the usage, the line **item/description** displays `t2.small`.

The description for a usage line item with an RI discount contains the pricing plan covered by the line item.

lineItem/LineItemType

The type of charge covered by this line item. Possible types are the following:

- **BundledDiscount** – A usage-based discount that provides free or discounted usage of a service or feature based on the usage of another service or feature.
- **Credit** – Any credits that Amazon applied to your bill. See the **Description** column for details. Amazon might update reports after they have been finalized if Amazon applies a credit to your account for the month after finalizing your bill.
- **DiscountedUsage** – The rate for any instances for which you had Reserved Instance (RI) benefits.
- **Fee** – Any upfront annual fee that you paid for subscriptions. For example, the upfront fee that you paid for an **All Upfront RI** or a **Partial Upfront RI**.
- **Refund** – The negative charges that Amazon refunded money for. Check the **Description** column for details. Amazon might update reports after they have been finalized if Amazon applies a refund to your account for the month after finalizing your bill.
- **RIFee** – The monthly recurring fee for subscriptions. For example, the recurring fee for **Partial Upfront RIs**, **No Upfront RIs**, and **All Upfronts** that you pay every month. Although the RIFee might be \$0 for all upfront reservations, this line is still populated for those reservation types to provide other columns such as **reservation/AmortizedUpfrontFeeForBillingPeriod** and **reservation/ReservationARN**.
- **Tax** – Any taxes that Amazon applied to your bill. For example, VAT or US sales tax.
- **Usage** – Any usage that is charged at On-Demand Instance rates.
- **SavingsPlanUpfrontFee** – Any one-time upfront fee from your purchase of an All Upfront or Partial Upfront Savings Plan.
- **SavingsPlanRecurringFee** – Any recurring hourly charges that correspond with your No Upfront or Partial Upfront Savings Plan. The Savings Plan recurring fee is initially added to your bill on the day that you purchase a No Upfront or Partial Upfront Savings Plan. After the initial purchase, Amazon adds the recurring fee to the first day of each billing period thereafter.
- **SavingsPlanCoveredUsage** – Any on-demand cost that is covered by your Savings Plan. Savings Plan covered usage line items are offset by the corresponding Savings Plan negation items.

- **SavingsPlanNegation** – Any offset cost through your Savings Plan benefit that’s associated with the corresponding Savings Plan covered usage item.

For more information and examples of Savings Plans line items, see [Understanding Savings Plans](#) (p. 63).

N

lineItem/NetUnblendedCost

The actual after-discount cost that you're paying for the line item. This column is included in your report only when your account has a discount in the applicable billing period.

lineItem/NetUnblendedRate

The actual after-discount rate that you're paying for the line item. This column is included in your report only when your account has a discount in the applicable billing period.

lineItem/NormalizationFactor

As long as the instance has shared tenancy, Amazon can apply all Regional Linux or Unix Amazon EC2 and Amazon RDS RI discounts to all instance sizes in an instance family and Amazon Region. This also applies to RI discounts for member accounts in an organization. All new and existing Amazon EC2 and Amazon RDS size-flexible RIs are sized according to a normalization factor, based on the instance size. The following table shows the normalization factor that Amazon applies to each instance size.

Normalization factors for Amazon EC2 size-flexible RIs

Instance size	Normalization factor
nano	0.25
micro	0.5
small	1
medium	2
large	4
xlarge	8
2xlarge	16
4xlarge	32
8xlarge	64
10xlarge	80
16xlarge	128
32xlarge	256

lineItem/NormalizedUsageAmount

The amount of usage that you incurred, in normalized units, for size-flexible RIs. The **NormalizedUsageAmount** is equal to **UsageAmount** multiplied by **NormalizationFactor**.

O

lineItem/Operation

The specific Amazon operation covered by this line item. This describes the specific usage of the line item. For example, a value of `RunInstances` indicates the operation of an Amazon EC2 instance.

P

lineItem/ProductCode

The code of the product measured. For example, Amazon EC2 is the product code for Amazon Elastic Compute Cloud.

R

lineItem/ResourceId

(Optional) If you chose to include individual resource IDs in your report, this column contains the ID of the resource that you provisioned. For example, an Amazon S3 storage bucket, an Amazon EC2 compute instance, or an Amazon RDS database can each have a resource ID. This field is blank for usage types that aren't associated with an instantiated host, such as data transfers and API requests, and line item types such as discounts, credits, and taxes. The following table shows a list of resource identifiers for common Amazon services.

Amazon resource identifiers

Amazon service	Resource identifier
Amazon CloudFront	Distribution ID
Amazon CloudSearch	Search domain
Amazon DynamoDB	DynamoDB table
Amazon Elastic Compute Cloud - Amazon EBS	Amazon EBS volume
Amazon Elastic Compute Cloud	Instance ID
Amazon Elastic Compute Cloud - CloudWatch	CloudWatch charges for an instance ID
Amazon EMR	MapReduce cluster
Amazon ElastiCache	Cache cluster
Amazon OpenSearch Service	Search domain
Amazon S3 Glacier	Vault
Amazon Relational Database Service	Database
Amazon Redshift	Amazon Redshift cluster
Amazon Simple Storage Service	Amazon S3 bucket
Amazon Virtual Private Cloud	VPN ID

Amazon service	Resource identifier
Amazon Lambda	Lambda function name

T

lineItem/TaxType

The type of tax that Amazon applied to this line item.

U

lineItem/UnblendedCost

The `UnblendedCost` is the `UnblendedRate` multiplied by the `UsageAmount`.

lineItem/UnblendedRate

In consolidated billing for accounts using Amazon Organizations, the unblended rate is the rate associated with an individual account's service usage.

For Amazon EC2 and Amazon RDS line items that have an RI discount applied to them, the `UnblendedRate` is zero. Line items with an RI discount have a `LineItemType` of `Discounted Usage`.

lineItem/UsageAccountId

The ID of the account that used this line item. For organizations, this can be either the management account or a member account. You can use this field to track costs or usage by account.

lineItem/UsageAmount

The amount of usage that you incurred during the specified time period. For size-flexible Reserved Instances, use the `reservation/TotalReservedUnits` column instead.

Note

Certain subscription charges will have a `UsageAmount` of 0.

lineItem/UsageEndDate

The end date and time for the corresponding line item in UTC, exclusive. The format is `YYYY-MM-DDTHH:mm:ssZ`.

lineItem/UsageStartDate

The start date and time for the line item in UTC, inclusive. The format is `YYYY-MM-DDTHH:mm:ssZ`.

lineItem/UsageType

The usage details of the line item. For example, `USW2-BoxUsage:m2.2xlarge` describes an M2 High Memory Double Extra Large instance in the US West (Oregon) Region.

Reservation details

The reservation columns in Amazon Cost and Usage Reports provide details about reserved resources.

[A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#) | [Y](#) | [Z](#)

A

reservation/AmortizedUpfrontCostForUsage

- **Description:** The initial upfront payment for all upfront RIs and partial upfront RIs amortized for usage time. The value is equal to: $RIAmortizedUpfrontFeeForBillingPeriod * \frac{\text{The normalized usage amount for DiscountedUsage line items}}{\text{The normalized usage amount for the RIFee}}$. Because there are no upfront payments for no upfront RIs, the value for a no upfront RI is 0. We do not provide this value for Dedicated Host reservations at this time. The change will be made in a future update.
- **Line items applicable:** DiscountedUsage
- **Sample values:** 0.05, 0.17, 0.15
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon DynamoDB
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

reservation/AmortizedUpfrontFeeForBillingPeriod

- **Description:** Describes how much of the upfront fee for this reservation is costing you for the billing period. The initial upfront payment for all upfront RIs and partial upfront RIs, amortized over this month. Because there are no upfront fees for no upfront RIs, the value for no upfront RIs is 0. We do not provide this value for Dedicated Host reservations at this time. The change will be made in a future update.
- **Line items applicable:** RIFee
- **Sample values:** 29.15, 200.67, 214.43
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon DynamoDB
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

reservation/AvailabilityZone

- **Description:** The Availability Zone of the resource that is associated with this line item.
- **Line items applicable:** Fee, Refund, RIFee

- **Sample values:** us-east-1, us-east-1b, eu-west-1b, ap-southeast-2a
- **Services:**
 - Amazon EC2

E

reservation/EffectiveCost

- **Description:** The sum of both the upfront and hourly rate of your RI, averaged into an effective hourly rate. `EffectiveCost` is calculated by taking the `amortizedUpfrontCostForUsage` and adding it to the `recurringFeeForUsage`. For more information, see [Amazon EC2 Reserved Instances Pricing](#).
- **Line items applicable:** `DiscountedUsage`
- **Sample values:** 0.23, 0.68, 0.10
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon DynamoDB
 - Amazon ElastiCache
 - Amazon RDS

reservation/EndTime

- **Description:** The end date of the associated RI lease term.
- **Line items applicable:** `RIFee`
- **Sample values:** 2019-05-15T04:23:14.000Z, 2020-02-08T17:32:15.000Z, 2019-07-14T00:00:33.000Z
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

M

reservation/ModificationStatus

- **Description:** Shows whether the RI lease was modified or if it is unaltered.
 - **Original:** The purchased RI was never modified.
 - **System:** The purchased RI was modified using the console or API.
 - **Manual:** The purchased RI was modified using Amazon Web Services Support assistance.
 - **ManualWithData:** The purchased RI was modified using Amazon Web Services Support assistance, and Amazon calculated estimates for the RI.
- **Line items applicable:** `RIFee`
- **Sample values:** `Original`, `System`, `Manual`, `ManualWithData`
- **Services:**

- Amazon EC2
- OpenSearch Service
- Amazon DynamoDB
- Amazon Redshift
- Amazon ElastiCache
- Amazon RDS

N

reservation/NetAmortizedUpfrontCostForUsage

The initial upfront payment for All Upfront RIs and Partial Upfront RIs amortized for usage time, if applicable. This column is included in your report only when your account has a discount in the applicable billing period.

reservation/NetAmortizedUpfrontFeeForBillingPeriod

The cost of the reservation's upfront fee for the billing period. This column is included in your report only when your account has a discount in the applicable billing period.

reservation/NetEffectiveCost

The sum of both the upfront fee and the hourly rate of your RI, averaged into an effective hourly rate. This column is included in your report only when your account has a discount in the applicable billing period.

reservation/NetRecurringFeeForUsage

The after-discount cost of the recurring usage fee. This column is included in your report only when your account has a discount in the applicable billing period.

reservation/NetUnusedAmortizedUpfrontFeeForBillingPeriod

The net unused amortized upfront fee for the billing period. This column is included in your report only when your account has a discount in the applicable billing period.

reservation/NetUnusedRecurringFee

The recurring fees associated with unused reservation hours for Partial Upfront and No Upfront RIs after discounts. This column is included in your report only when your account has a discount in the applicable billing period.

reservation/NetUpfrontValue

The upfront value of the RI with discounts applied. This column is included in your report only when your account has a discount in the applicable billing period.

reservation/NormalizedUnitsPerReservation

- **Description:** The number of normalized units for each instance of a reservation subscription.

- **Line items applicable:** RIFee
- **Sample values:** 1316, 54.5, 319
- **Services:**
 - Amazon RDS

reservation/NumberOfReservations

- **Description:** The number of reservations that are covered by this subscription. For example, one RI subscription might have four associated RI reservations.
- **Line items applicable:** Fee, RIFee, Refund, Credit
- **Sample values:** 5, 50, 500
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon DynamoDB
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

R

reservation/RecurringFeeForUsage

- **Description:** The recurring fee amortized for usage time, for partial upfront RIs and no upfront RIs. The value is equal to: $\frac{\text{The unblended cost of the RIFee} * \text{The sum of the normalized usage amount of Usage line items}}{\text{The normalized usage amount of the RIFee for size flexible Reserved Instances}}$. Because all upfront RIs don't have recurring fee payments greater than 0, the value for all upfront RIs is 0.
- **Line items applicable:** DiscountedUsage
- **Sample values:** 0.139, 0.729, 0.018
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon DynamoDB
 - Amazon ElastiCache
 - Amazon RDS

reservation/ReservationARN

- **Description:** The Amazon Resource Name (ARN) of the RI that this line item benefited from. This is also called the "RI Lease ID". This is a unique identifier of this particular Amazon Reserved Instance. The value string also contains the Amazon service name and the Region where the RI was purchased.
- **Line items applicable:** Fee, RIFee, DiscountedUsage, Refund, Credit
- **Sample values:** arn:aws:ec2:us-east-1:074108124787:reserved-instances/1d3fbc13-f181-4c40-9dd6-12b345678de9, arn:aws:ec2:us-east-1:499958231354:reserved-instances/be41234c3-b5c0-403e-a80c-1cfd12345678

- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon DynamoDB
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

S

reservation/StartTime

- **Description:** The start date of the term of the associated Reserved Instance.
- **Line items applicable:** RIFee
- **Sample values:** 2018-07-29T02:56:10.000Z, 2017-08-21T15:58:47.000Z, 2019-02-01T22:01:34.000Z
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

reservation/SubscriptionId

- **Description:** A unique identifier that maps a line item with the associated offer. We recommend you use the RI ARN as your identifier of an Amazon Reserved Instance, but both can be used.
- **Line items applicable:** Fee, RIFee, DiscountedUsage, Refund, Credit, and Usage
- **Sample values:** 123456789, 111122222, 333344444
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon DynamoDB
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

T

reservation/TotalReservedNormalizedUnits

- **Description:** The total number of reserved normalized units for all instances for a reservation subscription. Amazon computes total normalized units by multiplying the `reservation/NormalizedUnitsPerReservation` with `reservation/NumberOfReservations`.
- **Line items applicable:** DiscountedUsage

- **Sample values:** 40320, 3647.99, 17928.77
- **Services:**
 - Amazon EC2
 - Amazon RDS

reservation/TotalReservedUnits

- **Description:** TotalReservedUnits populates for both Fee and RIFee line items with distinct values.
 - Fee line items: The total number of units reserved, for the total quantity of leases purchased in your subscription for the entire term.

This is calculated by multiplying the NumberOfReservations with UnitsPerReservation. For example, 5 RIs x 744 hours per month x 12 months = 44,640.

- RIFee line items (monthly recurring costs): The total number of available units in your subscription, such as the total number of Amazon EC2 hours in a specific RI subscription.

For example, 5 RIs x 744 hours = 3,720.

- **Line items applicable:** Fee, RIFee, Refund, Credit
- **Sample values:** 26208, 98.19, 15796
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon DynamoDB
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

U

reservation/UnitsPerReservation

- **Description:** UnitsPerReservation populates for both Fee and RIFee line items with distinct values.
 - Fee line items: The total number of units reserved for the subscription, such as the total number of RI hours purchased for the term of the subscription.

For example 744 hours per month x 12 months = 8,928 total hours/units.

- RIFee line items (monthly recurring costs): The total number of available units in your subscription, such as the total number of Amazon EC2 hours in a specific RI subscription.

For example, 1 unit x 744 hours = 744.

- **Line items applicable:** Fee, RIFee, Refund, Credit
- **Sample values:** 334.0, 486.72, 18455
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon DynamoDB
 - Amazon Redshift
 - Amazon ElastiCache

- Amazon RDS

reservation/UnusedAmortizedUpfrontFeeForBillingPeriod

- **Description:** The amortized-upfront-fee-for-billing-period-column amortized portion of the initial upfront fee for all upfront RIs and partial upfront RIs. Because there are no upfront payments for no upfront RIs, the value for no upfront RIs is 0. We do not provide this value for Dedicated Host reservations at this time. The change will be made in a future update.
- **Line items applicable:** RIFee
- **Sample values:** 6.05, 1.97, 0.17
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

reservation/UnusedNormalizedUnitQuantity

- **Description:** The number of unused normalized units for a size-flexible Regional RI that you didn't use during this billing period.
- **Line items applicable:** RIFee
- **Sample values:** 25.00, 3.50, 274.33
- **Services:**
 - Amazon RDS

reservation/UnusedQuantity

- **Description:** The number of RI hours that you didn't use during this billing period.
- **Line items applicable:** RIFee line item
- **Sample values:** 209.65110408, 191.00000000, 176.00000000
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

reservation/UnusedRecurringFee

- **Description:** The recurring fees associated with your unused reservation hours for partial upfront and no upfront RIs. Because all upfront RIs don't have recurring fees greater than 0, the value for All Upfront RIs is 0.
- **Line items applicable:** RIFee
- **Sample values:** 0.02971114, 0.19190000, 1.37280000
- **Services:**

- Amazon EC2
- OpenSearch Service
- Amazon Redshift
- Amazon ElastiCache
- Amazon RDS

reservation/UpfrontValue

- **Description:** The upfront price paid for your Amazon Reserved Instance. For no upfront RIs, this value is 0.
- **Line items applicable:** RIFee
- **Sample values:** 150.00, 1000.00, 2000.00
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon Redshift
 - Amazon ElastiCache
 - Amazon RDS

Pricing details

The pricing columns in Amazon Cost and Usage Reports contain the prices for a line item. The pricing columns are based off of the Amazon Price List Service API. Amazon Price List Service API doesn't include Free Tier pricing, Spot Instances, products in Amazon Web Services Marketplace, upfront annual subscription fees (Fee), and monthly recurring fees (RIFee). The columns include but are not limited to the following:

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

L

pricing/LeaseContractLength

The length of time that your RI is reserved for.

P

pricing/publicOnDemandCost

The total cost for the line item based on public On-Demand Instance rates. If you have SKUs with multiple On-Demand public costs, the equivalent cost for the highest tier is displayed. For example, services offering free-tiers or tiered pricing.

pricing/publicOnDemandRate

The public On-Demand Instance rate in this billing period for the specific line item of usage. If you have SKUs with multiple On-Demand public rates, the equivalent rate for the highest tier is displayed. For example, services offering free-tiers or tiered pricing.

pricing/PurchaseOption

How you chose to pay for this line item. Valid values are All Upfront, Partial Upfront, and No Upfront.

R

pricing/RateCode

A unique code for a product/ offer/ pricing-tier combination. The product and term combinations can have multiple price dimensions, such as a free tier, low-use tier, and high-use tier.

pricing/RateId

The ID of the rate for a line item.

T

pricing/term

Whether your Amazon usage is Reserved or On-Demand.

U

pricing/unit

The pricing unit that Amazon used for calculating your usage cost. For example, the pricing unit for Amazon EC2 instance usage is in hours.

Product details

The product columns provide metadata about the product that incurred the expense, and the line item. The product columns are dynamic, and their visibility in Cost and Usage Reports depends on the usage of product in the billing period. The pricing columns are based off of the Amazon Price List Service API. Amazon Price List Service API doesn't include free tier pricing, Spot Instances, products in Amazon Marketplace, upfront annual subscription fees (Fee), and monthly recurring fees (RFee).

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | XYZ

A

product/availability

- **Description:** Describes the availability of your various Amazon storage options.
- **Sample values:** 99.99%, 99.5%
- **Services:**
 - Amazon S3 Glacier
 - Amazon S3

- AWS Elemental MediaStore
- Amazon RoboMaker

C

product/capacitystatus

- **Description:** Describes the status of your capacity reservations.
- **Sample values:** UnusedCapacityReservation, AllocatedCapacityReservation, Used
- **Services:**
 - Amazon EC2

product/clockspeed

- **Description:** Describes the operating speed of your Amazon instances.
- **Sample values:** 2.4 GHz, 2.6 GHz
- **Services:**
 - Amazon DocumentDB
 - Amazon EC2
 - Amazon MQ
 - Amazon Neptune
 - Amazon RDS
 - Amazon Database Migration Service

D

product/databaseedition

- **Description:** Describes the database software suitable for different development, deployment scenarios, and specific application purposes.
- **Services:**
 - Amazon RDS

product/dataTransfer

- **Description:** An Amazon data transfer occurs whenever data is moved to from Amazon to the internet, or moved between Amazon instances across their respective Regions or Availability Zones. Interregional and inter availability zone data transfers incur costs, metered per Gigabyte.
- **Services:**
 - Amazon Systems Manager

product/dedicatedEbsThroughput

- **Description:** Describes the dedicated throughput between your instances (for example, Amazon EC2 instances and Amazon EBS volumes), with options between 500 and 10,000 megabits per second

(Mbps) depending on the instance type used. The dedicated throughput minimizes contention between Amazon EBS I/O and other traffic from your EC2 instance, providing the best performance for your Amazon EBS volumes.

- **Sample values:** 200 Mbps, Up to 5000 Mbps
- **Services:**
 - Amazon EC2
 - Amazon Neptune
 - Amazon RDS

product/deploymentoption

- **Description:** Describes where the infrastructure for the environment is located. The deployment models for Amazon cloud are public, on-premise, and hybrid.
- **Sample values:** Multi-AZ, Single-AZ
- **Services:**
 - Amazon MQ
 - Amazon Neptune
 - Amazon RDS

product/description

- **Description:** The description of the specific Amazon service.
- **Services:**
 - Amazon CodePipeline
 - Amazon Device Farm
 - AWS Elemental MediaConvert
 - AWS Elemental MediaStore

product/directconnectlocation

- **Description:** Specifies the location where a private dedicated network connection from the customer to Amazon exists.
- **Sample values:** Equinix DC1 - DC6, Equinix DC10 - DC11, Global Switch Singapore
- **Services:**
 - Amazon Direct Connect

product/directorysize

- **Description:** The space on the disk that is used to store the meta information for the directory or folder.
- **Services:**
 - Amazon Direct Connect

product/directorytype

- **Description:** Specifies if the directory is a file or another directory.

- **Services:**
 - Amazon Direct Connect

product/directorytypedescription

- **Description:** The meaningful name given to the directory.
- **Services:**
 - Amazon Direct Connect

product/disableactivationconfirmationemail

- **Description:** Active or deactivate the ability to send an email to confirm the activation of a service.

product/durability

- **Description:** Describes the durability of objects over a given year.
- **Sample values:** 99.999999999%, N/A, 99.99%
- **Services:**
 - Amazon S3 Glacier
 - Amazon S3
 - AWS Elemental MediaStore

E

product/ebsOptimized

- **Description:** Describes whether your Amazon EC2 instances are Amazon EBS-optimized.
- **Sample values:** Yes, No
- **Services:**
 - Amazon EC2

product/ecu

- **Description:** Describes the EC2 Compute Unit (ECU) that provides the relative measure of the integer processing power of an Amazon EC2 instance.
- **Sample values:** 9, 100, variable
- **Services:**
 - Amazon EC2
 - OpenSearch Service
 - Amazon GameLift
 - Amazon Redshift

product/endpointtype

- **Description:** Describes the characteristics of the remote connection that a device connects to.

For example, `REST` (representational state transfer) endpoints. A `REST` API (or `RESTful` API) is an application programming interface that conforms to the constraints of `REST` architectural style and you can interact with `RESTful` web services.

- **Sample values:** `Ipsec`, `Amazon SQS`, `Amazon Lambda`
- **Services:**
 - Amazon SNS
 - Amazon VPC
 - Storage Gateway
 - S3 Glacier

product/enhancedNetworkingSupported

- **Description:** Describes whether your instance supports enhanced networking. Enhanced networking uses single root I/O virtualization (SR-IOV) to provide high-performance networking capabilities on supported instance types.
- **Sample values:** `Yes`, `No`
- **Services:**
 - Amazon DocumentDB
 - Amazon EC2
 - Amazon Neptune
 - Amazon RDS
 - Amazon Database Migration Service

F

product/filesystemtype

- **Description:** Describes the details of the local or remote storage device, and specifications of the operating system.

product/freeUsageIncluded

- **Description:** Free usage under Amazon Free Tier is calculated each month across all Regions, and automatically applied to your bill. For example, you receive 750 Amazon EC2 Linux Micro Instance hours for free across all of the Regions you use. Not 750 hours per Region.
- **Services:**
 - Amazon Inspector

product/fromLocation

- **Description:** Describes the location where the usage originated from.
- **Sample values:** `External`, `US East (N. Virginia)`, `Global`
- **Services:**
 - Amazon CloudFront
 - Amazon DataTransfer

product/fromRegionCode

- **Description:** Describes the source Region code for the Amazon service. For more information, see [product/regioncode \(p. 51\)](#).
- **Sample values:** ap-northeast-1
- **Services:**
 - Amazon RDS
 - Amazon EC2
 - Amazon VPC
 - Amazon Direct Connect

product/fromLocationType

- **Description:** Describes the location type where the usage originated from.
- **Sample values:** Amazon Region, Amazon Edge Location
- **Services:**
 - Amazon Direct Connect
 - AWS Elemental MediaConnect
 - Amazon CloudFront
 - Amazon Lightsail
 - Amazon Shield

G

product/gpu

- **Description:** Describes the number of GPUs.
- **Sample values:** 16, 32
- **Services:**
 - Amazon SageMaker
 - Amazon EC2

product/gpuMemory

- **Description:** Describes your GPU memory details.
- **Sample values:** 16, 32
- **Services:**
 - Amazon SageMaker
 - Amazon EC2

product/group

- **Description:** A construct of several products that are similar by definition, or grouped together. For example, the Amazon EC2 team can categorize their products into shared instances, dedicated host, and dedicated usage.
- **Services:**

- Amazon Certificate Manager
- Amazon CodeCommit
- Amazon Glue
- Amazon IoT Analytics
- Amazon Lambda

product/groupdescription

- **Description:** A simplified name given to a product group.
- **Services:**
 - Amazon Budgets
 - Amazon Certificate Manager
 - Amazon Lambda
 - Amazon SQS

|

product/instance

- **Description:** An Amazon EC2 instance is a virtual server in Amazon Elastic Compute Cloud (Amazon EC2) for running applications on the Amazon infrastructure. You can choose an AMI provided by Amazon, the user community, or through the Amazon Web Services Marketplace.
- **Sample values:** T3
- **Services:**
 - Amazon EC2

product/instanceFamily

- **Description:** Describes your Amazon EC2 instance family. Amazon EC2 provides you with a large number of options across 10 different instance types, each with one or more size options, organized into distinct instance families optimized for different types of applications.
- **Sample values:** General Purpose, Memory Optimized, Accelerated Computing
- **Services:**
 - Amazon EC2
 - Amazon RDS
 - OpenSearch Service
 - Amazon ElastiCache
 - Amazon EMR

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/instanceType

- **Description:** Describes the instance type, size, and family, which define the CPU, networking, and storage capacity of your instance.
- **Sample values:** t2.small, m4.xlarge, t2.micro, m4.large, t2.large

- **Services:**
 - Amazon EC2
 - Amazon RDS
 - OpenSearch Service
 - Amazon ElastiCache
 - Amazon EMR

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/instanceTypeFamily

- **Description:** The instance family that is associated with the given usage.
- **Sample values:** t2, m4, m3
- **Services:**
 - Amazon DocumentDB
 - Amazon RDS

product/integratingApi

- **Description:** Application integration on Amazon using service like Amazon API Gateway or no-code integration using Amazon AppFlow.

product/integratingService

- **Description:** Application integration on Amazon is a suite of services used to communicate between decoupled components within micro services, distributed systems, and serverless applications. You don't need to refactor your entire architecture. Decoupling applications at any scale can reduce the impact of changes, making it easier to update, and faster to release new features.

product/intelAvxAvailable

- **Description:** Describes whether your process has the Intel Advanced Vector Extension instruction set.
- **Sample values:** Yes, No.
- **Services:**
 - Amazon EC2

product/intelAvx2Available

- **Description:** Describes whether your process has the Intel Advanced Vector Extension instruction set two.
- **Sample values:** Yes, No
- **Services:**
 - Amazon EC2

product/intelTurboAvailable

- **Description:** Describes whether your core is allowed to use Intel Turbo Technology to increase frequency.
- **Sample values:** Yes, No
- **Services:**
 - Amazon EC2

L

product/licenseModel

- **Description:** Describes the license model for your instance.
- **Sample value:** license-included, bring-your-own-license, general-public-license
- **Services:**
 - Amazon AppStream
 - Amazon EC2
 - Amazon MQ
 - Amazon Neptune
 - Amazon RDS

product/location

- **Description:** Describes the Region that your Amazon S3 bucket resides in.
- **Sample values:** Asia Pacific (Mumbai), Asia Pacific (Seoul), Canada (Central), EU (London), US West (Oregon)
- **Services:**
 - Amazon EC2
 - Amazon Certificate Manager
 - Amazon S3
 - Amazon RDS
 - Amazon DynamoDB

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/locationType

- **Description:** Describes the endpoint of your task.
- **Sample values:** Amazon Region, Amazon Edge Location, Other
- **Services:**
 - Amazon EC2
 - Amazon Certificate Manager
 - Amazon S3
 - Amazon RDS
 - Amazon DynamoDB

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/logsDestination

- **Description:** The `AWS::Logs::Destination` resource specifies a CloudWatch logs destination. A destination includes a physical resource (for example, Amazon Kinesis data stream) and you can subscribe the resource to a stream of log events.
- **Sample values:** Amazon Region, Amazon Edge Location, Other
- **Services:**
 - Amazon EC2
 - Amazon Certificate Manager
 - Amazon S3
 - Amazon RDS
 - Amazon DynamoDB

and more. For the full service list, download [Column_Attribute_Service.zip](#).

M

product/maxIOPSBurstPerformance

- **Description:** Describes the max IOPS burst performance of your Amazon EBS volume.
- **Sample value:** 3000 IOPS for volumes <= 1TB
- **Services:**
 - Amazon EC2

product/maxIOPSVolume

- **Description:** Describes maximum input/output per second of your Amazon EBS volume.
- **Sample value:** 16,000 (maxIOPS for a General Purpose SSD (gp2))
- **Services:**
 - Amazon EC2

product/maxThroughputVolume

- **Description:** Describes the max network throughput volume of your Amazon EBS volume.
- **Sample values:** 500 MiB/s, 250 MiB/s, 1000 MiB/s, 40 - 90 MB/sec
- **Services:**
 - Amazon EC2
 - Amazon SageMaker

product/memory

- **Description:** The placeholder electronics for instructions and data a computer needs to respond quickly. Computer bytes indicate the storage units.
- **Services:**
 - Amazon Database Migration Service
 - DynamoDB Accelerator

- Amazon DocumentDB
- Amazon EC2

N

product/networkPerformance

- **Description:** Describes the network throughput of your Amazon EC2 instances.
- **Sample values:** moderate, high, up to 10 GB
- **Services:**
 - Amazon EC2
 - Amazon RDS
 - Amazon ElastiCache
 - Amazon SageMaker
 - Amazon Database Migration Service

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/normalizationSizeFactor

- **Description:** Describes the normalization factor of the instance size.
- **Sample values:** nano - 0.25, micro - 0.5, medium - 2, xlarge - 8, 16xlarge - 128
- **Services:**
 - Amazon DocumentDB
 - Amazon EC2
 - Amazon MQ
 - Amazon Neptune
 - Amazon RDS

O

product/OfferingClass

- **Description:** Describes the type of Reserved Instances. When you purchase a Reserved Instance, you can choose between a Standard or Convertible offering class.
- **Sample values:** Standard, Convertible
- **Services:**
 - Amazon DynamoDB
 - Amazon EC2
 - Amazon ElastiCache
 - OpenSearch Service
 - Amazon RDS
 - Amazon Redshift

product/operatingSystem

- **Description:** Describes the operating system of your Amazon EC2 instance.
- **Sample values:** Amazon Linux, Ubuntu, Windows Server, Oracle Linux, FreeBSD
- **Services:**
 - Amazon AppStream
 - Amazon EC2
 - Amazon GameLift
 - Amazon Lightsail
 - Amazon WorkSpaces
 - Amazon CodeBuild

product/operation

- **Description:** Describes the specific Amazon operation that this line item covers.
- **Sample values:** RunInstances (indicates the operation of an Amazon EC2 instance)
- **Services:**
 - Amazon EC2
 - Amazon S3
 - Amazon RDS
 - Amazon DynamoDB
 - Amazon CloudWatch
 - Amazon Redshift

and more. For the full service list, download [Column_Attribute_Service.zip](#).

P

product/parameterType

- **Description:** Use parameters in Amazon CloudFormation to enter custom values to your template when you create or update a stack. For Example, InstanceTypeParameter. You can use this parameter to specify the Amazon EC2 instance type when you create or update the stack.

product/physicalCores

- **Description:** Describes the number of physical cores an instance provides.
- **Sample values:** 4, 8
- **Services:**
 - Amazon EC2

product/pricingUnit

- **Description:** The smallest billing unit for an Amazon service. For example, 0.01c per API call.
- **Services:**
 - Amazon Directory Service

product/primaryplaceofuse

- **Description:** The primary business or residential street address location where a customer's use of the service primarily occurs.

product/physicalProcessor

- **Description:** Describes the processor on your Amazon EC2 instance.
- **Sample values:** High Frequency Intel Xeon E7-8880 v3 (Haswell), Intel Xeon E5-2670, AMD EPYC 7571
- **Services:**
 - Amazon DocumentDB
 - Amazon EC2
 - Amazon Neptune
 - Amazon RDS
 - Amazon Database Migration Service

product/processorArchitecture

- **Description:** Describes your processor architecture.
- **Sample values:** 32-bit, 64-bit
- **Services:**
 - Amazon DocumentDB
 - Amazon EC2
 - Amazon Neptune
 - Amazon RDS
 - Amazon Database Migration Service

product/processorFeatures

- **Description:** Describes the processor features of your instances.
- **Sample values:** Intel AVX, Intel AVX2, Intel AVX512, Intel Turbo
- **Services:**
 - Amazon Database Migration Service
 - Amazon DocumentDB
 - Amazon EC2
 - Amazon Neptune
 - Amazon RDS

product/ProductFamily

- **Description:** The category for the type of product.
- **Sample values:** Alarm, Amazon Budgets, Stopped Instance, Storage Snapshot, Compute
- **Services:**
 - Amazon EC2
 - Amazon Certificate Manager

- Amazon S3
- Amazon RDS
- Amazon DynamoDB

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/ProductName

- **Description:** The full name of the Amazon service. Use this column to filter Amazon usage by Amazon service.
- **Sample values:** Amazon Backup, Amazon Config, Amazon Registrar, Amazon Elastic File System, Amazon Elastic Compute Cloud

product/productSchemaDescription

- **Description:** A blueprint of how your product is constructed. This contains the various attributes that make up your product.

product/provisioned

- **Description:** Indicates whether Amazon EBS usage was related to provisioned Amazon EBS storage.
- **Sample values:** Yes, No
- **Services:**
 - Amazon EC2
 - Amazon MQ

product/PurchaseOption

- **Description:** Describes the available purchasing models for an Amazon service. For example, Amazon provides four main Amazon EC2 instance purchasing options: On-Demand, Reserved Instances, Spot Instances, with the added option of Dedicated Hosts.

product/purchaseterm

- **Description:** In Amazon EC2, this specifies a commitment to a consistent instance configuration. This includes instance type and Region for a period of 1 to 3 years.

R

product/region

- **Description:** The geographical area that hosts your Amazon services. Use this field to analyze spend across a particular Region.
- **Sample values:** eu-west-3, us-west-1, us-east-1, ap-northeast-2, sa-east-1
- **Services:**
 - Amazon EC2
 - Amazon Certificate Manager

- Amazon S3
- Amazon RDS
- Amazon DynamoDB

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/regioncode

- **Description:** A Region is a physical location around the world where data centers are clustered. Amazon calls each group of logical data centers an Availability Zone (AZ). Each Amazon Region consists of multiple, isolates, and physically separate AZs within a geographical area. The Region code attribute has the same name as an Amazon Region, and specifies where the Amazon service is available.
- **Sample values:** `us-west-2`, `us-east-1`, `ap-southeast-2`
- **Services:**
 - Amazon SageMaker

product/replicationType

- **Description:** Specifies that the service is free to use. For example, Amazon Server Migration Service is free to use, and you only pay for the storage resources used during the migration process.
- **Sample values:** `Free`
- **Services:**
 - Amazon Application Migration Service

product/resourceAssessment

- **Description:** A process that collects, stores, and manages evidence. You can use this to assess risk and compliance with industry standards and regulations.
- **Sample values:** `All assessment`
- **Services:**
 - Amazon Audit Manager

S

product/servicecode

- **Description:** This identifies the specific Amazon service to the customer as a unique short abbreviation.
- **Sample values:** `Amazon EC2`, `Amazon KMS`
- **Services:**
 - Amazon Budgets
 - Amazon Backup
 - Amazon Certificate Manager
 - Amazon Cloud Map
 - Amazon CloudTrail

product/servicename

- **Description:** A simplified description about the Amazon service.
- **Services:**
 - Amazon EC2 Budgets
 - Amazon ECR
 - Amazon ECS
 - Amazon EFS
 - Amazon Elastic Inference
 - Amazon EKS

product/singleOrDualPass

- **Description:** Terms used to decide the type of encoding that happens for videos. In one-pass encoding, the encoding is done in the first pass itself. For 2-pass encoding, the file is analyzed thoroughly in the first pass, and an intermediate file is created. In the second pass the encoder finds the intermediate file and allocates bits. The actual encoding takes place in the second pass.
- **Services:**
 - AWS Elemental MediaConvert

product/sku

- **Description:** A unique code for a product. The SKU is created by combining the `ProductCode`, `UsageType`, and `Operation`. For size-flexible RIs, the SKU uses the instance that was used. For example, if you used a `t2.micro` instance and Amazon applied a `t2.small` RI discount to the usage, the line item SKU is created with the `t2.micro`.
- **Sample values:** `FFNT87MQSCR328W6, VBYCEU494XUAHCA7`
- **Services:**
 - Amazon EC2
 - Amazon Certificate Manager
 - Amazon S3
 - Amazon RDS
 - Amazon DynamoDB

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/storage

- **Description:** Describes the disk storage attached to your instance.
- **Sample values:** `60GB, True, EBS Only, 1 x 900 NVMe SSD, 1 x 150 NVMe SSD`
- **Services:**
 - Amazon EC2
 - Amazon RDS
 - Amazon Redshift
 - OpenSearch Service
 - Amazon WorkSpaces

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/storageclass

- **Description:** Describes the storage class of your Amazon S3 bucket.
- **Sample values:** Archive, General Purpose, Infrequent Access, Intelligent-Tiering, Non-Critical Data
- **Services:**
 - AWS Elemental MediaStore
 - Amazon Storage Gateway
 - Amazon Cloud Directory
 - Amazon EFS
 - Amazon MQ
 - Amazon S3

product/storagemedia

- **Description:** A storage medium is any technology, including device and material used to place, keep, and retrieve electronic data.
- **Services:**
 - Amazon Database Migration Service
 - Amazon CloudWatch
 - Amazon DocumentDB
 - Amazon EC2
 - Amazon ES

product/storagetype

- **Description:** Describes how and where the information is stored by a computer. This might be internal or external to a computer, server, or computing device.
- **Sample values:** Amazon S3, SSD, SSD-backed
- **Services:**
 - Amazon Backup
 - Amazon ECR

T

product/tenancy

- **Description:** The type of tenancy allowed on the Amazon EC2 instance.
- **Sample values:** Dedicated, Reserved, Shared, NA, Host
- **Services:**
 - Amazon EC2
 - Amazon ECS

product/tier

- **Description:** With Amazon, you can get volume based discounts and savings as your usage increases. For services like Amazon S3, pricing is tiered. This means the more you use, the less you pay per GB. Amazon provides options to acquire services that assist your business needs.
- **Services:**
 - AWS Elemental MediaConvert

product/toLocation

- **Description:** Describes the location usage destination.
- **Sample values:** External, US East (N. Virginia)
- **Services:**
 - Amazon CloudFront
 - Amazon Data Transfer

product/toLocationType

- **Description:** Describes the destination location of the service usage.
- **Sample values:** Amazon Region, Amazon Edge Location
- **Services:**
 - Amazon Direct Connect
 - AWS Elemental MediaConnect
 - Amazon Shield
 - Amazon CloudFront
 - Amazon Lightsail

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/toRegionCode

- **Description:** Describes the source Region code for the Amazon service. For more information, see [product/regioncode](#) (p. 51).
- **Sample values:** eu-west-1
- **Services:**
 - Amazon RDS
 - Amazon EC2
 - Amazon VPC
 - Amazon Direct Connect

product/transcodingResult

- **Description:** The output of decoding an encoded video source to an intermediate uncompressed format, and re-encoding it into the target format.
- **Services:**
 - AWS Elemental MediaConvert

product/trialProduct

- **Description:** Describes if Amazon CloudHSM allows free hours.
- **Services:**
 - Amazon CloudHSM

U

product/upfrontCommitment

- **Description:** Describes if any usage commitment is required for Amazon CloudHSM. You will be charged an hourly fee for each hour (or partial hour) that an HSM is provisioned to a Amazon CloudHSM cluster. A cluster with no HSMs is not billed, and you aren't billed for automatic storage of encrypted backups. For more information, see [Amazon CloudHSM Pricing](#).

Network data transfers to and from your HSMs are charged separately. For more information, see [Amazon EC2 Pricing](#).

- **Services:**
 - Amazon CloudHSM

product/usagetype

- **Description:** Describes the usage details of the line item.
- **Sample values:** EU-BoxUsage:c5d.9xlarge, EU-BoxUsage:m4.16xlarge, SAE1-InstanceUsage:db.t2.medium, USW2-AW-SW-19, SAE1-BoxUsage:c4.large,
- **Services:**
 - Amazon EC2
 - Amazon Certificate Manager
 - Amazon S3
 - Amazon RDS
 - Amazon DynamoDB

and more. For the full service list, download [Column_Attribute_Service.zip](#).

V

product/vcpu

- **Description:** Describes the number of threads concurrently running on a single CPU core. Amazon EC2 instances support multithreading, which enables multiple threads to run concurrently on a single CPU core. Each thread is represented as a virtual CPU (vCPU) on the instance.
- **Sample values:** 8, 16, 36, 72, 128
- **Services:**
 - Amazon EC2
 - Amazon RDS
 - Amazon Redshift
 - OpenSearch Service
 - Amazon ElastiCache

and more. For the full service list, download [Column_Attribute_Service.zip](#).

product/videoCodec

- **Description:** A software or hardware that compresses and decompresses digital video. In the context of video compression, codec is a blending of encoder and decoder. A device that only compresses is typically called an encoder, and one that only decompresses is a decoder.
- **Services:**
 - AWS Elemental MediaConvert;

product/videoFrameRate

- **Description:** A video frame rate (shown as frames per second (FPS)) is the frequency rate which consecutive images (frames) are captured or displayed by video cameras, computer graphics, and motion capture systems.
- **Services:**
 - AWS Elemental MediaConvert;

product/volumeType

- **Description:** Describes your Amazon EBS volume types.
- **Sample values:** Standard, General Purpose, General Purpose-Aurora, Amazon Glacier, Amazon SimpleDB - Standard,
- **Services:**
 - Amazon EC2
 - Amazon S3
 - Amazon RDS
 - Amazon DynamoDB
 - Amazon S3 Glacier

and more. For the full service list, download [Column_Attribute_Service.zip](#).

W

product/workforceType

- **Description:** The segmentation of the employed or unemployed labour pool. For example, Full Time Employees (FTE), or Temporary.
- **Services:**
 - Amazon SageMaker

Resource tags details

You can use the resource columns in Amazon Cost and Usage Reports to find information about the specific resources covered by a line item. These columns include user-defined cost allocation tags. Examples include the following:

resourceTags/user:creator

Use a **user:Creator** tag to track which user created a resource.

resourceTags/user:name

Use a **user:Name** tag to track which resources are associated with a specific user.

resourceTags/user:owner

Use a **user:Owner** tag to track which user owns a resource.

resourceTags/user:purpose

Use a **user:Purpose** tag to track why a resource was created.

Savings Plans details

The SavingsPlan columns in Amazon Cost and Usage Reports provide details about the ComputeSavingsPlans. SavingsPlan columns are visible if you have purchased Savings Plans. For more information about Savings Plans, see [What are Savings Plans?](#) in the *Savings Plans User Guide*.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | VWXYZ

A

savingsPlan/AmortizedUpfrontCommitmentForBillingPeriod

- **Description:** The amount of upfront fee a Savings Plan subscription is costing you for the billing period. The initial upfront payment for **All Upfront Savings Plan** and **Partial Upfront Savings Plan** amortized over the current month. For **No Upfront Savings Plan**, the value is 0.
- **Line items applicable:** SavingsPlanRecurringFee
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

E

savingsPlan/EndTime

- **Description:** The expiration date for the Savings Plan agreement.
- **Line items applicable:** SavingsPlanCoveredUsage, SavingsPlanNegation, SavingsPlanUpfrontFee, SavingsPlanRecurringFee
- **Services:**

- Amazon EC2
- Fargate
- Amazon Lambda
- Amazon SageMaker

I

savingsPlan/InstanceTypeFamily

- **Description:** The instance family that is associated with the specified usage.
- **Line items applicable:** SavingsPlanCoveredUsage
- **Sample values:** m4, g2
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

N

savingsPlan/NetAmortizedUpfrontCommitmentForBillingPeriod

The cost of a Savings Plan subscription upfront fee for the billing period. This column is included in your report only when your account has a discount in the applicable billing period.

savingsPlan/NetRecurringCommitmentForBillingPeriod

The net unblended cost of the Savings Plan fee. This column is included in your report only when your account has a discount in the applicable billing period.

savingsPlan/NetSavingsPlanEffectiveCost

The effective cost for Savings Plans, which is your usage divided by the fees. This column is included in your report only when your account has a discount in the applicable billing period.

O

savingsPlan/OfferingType

- **Description:** Describes the type of Savings Plan purchased.
- **Line items applicable:** SavingsPlanCoveredUsage, SavingsPlanNegation, SavingsPlanUpfrontFee, SavingsPlanRecurringFee
- **Sample values:** ComputeSavingsPlans, EC2InstanceSavingsPlans, SageMakerSavingsPlans
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda

- Amazon SageMaker

P

savingsPlan/PaymentOption

- **Description:** The payment options available for your Savings Plan.
- **Line items applicable:** SavingsPlanCoveredUsage, SavingsPlanNegation, SavingsPlanUpfrontFee, SavingsPlanRecurringFee
- **Sample values:** Partial Upfront, All Upfront, No Upfront
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

savingsPlan/PurchaseTerm

- **Description:** Describes the duration, or term, of the Savings Plan.
- **Line items applicable:** SavingsPlanCoveredUsage, SavingsPlanNegation, SavingsPlanUpfrontFee, SavingsPlanRecurringFee
- **Sample values:** 1yr, 3yr
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

R

savingsPlan/RecurringCommitmentForBillingPeriod

- **Description:** The monthly recurring fee for your Savings Plan subscriptions. For example, the recurring monthly fee for a **Partial Upfront Savings Plan** or **No Upfront Savings Plan**.
- **Line items applicable:** SavingsPlanRecurringFee
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

savingsPlan/Region

- **Description:** The Amazon Region (geographic area) that hosts your Amazon services. You can use this field to analyze spend across a particular Amazon Region.

- **Line items applicable:** SavingsPlanCoveredUsage, SavingsPlanNegation, SavingsPlanUpfrontFee, SavingsPlanRecurringFee
- **Sample values:** US East (N. Virginia), US West (N. California), US East (Ohio), Asia Pacific (Mumbai), Europe (Ireland)
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

S

savingsPlan/SavingsPlanArn

- **Description:** The unique Savings Plan identifier.
- **Line items applicable:** SavingsPlanUpfrontFee
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

savingsPlan/SavingsPlanEffectiveCost

- **Description:** The proportion of the Savings Plan monthly commitment amount (upfront and recurring) that is allocated to each usage line.
- **Line items applicable:** SavingsPlanCoveredUsage
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

savingsPlan/SavingsPlanRate

- **Description:** The Savings Plan rate for the usage.
- **Line items applicable:** SavingsPlanCoveredUsage
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

savingsPlan/StartTime

- **Description:** The start date of the Savings Plan agreement.

- **Line items applicable:** SavingsPlanCoveredUsage, SavingsPlanNegation, SavingsPlanUpfrontFee, SavingsPlanRecurringFee
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

T

savingsPlan/TotalCommitmentToDate

- **Description:** The total amortized upfront commitment and recurring commitment to date, for that hour.
- **Line items applicable:** SavingsPlanRecurringFee
- **Services:**
 - Amazon EC2
 - Fargate
 - Amazon Lambda
 - Amazon SageMaker

U

savingsPlan/UsedCommitment

- **Description:** The total dollar amount of the Savings Plan commitment used. (SavingsPlanRate multiplied by usage)
- **Line items applicable:** SavingsPlanRecurringFee
- **Services:**
 - Amazon EC2
 - Amazon Lambda
 - Fargate
 - Amazon SageMaker

Cost Categories details

You can use the Cost Category columns in Amazon Cost and Usage Reports to find information about the specific line items categorized by the Cost Categories feature. Cost Category columns are automatically populated when you create a Cost Category and categorization rule. These columns include user-defined Cost Category names as the column headers, and corresponding Cost Category values. Examples include the following:

costCategory/Project

Use a **Project** Cost Category to categorize cost and usage information associated with the project. Project names such as "alpha" or "beta" will show as values under **costCategory/Project**.

costCategory/Team

Use a **Team** Cost Category to categorize cost and usage information associated with the team. Team names such as "Marketing" will show as values under **costCategory/Team**.

costCategory/Environment

Use a **Environment** Cost Category to categorize cost and usage information associated with the environment. Environment names such as "Development" will show as values under **costCategory/Environment**.

Discount details

Columns under the **discount** header are included in Amazon Cost and Usage Reports only when the account has a discount applied during the report's billing period. This includes the following columns:

discount/BundledDiscount

The bundled discount applied to the line item. A bundled discount is a usage-based discount that provides free or discounted usage of a service or feature based on the usage of another service or feature. Examples of bundled discounts include:

- If you use Amazon Shield Advanced, then you don't have to pay for Amazon WAF separately. Amazon WAF usage is bundled with Amazon Shield Advanced. For more information about Amazon Shield Advanced, see [Amazon CloudFront Pricing](#).
- If you create a NAT gateway with Amazon Network Firewall, then the standard NAT gateway processing and per-hour usage charges are waived on a one-to-one basis with the firewall's processing per GB and usage hours.

discount/TotalDiscount

The sum of all the discount columns for the corresponding line item.

Use cases

You can use the Amazon Cost and Usage Reports (Amazon CUR) to suit your reports management needs. This section goes in-depth to help you understand use cases such as tracking your Savings Plans and Reserved Instance (RI) utilization, charges, and allocations.

Topics

- [Understanding Savings Plans \(p. 63\)](#)
- [Understanding your reservations \(p. 65\)](#)
- [Understanding your data transfer costs \(p. 74\)](#)

Understanding Savings Plans

You can use Cost and Usage Reports (Amazon CUR) to track your Savings Plans utilization, charges, and allocations.

Savings Plans line items

Savings Plans provide a flexible pricing model that offers low prices on Amazon EC2, Amazon Fargate, Amazon Lambda, and Amazon SageMaker in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1-year or 3-year term.

The following line items in Amazon CUR help you track and manage your spend with Savings Plans.

Note

In the following tables, the columns and rows from Amazon CUR are transposed for clarity. The values in the first column represent the headers of a report. These examples include only a few key Amazon CUR columns. To learn more about other Amazon CUR columns, see the [Data dictionary \(p. 22\)](#).

Upfront fee

The **SavingsPlanUpfrontFee** line item is added to your bill when you purchase an All Upfront or Partial Upfront Savings Plan. The following table shows how this one-time fee appears in some Amazon CUR columns.

lineItem/Linetype	SavingsPlanUpfrontFee	SavingsPlanUpfrontFee
lineItem/ProductCode	ComputeSavingsPlans	ComputeSavingsPlans
lineItem/UsageStartDate	2019-10-10T00:03:54Z	2019-10-10T00:12:15Z
lineItem/LinetypeDescription	USD \$43.80 one-time fee for 1 year All Upfront Compute Savings Plan ID: 70352035	USD \$43.80 one-time fee for 3-year Partial Upfront Any Region M5 Instance Type EC2 Savings Plan ID: 12355516
lineItem/UnblendedCost	43.8	43.8
savingsPlan/SavingsPlanARN	arn:aws:savingsplans::555555555555:savingsplan/bc1d08fd	arn:aws:savingsplans::555555555555:savingsplan/67b0ef20

Savings Plans recurring monthly fee

The **SavingsPlanRecurringFee** line item describes the recurring hourly charges that correspond to **No Upfront** or **Partial Upfront** Savings Plans. Initially, the **SavingsPlanRecurringFee** is added to your bill on the day of purchase and on the first day of each billing period thereafter.

The **SavingsPlanRecurringFee** allocated to the hour (applicable to Hourly cost and usage) or day (applicable to Daily cost and usage) is added to your bill at the hour of purchase. It is added every hour/day of the billing period subsequently.

The following table shows how the recurring hourly charges appear in some Amazon CUR columns.

lineItem/Linetype	SavingsPlanRecurringFee	SavingsPlanRecurringFee
lineItem/UsageStartDate	2019-20-10T00:00:00Z	2019-20-10T00:00:00Z
lineItem/ProductCode	Compute Savings Plans	Compute Savings Plans
lineItem/UsageType	ComputeSP:1yrPartialUpfront	USE2-EC2SP:t3.1yrPartialUpfront
lineItem/UnblendedCost	0.01	0.01
lineItem/LinetypeDescription	1 year Partial Upfront Compute Savings Plan	1 year Partial Upfront t3 EC2 Instance Savings Plan in us-east-2
savingsPlan/SavingsPlanARN	arn:aws:savingsplans::555555555555:savingsplan/bc1d08fd	arn:aws:savingsplans::555555555555:savingsplan/bc1d08fd
savingsPlan/PaymentOption	Partial Upfront	Partial Upfront
savingsPlan/OfferingType	ComputeSavingsPlans	EC2InstanceSavingsPlans
savingsPlan/PurchaseTerm	1yr	1yr

Note

The **SavingsPlanRecurringFee** is calculated differently than the recurring RI fee. The recurring RI fee is a monthly charge while the **SavingsPlanRecurringFee** is an hourly charge. For information on the recurring RI fee, see [Recurring monthly RI fee \(p. 67\)](#).

Savings Plans discount benefits

The **SavingsPlanCoveredUsage** line item describes the instance usage that received Savings Plans benefits. A **SavingsPlanCoveredUsage** line item shows an unblended cost of what the On-Demand charge would have been without the Savings Plan benefit. This unblended cost is offset by the corresponding **SavingsPlanNegation** line item.

In each **SavingsPlanCoveredUsage** line item, you can see how that usage was billed against your Savings Plans hourly commitment by using the **savingsPlan/SavingsPlanRate** and **savingsPlan/SavingsPlanEffectiveCost** fields.

You'll see a corresponding **SavingsPlanNegation** for each **SavingsPlanCoveredUsage** line item. **SavingsPlanNegation** line items offset the unblended cost of **SavingsPlanCoveredUsage**, and grouped at the hourly level by SavingsPlanARN, Operation, Usage Type, and Availability Zone. Therefore, one **SavingsPlanNegation** line item might correspond to multiple **SavingsPlanCoveredUsage** line items.

The following table shows how the covered usage and the negation line items appear in some Amazon CUR columns.

lineItem/ LineItemType	SavingsPlanCoveredUsage	SavingsPlanCoveredUsage	SavingsPlanNegation
lineItem/ UsageStartDate	2019-10-10T00:00:00Z	2019-10-10T00:00:00Z	2019-10-10T00:00:00Z
lineItem/ UsageEndDate	2019-10-10T01:00:00Z	2019-10-10T01:00:00Z	2019-10-10T01:00:00Z
lineItem/ ProductCode	AmazonEC2	AmazonEC2	AmazonEC2
lineItem/ UsageType	BoxUsage:t3.nano	BoxUsage:t3.nano	BoxUsage:t3.nano
lineItem/ UsageAmount	1	0.5	-1.5
lineItem/ UnblendedCost	0.0052	0.0026	-0.0078
lineItem/ LineItemDescription	\$0.0052 per On Demand Linux t3.nano Instance Hour	\$0.0052 per On Demand Linux t3.nano Instance Hour	SavingsPlanNegation used by AccountId : 555555555555 and UsageSku : K7ERD2Q28HHU97DT
SavingsPlan/ SavingsPlanARN	arn:aws:savingsplans::555555555555:savingsplan/bc1d08fd	arn:aws:savingsplans::555555555555:savingsplan/bc1d08fd	arn:aws:savingsplans::555555555555:savingsplan/bc1d08fd
savingsPlan/ SavingsPlanRate	0.0026	0.0026	
savingsPlan/ SavingsPlanEffectiveCost	0.0026	0.0013	

When you have more usage than your Savings Plans commitment can cover, your uncovered usage still appears as a Usage Line Item and the covered usage appears as **SavingsPlanCoveredUsage** with the corresponding **SavingsPlanNegation** line items.

Understanding your reservations

You can use the Amazon Cost and Usage Reports (Amazon CUR) to track your Reserved Instance (RI) utilization, charges, and allocations. This section is an in-depth description to understand your reservations.

Topics

- [Understanding your reservation line items \(p. 66\)](#)
- [Understanding your amortized reservation data \(p. 69\)](#)
- [Monitoring your size flexible reservations for Amazon EC2 \(p. 71\)](#)

- [Monitoring your On-Demand capacity reservations \(p. 72\)](#)

Understanding your reservation line items

RIs provide you a significant discount compared to On-Demand Instance pricing. RIs aren't physical instances. They're a billing discount applied to the use of On-Demand Instances in your account. These On-Demand Instances must match certain attributes to benefit from the billing discount.

Topics

- [Upfront fee \(p. 66\)](#)
- [True-up fee \(p. 66\)](#)
- [Recurring monthly RI fee \(p. 67\)](#)
- [RI discount benefits \(p. 68\)](#)
- [Reserved Instance type \(p. 68\)](#)
- [Reserved Instance benefits applied to instance usage \(p. 69\)](#)

Note

In the following tables, the columns and rows from Amazon CUR are transposed for clarity. The values in the first column represent the headers of a report. These examples include only a few key Amazon CUR columns. To learn more about other Amazon CUR columns, see the [Data dictionary \(p. 22\)](#).

Upfront fee

The **Fee** line item is added to your bill when you purchase an `All Upfront` or `Partial Upfront` RI.

The following table shows how this one-time fee appears in some Amazon CUR columns.

lineitem/LineItemType	Fee
lineitem/ProductCode	AmazonEC2
lineitem/UsageStartDate	2016-01-01T00:00:00Z
lineitem/description	Sign up charge for subscription: 363836886, planId: 1026576
lineitem/UnblendedCost	68
Reservation/ReservationARN	arn:aws:ec2:us-east-1:572481847476:reserved-instances/f8c204c1-dd48-43f1-adb8-f88aa61e0dea

True-up fee

If you exchange a Convertible Reserved Instance, any cost associated with the exchange of the original Reserved Instance and the new Reserved instance (true-up fee) is also added to your bill as a **Fee** line item. For a true-up fee, the **reservation/ReservationARN** column contains **reserved-instance-exchange/riex**.

The following table shows a true-up fee from exchanging a Convertible Reserved Instance.

lineitem/ LineItemType	lineitem/ ProductCode	lineitem/ UsageStartDate	lineitem/ description	lineitem/ UnblendedCost	Reservation/ ReservationARN
Fee	AmazonEC2	2016-01-01T00:00:00Z			arn:aws:ec2:eu-west-1:012345678901:reservation-exchange/riexamplef-5d71-4215-886f-1

Recurring monthly RI fee

The **RI Fee** line item describes the recurring monthly charges that are associated RIs applied that month. The **RI Fee** initially is added to your bill on the day of purchase and on the first day of each billing period thereafter.

The **RI Fee** is calculated by multiplying your discounted hourly rate and the number of hours in the month.

The following table shows how the recurring monthly charges appear in the report.

lineitem/LineItemType	RI fee
lineitem/ProductCode	AmazonEC2
lineitem/UsageStartDate	2016-01-01T00:00:00Z
lineitem/UsageType	HeavyUsage: m4.large
lineitem/description	USD 0.0309 hourly fee per Linux/UNIX (Amazon VPC), m4.large instance
lineitem/ normalization factor	4
lineitem/UnblendedCost	23
Reservation/AvailabilityZone	
Reservation/ReservationARN	arn:aws:ec2:us-east-1:572481847476:reserved-instances/f8c204c1-dd48-43f1-adb8-f88aa61e0dea
Reservation/TotalReservedunits	744
Reservation/TotalReservedNormalizedUnits	2976

Recurring monthly charges are recorded differently for RIs that have an Availability Zone or Amazon Web Services Region Region scope. For RIs that have an Availability Zone scope, the corresponding Availability Zone is shown in the **reservation/AvailabilityZone** column. For RIs that have a Region scope, the **reservation/AvailabilityZone** column is empty. RIs with a Region scope have values for the **lineitem/NormalizationFactor** and **reservation/TotalReservedNormalizedUnits** columns that show the instance size.

Note

The recurring RI fee is calculated differently than the SavingsPlanRecurringFee. The recurring RI fee is a monthly charge while the SavingsPlanRecurringFee is an hourly charge. For information on the SavingsPlanRecurringFee, see [Understanding Savings Plans \(p. 63\)](#).

RI discount benefits

The **Discounted Usage** line item describes the instance usage that received a matching RI discount benefit, and is added to your bill when you have usage that matches one of your RIs. Amazon calculates RI discount benefits based on matching usage: for example, the use of an instance that matches the instance reservation. If you have matching usage, the cost associated with the usage line item is always zero because the charges associated with RIs are already accounted for in the two other line items (the upfront fee and the recurring monthly charges).

The following table shows an example of usage that received an RI discount benefit.

lineitem/LineItemType	DiscountedUsage
lineitem/ProductCode	AmazonEC2
lineitem/UsageStartDate	2016-01-01T00:00:00Z
lineitem/UsageType	BoxUsage:m4.large
lineitem/description	Linux/UNIX (Amazon VPC), m4.large Reserved Instance applied
lineitem/ResourceID	i-1bd250bc
lineitem/AvailabilityZone	us-east-1b
lineitem/normalizationfactor	4
lineitem/NormalizedUsageAmount	4
lineitem/UnblendedRate	0
lineitem/UnblendedCost	0
Reservation/ReservationARN	arn:aws:ec2:us-east-1:572481847476:reserved-instances/f8c204c1-dd48-43f1-adb8-f88aa61e0dea

The value for **UsageAmount** in the Amazon EC2 **DiscountedUsage** line is the actual number of hours used. The value for **NormalizedUsageAmount** is the value for **UsageAmount** multiplied by the value for **NormalizationFactor**. The value for **NormalizationFactor** is determined by the instance size. When an RI benefit discount is applied to a matching line item of usage, the Amazon Resource Name (ARN) value in the **reservation/ReservationARN** column for the initial upfront fees and recurring monthly charges matches the ARN value in the discounted usage line items.

For more information about mapping instance size to normalization factor, see [Support for modifying instance sizes](#) in the *Amazon EC2 User Guide for Linux Instances*.

Reserved Instance type

To determine if your report line items are associated with a Standard Reserved Instance or a Convertible Reserved Instance, filter the **lineitem/LineItemType** column by **Fee** or **RI fee**. Then, review the **product/OfferingClass** column, which indicates the Reserved Instance type.

The following table shows a **Fee** line item from a Standard Reserved Instance:

The following table shows an **RI fee** line item from a Convertible Reserved Instance:

To determine if your report line items are associated with a zonal or regional Reserved Instance, review the **reservation/AvailabilityZone** column. For zonal Reserved Instances, this column shows the corresponding Availability Zone. For regional Reserved Instances, this column is empty.

Reserved Instance benefits applied to instance usage

To understand which instance usage line items benefitted from which Reserved Instances, you can filter your report by one or more of the following columns:

- **reservation/reservationARN**: Filter this column by a reservation ARN to identify which Reserved Instance lease is associated with each line item.
- **lineitem/ResourceId**: Review this column for the ID of the resource that's covered by the Reserved Instance.
- **lineitem/LineItemType**: Filter this column by **Fee**, **RI fee**, or **DiscountedUsage** to determine the associated fees or benefits.
- **lineitem/UsageType**: Filter this column by **HeavyUsage** to identify **RI fee** line items. Or, filter this column by **BoxUsage** to identify **DiscountedUsage** line items.
- **lineitem/UsageAmount**: For **RI fee** line items, this column shows the total number of hours in the month that the Reserved Instance was applied. For **DiscountedUsage** line items, this column shows the total number of hours that the Reserved Instance was applied to a specific instance at the daily or monthly level, depending on how you configured your report.

To understand a size flexible Reserved Instance's number of normalized units applied to instance usage, review the **lineitem/NormalizedUsageAmount** column in your report. The value in this column equals the product of the following columns:

- **lineitem/UsageAmount**: This column shows the metered instance usage measured in hours.
- **lineitem/NormalizationFactor**: For **DiscountedUsage** and **RI fee** line items, this column shows the associated normalization factor of the instance. For more information on the normalization factor, see [Instance size flexibility determined by normalization factor](#) in the *Amazon EC2 User Guide for Linux Instances*.

For Amazon Organizations with multiple accounts, to see which accounts purchased or benefitted from a Reserved Instance, review the following columns:

- **reservation/reservationARN**: Review the reservation ARNs to see which accounts purchased the Reserved Instance. The ARN includes the account ID.
- **lineitem/UsageAccountID**: For **DiscountedUsage** line items, this column identifies the account IDs that received benefits from the purchased Reserved Instances.

Note

A Reserved Instance is a billing subscription and not a resource like an Amazon EC2 instance. Because of this, Reserved Instances that are tagged don't populate line items like a tagged resource. For line items with **DiscountedUsage**, tags populate for the tagged resources and not for the Reserved Instance.

To identify costs associated with a specific Reserved Instance lease, you can filter **Fee** or **RI fee** line items by the Reserved Instance ARN, which is the lease ID. To organize your cost data for Reserved Instances, consider using Amazon Cost Categories. For more information, see [Managing your costs with Amazon Cost Categories](#) in the *Amazon Billing User Guide*

Understanding your amortized reservation data

Amortizing is when you distribute one-time reservation costs across the billing period that is affected by that cost. Amortizing enables you to see your costs in accrual-based accounting as opposed to

cash-based accounting. For example, if you pay \$365 for an All Upfront RI for one year and you have a matching instance that uses that RI, that instance costs you \$1 a day, amortized.

You can see the data that Billing and Cost Management uses to calculate your amortized costs in the following Cost and Usage Reports columns.

Topics

- [Reserved Instance inventory \(p. 70\)](#)
- [Amortization data for the billing period \(p. 70\)](#)
- [Reserved Instance effective costs \(p. 70\)](#)

Note

Not all **reservation/** columns are populated for every Reserved Instance line item. The **reservation/** columns in your report are populated based on the line item type. For example, **RI fee** line items populate the **reservation/UnusedAmortizedUpfrontFeeForBillingPeriod** column. Meanwhile, **DiscountedUsage** line items populate the **reservation/effectivecost** column.

Reserved Instance inventory

You can use the following columns to track your RI inventory. The values for these columns appear only for RI subscription line items (also known as **RI Fee** line items) and not for the actual instances using the RIs.

For more information about column descriptions and sample values, see [the section called "Reservation details" \(p. 29\)](#).

- reservation/UpfrontValue
- reservation/startTime
- reservation/endTime
- reservation/modificationStatus

Amortization data for the billing period

You can use the following columns to understand the amortized costs of your RIs for the billing period. The values for these columns appear only for RI subscription line items (also known as **RI Fee** line items) and not for the actual instances using the RIs.

For more information about column descriptions and sample values, see [the section called "Reservation details" \(p. 29\)](#).

- reservation/amortizedUpfrontFeeForBillingPeriod
- reservation/unusedQuantity
- reservation/unusedNormalizedUnitQuantity
- reservation/unusedRecurringFee
- reservation/unusedAmortizedUpfrontFeeForBillingPeriod

Reserved Instance effective costs

You can use the following columns to understand your effective cost at the instance level. The values for these columns appear only for instance usage line items (also known as **Discounted Usage boxUsage** line items).

For more information about column descriptions and sample values, see [the section called “Reservation details” \(p. 29\)](#).

- reservation/amortizedUpfrontCostForUsage
- reservation/recurringFeeForUsage
- reservation/effectiveCost

Monitoring your size flexible reservations for Amazon EC2

Amazon EC2 Reserved Instances that apply to a Region provide Availability Zone flexibility and instance size flexibility. Reserved Instances that provide Availability Zone flexibility provide a discount on usage in any Availability Zone in the Region. Reserved Instances that provide instance size flexibility provide a discount on usage, regardless of instance size in that family. Size flexible Reserved Instances apply to the smallest instance sizes first. For more information, see [How Reserved Instances are applied](#) in the *Amazon EC2 User Guide for Linux Instances*.

To understand how instance size flexibility provided by your Reserved Instance is applied to your usage, refer to the **lineItem/NormalizationFactor** and **lineItem/NormalizedUsageAmount** columns.

Note

Instance size flexibility is supported only by Linux or Unix Reserved Instances with default tenancy that are assigned to a Region. For more information on the limitations of instance size flexibility for Regional Reserved Instances, see [How regional Reserved Instances are applied](#) in the *Amazon EC2 User Guide for Linux Instances*.

In a Cost and Usage Report, the Reserved Instance usage is applied by default to the account that purchased the Reserved Instance. Any available Reserved Instance benefit that the purchasing account can't use within the hour is then applied to other linked accounts based on the available matching On-Demand Instance usage.

Example

You purchase one `m4.xlarge` RI in a given Region. This `m4.xlarge` RI can be applied automatically to all `m4` instance usage in the same Region. In the following table, Amazon applied the `m4.xlarge` to two separate `m4.large` instances.

lineItem/LineItemType	RIFee	Discounted Usage	Discounted Usage
lineItem/ProductCode	AmazonEC2	AmazonEC2	AmazonEC2
lineItem/UsageStartDate	2016-01-01T00:00:00Z	2016-01-01T00:00:00Z	2016-01-01T00:00:00Z
lineItem/UsageType	HeavyUsage:m4.xlarge	BoxUsage:m4.large	BoxUsage:m4.large
lineItem/description	USD 0.0618 hourly fee per Linux/UNIX (Amazon VPC), m4.xlarge instance	Linux/UNIX (Amazon VPC), m4.large Reserved Instance applied	Linux/UNIX (Amazon VPC), m4.large Reserved Instance applied
lineItem/ResourceID		i-1bd250bc	i-1df340ed
lineItem/UsageAmount		1	1

lineItem/ normalizationfactor	4	4	4
lineItem/ NormalizedUsageAmount		4	4
lineItem/ UnblendedRate		0	0
lineItem/ UnblendedCost	46	0	0
Reservation/ ReservationARN	arn:aws:ec2:us-east-1: 572481847476:reserved- instances /f8c204c1	arn:aws:ec2:us-east-1: 572481847476:reserved- instances /f8c204c1	arn:aws:ec2:us-east-1: 572481847476:reserved- instances /f8c204c1
Reservation/ TotalReservedUnits	744		
Reservation/ TotalReserved NormalizedUnits	5952		

The two `m4.large` usage line items have different **ResourceIDs**, and both received a discount benefit from the single `m4.xlarge` RI. This is shown by matching the **reservationARN** value across the usage and recurring monthly charge line items.

For more information about RI purchase options, see [How you are billed](#) in the *Amazon EC2 User Guide for Linux Instances*.

Monitoring your On-Demand capacity reservations

Capacity reservations enable you to reserve capacity for your Amazon EC2 instances for any duration in a specific Availability Zone. This enables you to create and manage capacity reservations separately from the billing discounts offered by Regional Reserved Instances (RI). To benefit from billing discounts, you can use Regional RIs in combination with capacity reservations.

Capacity reservation line items

You can use some columns defined in the Amazon CUR data dictionary to track your capacity reservations. The following columns are also used for capacity reservations.

This section defines these line items with supplementary definitions specific to capacity reservations.

For more information about Cost and Usage Reports column descriptions, see [the section called "Line item details"](#) (p. 24).

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z

B

[lineItem/BlendedRate](#)

For capacity reservations with a **UsageType** of **Reservation** or **DedicatedRes**, the **BlendedRate** is 0. This is because the capacity reservation costs are associated with the instance that provides the capacity, instead of with the capacity reservation itself.

R

lineItem/ResourceId

If you included `lineItem/ResourceId` when you created your Cost and Usage Reports, you can identify and track your capacity reservations using the **ResourceId** column. The capacity reservation **ResourceId** is captured only for the **UnusedBox**, **UnusedDed**, **Reservation**, and **DedicatedRes** **UsageTypes**.

Capacity reservations always include a `cr-` in their resource ID, and the resource ID has the following format:

```
arn:aws:ec2:<region>:<account id>:<capacity-reservation>/cr-0be443example1db6f
```

U

lineItem/UnblendedCost

The **BlendedRate** multiplied by the **UsageAmount**.

lineItem/UnblendedRate

For capacity reservations with a **UsageType** of **Reservation** or **DedicatedRes**, the **UnblendedRate** is 0. This is because the costs for capacity reservations are associated with the instance that provides the capacity, instead of with the capacity reservation itself.

lineItem/UsageAmount

How much of a capacity reservation you've used. Each capacity reservation can have multiple slots for an hour, enabling you to run more than one instance that uses the reservation during an hour. Therefore, it's possible to use more than one instance-hour in an hour. **UsageAmount** is calculated by multiplying the number of instance slots covered by the line item with the number of hours covered by the line item.

lineItem/UsageType

How much of a specific reservation you've used. For Amazon EC2, the options are as follows:

lineItem/lineitemtype = BoxUsage

For this **UsageType**, the **UsageAmount** column is the amount of instance-hours of an instance you've used.

For example, a report covers 1 hour and has a capacity reservation line item that can cover 10 instances. If you use two instance-slots during the time period covered by the report, the **BoxUsage UsageAmount** covers the number of instance hours that you reserved and used. In this case, this is two (the number of used instance slots) multiplied by 1 hour (the time covered by the report) for a total of two. For a report that covers 1 day, the **UsageAmount** is two multiplied by 24, for a total of 48.

DedicatedRes

For a **UsageType** of **DedicatedRes**, the **UsageAmount** column describes how many instance-hours of a dedicated capacity reservation you reserved.

Reservation

For a **UsageType** of **Reservation**, the **UsageAmount** column describes how many instance-hours of a capacity reservation you reserved.

For example, if a report covers one hour and has a capacity reservation line item that can cover 10 instances, the **Reservation UsageAmount** covers the number of instance slots that you reserved. In this

case, that's 10 (the number of available instance slots) multiplied by 1 hour (the time covered by the report) for a total of 10. For a report that covers 1 day, the **UsageAmount** would be 10 multiplied by 24, for a total of 240.

UnusedBox

For a **UsageType** of **UnusedBox**, the **UsageAmount** column describes how many instance-hours of a capacity reservation you reserved, but didn't use.

For example, a report covers 1 hour and has a capacity reservation line item that can cover 10 instances. If you didn't use eight instance-slots during the time period covered by the report, the **UnusedBox UsageAmount** covers the number of instance hours that you reserved but didn't use. In this case, that's eight (the number of unused instance slots) multiplied by 1 hour (the time covered by the report) for a total of eight. For a report that covers 1 day, the **UsageAmount** is eight multiplied by 24, for a total of 192.

UnusedDed

For a **UsageType** of **UnusedDed**, the **UsageAmount** column describes how many instance-hours of a dedicated capacity reservation that you reserved, but didn't use.

Understanding your data transfer costs

You can use your Cost and Usage Report to analyze the data transfer costs from your Amazon usage.

All data transfer costs

To find all data transfer costs in your report, filter the **lineItem/UsageType** column by values that contain the following strings:

- For Region-to-Region data transfers, filter the column for values that contain **AWS-In-Bytes** and **AWS-Out-Bytes**.
- For Amazon Direct Connect data transfers, filter the column for values that contain **Dataxfer-In** and **Dataxfer-Out**.
- For internet data transfers, filter the column for values that contain **DataTransfer-In-Bytes** and **DataTransfer-Out-Bytes**.
- For Amazon CloudFront data transfers, filter the column for values that contain **CloudFront-In-Bytes** and **CloudFront-Out-Bytes**.
- For data transfers between Availability Zones within a Region, filter the column for values that contain **DataTransfer-Regional-Bytes**.

Amazon EC2 data transfer costs

To find data transfer costs from Amazon EC2, filter these columns by the following values:

- **product/ProductFamily**: Data Transfer
- **product/servicecode**: AWSDataTransfer
- **product/usagetype**: DataTransfer-Out-Bytes
- **product/servicename**: Amazon Data Transfer
- **lineItem/UsageType**: DataTransfer-Out-Bytes
- **lineItem/ProductCode**: AmazonEC2

Tip

You can also search the **lineitem/LineitemDescription** column for strings that contain “data transfer”.

VPC peering data transfer costs

To find data transfer costs from VPC peering between Availability Zones within a Region, filter the **lineitem/Operation** column by the values **VPCPeering-In** and **VPCPeering-Out**.

To learn more about the data transfer associated with the line item, review the description in the **lineitem/LineitemDescription** column. To determine the source and destination Regions associated with the data transfer, review the values in the **product/fromLocation** and **product/toLocation** columns.

Legacy reports

This section describes legacy billing reports offered outside of Amazon Cost and Usage Reports. These pages are available for reference, but we strongly recommend you use Amazon Cost and Usage Reports since these reporting methods will be unavailable at a later date.

Topics

- [Detailed Billing Reports \(p. 76\)](#)
- [Migrating from Detailed Billing Reports to Cost and Usage Reports \(p. 76\)](#)
- [Understanding unused reservation costs \(p. 81\)](#)
- [Monthly Report \(p. 86\)](#)
- [Monthly cost allocation \(p. 87\)](#)
- [Amazon Usage Report \(p. 87\)](#)

Detailed Billing Reports

Important

The Detailed Billing Report feature is unavailable for new customers as of 07/08/2019.

Detailed Billing Reports (DBR) contain similar information as Amazon CUR regarding your charges, but calculates the individual line items differently. If you've signed up for both the DBR and Amazon CUR, the line items will not match. However, when the reports are finalized at the end of the month, the total cost will align.

Amazon stores the DBR in Amazon S3 as csv files, using the following naming convention:

```
Amazon account number-aws-billing-detailed-line-items-yyyy-mm.csv.zip
```

Amazon recreates the DBR multiple times a day, overwriting the report. When Amazon overwrites a report, the line items might be in a different order than they were in the previous report. A final report is created at the end of the month. For the following month, Amazon creates a new report file instead of overwriting the final report from the previous month. Reports for previous months remain in your S3 bucket until you delete them.

For information on how to migrate your DBR to Amazon CUR, see [the section called "Migrating From DBR to Amazon CUR" \(p. 76\)](#).

Migrating from Detailed Billing Reports to Cost and Usage Reports

The Detailed Billing Report (DBR) and Amazon Cost and Usage Reports (CUR) both provide information about your charges. If you are using a legacy DBR, we strongly recommend you transfer your report to Cost and Usage Reports.

Topics

- [Comparing benefits of the Cost and Usage Reports \(Amazon CUR\) \(p. 77\)](#)

- [Key differences between the Detailed Billing Report and the Cost and Usage Reports \(p. 78\)](#)
- [Reporting on advanced charge types \(p. 80\)](#)

Comparing benefits of the Cost and Usage Reports (Amazon CUR)

Amazon CUR provides the most comprehensive source of information. It allows you to understand individual costs in depth, and to analyze them in greater detail, which is especially useful at an enterprise scale. Amazon CUR is best suited for customers with complex cost management needs, for example, those with dedicated query or analytic-based systems. Amazon CUR is also your best source for Reserved Instance (RI) information, especially if you want to view amortized costs.

Comprehensive reservation information

Reserved Instances (RI), or reservations, offer a discounted hourly rate compared to On-Demand usage in exchange for committing to a one- or three-year term of service. This can result in significant savings. Amazon CUR helps you monitor and manage your reservation portfolio by providing comprehensive information like reservation Amazon Resource Numbers (ARNs), numbers of reservations, and total RIs. You can track your reservation-related discounts to specific resources, which help you better understand your savings.

DBR provides a subset of this metadata, but work is required to transform the required columns.

Amazon CUR provides additional columns not available in DBR, such as information regarding your amortized reservation costs. For more information, see [the section called "Understanding your amortized reservation data" \(p. 69\)](#).

On-Demand pricing availability

Amazon CUR provides information regarding the On-Demand rates for each individual line item of usage. This information enables you to quantify your savings by subtracting the amount you paid from the On-Demand rate, enabling you to compute your savings as compared to On-Demand prices. This also gives you the flexibility of choosing to allocate your costs using public On-Demand rates.

DBR doesn't contain information for On-Demand rates, but only the billed amount. This makes it difficult to calculate your overall savings or to allocate costs using On-Demand rates.

Granular breakdown of discounts

Amazon CUR can access a granular view of the usage-based discounts. If discounts were applied, you can use Amazon CUR to view the following:

- Cost prior to being discounted
- Discounted amount
- Total cost after the discount was applied at the line item level

DBR does not contain a granular breakdown of your discounts.

Automated data ingestion at scale

When you use Amazon CUR, you can easily configure an event to trigger an automated data ingestion process, streamlining the process of refreshing the billing data in your in-house systems. Amazon CUR data can automatically be refreshed when charges related to previous months are detected.

Additionally, Amazon CUR is generated as multiple files, providing the added benefit of segmenting the data into smaller pieces. This makes it easier to ingest the data according to the processes used by multiple workers. It also allows you to retry data downloads in smaller pieces.

Amazon CUR is formatted in a way that enables you to locate and extract data quickly. This report is modeled from a manifest file that contains information for the overall structure of the data, including a list of every column that is contained in the report. This enables you to extend the report and include new information regarding your usage when it becomes available.

Cross-product integration

Amazon CUR supports integration with Amazon Redshift, Amazon QuickSight, and Amazon Athena, making it easy to quickly build out an Amazon based cost management solution. Amazon CUR also provides data in Parquet format, broadening your options when it comes to building out your own cost and usage reporting system. For more information, see [Amazon Cost and Usage Reports Manifest Files](#) in the *Amazon Billing User Guide*.

Key differences between the Detailed Billing Report and the Cost and Usage Reports

There are a few differences between DBR and Amazon CUR that you should be aware of after you migrate to using Amazon CUR. You may need to adjust how you ingest the data into your systems accordingly.

File structure

DBR is delivered as a single file, while Amazon CUR provides as a consolidated set of files. In Amazon CUR, you can view the following files in your Amazon S3 bucket:

- Set of data files containing all of your usage line items
- Separate data file containing all of your discounts (if applicable)
- Manifest file that lists all of the data files that belong to a single report

Column structure

DBR has a fixed list of columns, limiting its flexibility. Amazon CUR does not have a fixed column structure, and instead allows you to freely add or remove columns as needed. When you begin using a new Amazon service, Amazon CUR can dynamically start to include new data in the report that may be useful in your case. The manifest file provides a map of all columns present in the report.

Equivalent Column Names for DBR and Amazon CUR

DBR Column Name	Amazon CUR Column Name
InvoiceId	bill/InvoiceId
PayerAccountId	bill/PayerAccountId
LinkedAccountId	lineItem/UsageAccountId
ProductName	product/ProductName
SubscriptionId	reservation/subscriptionid

Amazon Cost and Usage Reports User Guide
Key differences between the Detailed Billing Report and the Cost and Usage Reports

DBR Column Name	Amazon CUR Column Name
UsageType	lineItem/UsageType
Operation	lineItem/Operation
AvailabilityZone	lineItem/AvailabilityZone
ReservedInstance	Not Supported
ItemDescription	lineItem/LineItemDescription
UsageStartDate	lineItem/UsageStartDate
UsageEndDate	lineItem/UsageEndDate
UsageQuantity	lineItem/UsageAmount
BlendedRate	lineItem/BlendedRate
BlendedCost	lineItem/BlendedCost
UnBlendedRate	lineItem/UnblendedRate
UnBlendedCost	lineItem/UnblendedCost
ResourceId	lineItem/ResourceId
RecordType	Not Supported
Pricingplanid	Not Supported
RateID	pricing/RateId

Note

There is no equivalent for RecordId in Amazon CUR, but you can gather this information by combining identity/LineItemId, identity/TimeInterval, and bill/BillType.

Retrieving DBR RecordType Values Through Amazon CUR

RecordType values in DBR	Syntax to Retrieve RecordType Through Amazon CUR	Use Case
LineItem	SELECT SUM(line_item_unblended_cost) FROM [CUR] WHERE line_item_line_item_type = 'Usage'	Usage line item partitions out usage costs from one-time charges. For example: upfront RI payment
InvoiceTotal	SELECT (bill_invoice_id), sum(line_item_unblended_cost) FROM [CUR] GROUP BY bill_invoice_id	Invoice total helps you reconcile your costs between Invoices and the Cost and Usage Report.
AccountTotal	SELECT line_item_usage_account_id, sum(line_item_unblended_cost) FROM [CUR] GROUP BY line_item_usage_account_id	Account total helps you isolate costs related to your member accounts for charge back purposes.

RecordType values in DBR	Syntax to Retrieve RecordType Through Amazon CUR	Use Case
StatementTotal	SELECT SUM(line_item_unblended_cost) FROM [CUR]	Statement total helps you understand your costs for the billing period.
Discount	SELECT SUM(line_item_unblended_cost) FROM [CUR] WHERE line_item_line_item_type = 'Discount'	Discount line items helps you identify all of your discount-related line items.
Rounding	Not yet supported	Not yet supported

Reporting on advanced charge types

Refunds

Amazon CUR: Refunds are identified by filtering for the `lineItem/LineItemDescription = 'Refund'` string.

DBR: Refunds can be identified by parsing the `ItemDescription` column for the 'Refund' substring.

Credits

Amazon CUR: Credits are identified by filtering for the `lineItem/LineItemDescription = 'Credit'` string.

DBR: Credits are identified through parsing the `ItemDescription` column for the 'Credit' substring.

Taxes

Amazon CUR: Taxes are identified by filtering the `lineItem/LineItemDescription = 'Tax'` string.

DBR: Taxes are identified by parsing the `ItemDescription` column for the 'Tax' substring.

Identifying reservation-related upfront costs

Amazon CUR: Reservation-related upfront fees can be identified by filtering `"lineitem/lineitemtype" = 'Fee'`.

DBR: Reservation-related upfront costs can be identified by examining the `UsageType` column for the 'HeavyUsage' substring, and whether the 'SubscriptionId' is null.

Identifying reservation-related monthly fee

Amazon CUR: Reservation-related monthly fees can be identified by filtering `"lineitem/lineitemtype" = 'RIfee'`.

DBR: Reservation-related monthly fees can be identified by examining the `UsageType` column for the 'HeavyUsage' substring.

Identifying instances that received reserved instance benefits

Amazon CUR: Reservation-related upfront fees can be identified by filtering `"lineitem/lineitemtype" = 'DiscountedUsage'`.

DBR: Reservation-related upfront fees can be identified by filtering 'ReservedInstance' = 'Y'.

Understanding unused reservation costs

Amazon Cost and Usage Reports (Amazon CUR) can be leveraged to better understand unused RI costs. Here's four scenarios showing how.

Note

In the following tables, the columns and rows from Amazon CUR and DBR/DBR-RT are transposed for clarity. The values in the first column represent the headers of a report.

Scenario 1: RI usage is 100%

RI Fee line item has \$0 unused cost and 0 usage hours.

Using the DBR/DBR-RT, you can understand your unused RI usage and costs by referring to the fields UsageQuantity and UnblendedCosts for RI Fee line items. RI Fee line items can be identified by the existence of 'purchased hours' information in the ItemDescription field. Table 1 illustrates the columns and information used to manage unused RI costs in the DBR and DBR-RT report.

Table 1 – Unused RI costs for a 100% RI usage in DBR and DBR-RT prior to June 17, 2019

ProductName	Amazon Elastic Compute Cloud	Amazon Elastic Compute Cloud
UsageType	HeavyUsage:c3.8xlarge	HeavyUsage:c3.8xlarge
Operation	RunInstances	RunInstances
Availability Zone	us-east-1a	us-east-1a
Reserved Instance	Y	Y
ItemDescription	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge (744 hours purchased, 744 hours used)	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge
Usage Quantity	0	744
Unblended Rate	0.1	0.1
Unblended Cost	0	74.4

Using Amazon CUR, you can understand your unused RI usage and costs by referring to the fields 'reservation/ UnusedQuantity' and 'reservation/ UnusedRecurringFee' for RI Fee line items. Table 4 below illustrates the current columns and information utilized to manage unused RI costs in Amazon CUR.

Table 2 – Unused RI costs for a 100% RI usage in Amazon CUR

lineitem/Productcode	Amazon EC2	Amazon EC2
UsageType	HeavyUsage:c3.8xlarge	USW2-BoxUsage:c3.8xlarge

lineitem/LinItemtype	RI Fee	DiscountedUsage
lineitem/LinItemDescription	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge	USD 0.00 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge
lineitem/UsageAmount	744	744
lineitem/NormalizedUsageAmount	47,616	47,616
lineitem/UnblendedRate	0.1	0
lineitem/UnblendedCost	74.4	0
reservation/UnusedQuantity	0	
reservation/UnusedRecurringFee	0	
reservation/UnusedAmortizedUpfrontFeeForBillingPeriod	0	
reservation/RecurringFeeForUsage		74.4
reservation/AmortizedUpfrontCostForUsage		5
reservation/EffectiveCost		79.4

In addition to matching the current functionality supported by DBR/DBR-RT, Amazon CUR has the following advantages:

- Using Amazon CUR, you are able to access information regarding the EffectiveCost for the DiscountedUsage line item, which includes both the recurring and upfront fees. The DBR only accounts for recurring fees.
- In Amazon CUR, the UsageType field is not transformed for the DiscountedUsage line items whereas DBR replaces the information with RI Fee line item information. This is because the user can group line items in Amazon CUR by ReservationARN in order to understand what usage was discounted by which RI.
- In Amazon CUR, the LinItemDescription field is not transformed for the RI Fee line item. DBR appends the hours purchased and hours used.

Scenario 2: Partial RI usage

RI Fee line item has unused cost and usage.

Using the DBR/DBR-RT, you can understand your unused RI usage and costs by referring to fields UsageQuantity and UnblendedCosts for RI Fee line items. Table 3 illustrates the columns and information used to manage unused RI costs in the DBR and DBR-RT report.

Table 3 – Unused RI costs for a partial RI usage in DBR and DBR-RT prior to June 17, 2019

ProductName	Amazon Elastic Compute Cloud	Amazon Elastic Compute Cloud
--------------------	------------------------------	------------------------------

UsageType	HeavyUsage:c3.8xlarge	HeavyUsage:c3.8xlarge
Operation	RunInstances	RunInstances
Availability Zone	us-east-1a	us-east-1a
Reserved Instance	Y	Y
ItemDescription	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge (744 hours purchased, 644 hours used)	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge
Usage Quantity	100	644
Unblended Rate	0.1	0.1
Unblended Cost	10	64.4

Using Amazon CUR, you can understand your unused RI usage and costs by referring to fields 'reservation/ UnusedQuantity' and 'reservation/ UnusedRecurringFee' for RI Fee line items. Table 4 illustrates the current columns and information utilized to manage unused RI costs in Amazon CUR.

Table 4 – Unused RI costs for a partial RI usage in Amazon CUR

lineitem/Productcode	Amazon EC2	Amazon EC2
UsageType	HeavyUsage:c3.8xlarge	USW2-BoxUsage:c3.8xlarge
lineitem/Linetype	RI Fee	DiscountedUsage
lineitem/LinetypeDescription	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge	USD 0.00 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge
lineitem/UsageAmount	744	644
lineitem/NormalizedUsageAmount	47,616	47,216
lineitem/UnblendedRate	0.1	0
lineitem/UnblendedCost	74.4	0
reservation/UnusedQuantity	100	
reservation/UnusedRecurringFee	0	
reservation/UnusedAmortizedUpfrontFeeForBillingPeriod	10	
reservation/RecurringFeeForUsage		64.4
reservation/AmortizedUpfrontCostForUsage		5
reservation/EffectiveCost		69.4

In addition to matching the current functionality supported by DBR/DBR-RT, Amazon CUR has the following advantages:

- Amazon CUR has a separate column representing UnusedQuantity for the RI Fee line item vs. DBR / DBR-RT which overloads the UsageQuantity column with the unused hours

Scenario 3: Capacity reservation

DBR/DBR-RT filters out Capacity Reservations related UnusedBox and UnusedDed usage type line items when covered by an RI because the RI Fee line item already covers the unused amount in the UsageQuantity and UnblendedCost fields. Table 5 illustrates the columns and information utilized to manage unused RI costs in the DBR and DBR-RT report.

Table 5 – Unused RI costs for Capacity Reservation scenario in DBR and DBR-RT prior to June 17 2019

ProductName	Amazon Elastic Compute Cloud	Amazon Elastic Compute Cloud
UsageType	HeavyUsage:c3.8xlarge	HeavyUsage:c3.8xlarge
Operation	RunInstances	RunInstances
Availability Zone	us-east-1a	us-east-1a
Reserved Instance	Y	Y
ItemDescription	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge (744 hours purchased, 734 hours used)	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge
Usage Quantity	10	734
Unblended Rate	0.1	0.1
Unblended Cost	1	73.4

Amazon CUR shows these line items as DiscountedUsage. Table 6 illustrates the current columns and information utilized to manage unused RI costs in Amazon CUR.

Table 6 – Unused RI costs for the Capacity Reservation scenario in Amazon CUR

lineitem/Productcode	Amazon EC2	Amazon EC2	Amazon EC2
UsageType	HeavyUsage:c3.8xlarge	USW2-Reservation:c3.8xlarge	USW2-BoxUsage:c3.8xlarge
lineitem/Linetype	RI Fee	Usage	DiscountedUsage
lineitem/LinetypeDescription	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge	USD 0.00 per Reservation Linux/UNIX (Amazon VPC), c3:8xlarge Instance Hour	USD 0.00 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge
lineitem/UsageAmount	744	744	744

lineitem/ NormalizedUsageAmount	47,616		47,216
lineitem/ UnblendedRate	0.1	0	0
lineitem/ UnblendedCost	74.4	0	0
reservation/ RecurringFeeForUsage			64.4
reservation/ AmortizedUpfrontCostForUsage			5
reservation/ EffectiveCost			69.4

Scenario 4: Size flexible reservations

Utilizing the DBR/DBR-RT, you can understand your unused RI usage and costs by referring to fields UsageQuantity and UnblendedCosts for RI Fee line items. RI Fee line items can be identified by the existence of 'purchased hours' information in the ItemDescription field. Table 9 illustrates the columns and information utilized to manage unused RI costs in the DBR and DBR-RT report.

Table 7 – Unused RI costs for a size flex RI scenario in DBR and DBR-RT prior to June 17, 2019

ProductName	Amazon Elastic Compute Cloud	Amazon Elastic Compute Cloud
UsageType	HeavyUsage:c3.8xlarge	HeavyUsage:c3.8xlarge
Operation	RunInstances	RunInstances
Availability Zone	us-east-1a	us-east-1a
Reserved Instance	Y	Y
ItemDescription	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge (744 hours purchased, 644 hours used)	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge; UsageType: BoxUsage:c3.large
Usage Quantity	100	644
Unblended Rate	0.1	0.1
Unblended Cost	10	64.4

Using Amazon CUR, you can understand your unused RI usage and costs by referring to fields 'reservation/ UnusedQuantity' and 'reservation/ UnusedRecurringFee' for RI Fee line items. Table 8 illustrates the current columns and information utilized to manage unused RI costs in the Amazon CUR.

Table 8 – Unused RI costs for a size flex RI scenario in Amazon CUR

lineitem/Productcode	Amazon EC2	Amazon EC2
-----------------------------	------------	------------

UsageType	HeavyUsage:c3.8xlarge	USW2-BoxUsage:c3.8xlarge
lineitem/Linetype	RI Fee	DiscountedUsage
lineitem/LinetypeDescription	USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge	USD 0.00 hourly fee per Linux/UNIX (Amazon VPC), c3:8large
lineitem/UsageAmount	744	644
lineitem/NormalizedUsageAmount	47,616	2,576
lineitem/UnblendedRate	0.1	0
lineitem/UnblendedCost	74.4	0
reservation/UnusedQuantity	100	
reservation/UnusedRecurringFee	70.37	
reservation/UnusedAmortizedUpfrontFeeForBillingPeriod	5.5	
reservation/RecurringFeeForUsage		4.03
reservation/AmortizedUpfrontCostForUsage		0.5
reservation/EffectiveCost		4.53

In addition to matching the current functionality supported by DBR/DBR-RT, Amazon CUR has the following advantages:

- Amazon CUR has the NormalizedUsageAmount and quantity. The DBR / DBR-RT do not have columns representing this.
- Amazon CUR UsageType and Operation are not transformed for the DiscountedUsage lineitem. The DBR / DBR-RT replaces these values with the RI Fee line item.
- Amazon CUR LinetypeDescription is not transformed for the DiscountedUsage line item. In DBR / DBR-RT, which replaces with the RI Fee line item description and appends the DiscountedUsage line item Usage Type to the end of the string i.e. "USD 0.10 hourly fee per Linux/UNIX (Amazon VPC), c3:8xlarge; UsageType: BoxUsage:c3.large"

Monthly Report

You can download a Monthly Report of your estimated Amazon charges from the **Bills** page of the Billing and Cost Management console.

If you use the consolidated billing feature in Amazon Organizations, this report is available only for a management account, and includes activity for all the member accounts. Member account owners can obtain the monthly report only from the management account. For more information, see [Consolidated Billing for Organizations](#) in the *Amazon Billing User Guide*.

The report contains line items for each unique combination of Amazon product, usage type, and operation that the account uses. The estimated report is updated up to several times per day. You can

get reports for previous months by selecting the statement period, starting with the report for the month when you signed up for monthly reports. Reports from before you signed up are not available.

Monthly cost allocation

Important

The Monthly Cost Allocation Report feature will be unavailable at a later date. We strongly recommend that you use the Amazon Cost and Usage Reports instead.

You can create custom cost allocation tag sets for your Amazon resources that can describe the business dimensions of your Amazon usage. These tag sets enable you to organize and track your Amazon costs. Many Amazon services expose tagging in their feature sets. You create the tags within those services by using the console, API, or the Amazon command line interface (CLI). For more information, see [Using Cost Allocation Tags](#) in the *Amazon Billing User Guide*.

After you create your tags, you can obtain a monthly cost allocation report/. This is essentially the Monthly Report including your cost allocation tag sets.

Amazon Usage Report

Important

The Amazon Usage Report feature will be unavailable at a later date. We strongly recommend that you use the Amazon Cost and Usage Reports instead.

You can download dynamically generated Amazon usage reports. Each report covers a single service, and you can choose which usage type, operation, and time period to include. You can also choose how the data is aggregated.

Troubleshooting Cost and Usage Reports

Use the following topics to help you troubleshoot common issues with Cost and Usage Reports.

Topics

- [There are no report files in the Amazon S3 bucket \(p. 88\)](#)
- [One of my report data partitions is empty \(p. 88\)](#)
- [My Cost and Usage Report data doesn't match the data in other Billing and Cost Management features \(p. 89\)](#)
- [I want to backfill data because I changed the settings of my report \(p. 89\)](#)
- [My report file folder in Amazon S3 is in an unnamed folder \(p. 89\)](#)
- [I can't select the option to include resource IDs on my report \(p. 90\)](#)
- [My Cost and Usage Report queries for Amazon Athena don't work on Amazon Redshift, or my Amazon Redshift queries don't work on Amazon Athena \(p. 90\)](#)
- [The columns included in my report have changed from a previous month \(p. 90\)](#)
- [The queries or tables based on my report don't work because the columns in my report have changed \(p. 91\)](#)
- [I need help querying my report \(p. 91\)](#)
- [I can't find the billing data for my Amazon EC2 Dedicated Host \(p. 91\)](#)
- [I don't understand the billing data for my Amazon EC2 Elastic IP addresses \(p. 91\)](#)
- [I use consolidated billing and I don't understand the difference between unblended and blended rates or costs \(p. 92\)](#)
- [Some line items in my report have a blended rate or blended cost of 0 \(p. 92\)](#)
- [I don't understand how All Upfront Reserved Instances are amortized in my report \(p. 92\)](#)

There are no report files in the Amazon S3 bucket

Confirm that the Amazon S3 bucket policy grants the **billingreports.amazonaws.com** service permission to put files in the bucket. For more information on the required bucket policy, see [Setting up an Amazon S3 bucket for Cost and Usage Reports \(p. 3\)](#).

One of my report data partitions is empty

If a report is larger than most applications can handle, then Amazon splits the report into multiple files. A report update might have fewer individual file partitions than an earlier report version.

Review the report's manifest file to find any empty files that you don't need to ingest.

My Cost and Usage Report data doesn't match the data in other Billing and Cost Management features

Other Billing and Cost Management features (Cost Explorer, Detailed Billing Reports, Billing and Cost Management console) might present your costs differently for the following reasons:

- The billing features round cost data in different ways.
- The billing features might have different data refresh settings. For example, you can choose whether or not your Cost and Usage Report automatically refreshes a previously closed bill with any refunds, credits, or Support fees applied after the bill is finalized. Cost Explorer automatically reflects the same items. In this scenario, if you don't activate the automatic refresh on your Cost and Usage Report, then the Cost and Usage Report data won't match the Cost Explorer data.
- The billing features can group charges differently. For example, the **Bills** page in the Billing and Cost Management console shows data transfer charges as a separate **Data Transfer** grouping within your **Amazon Service Charges**. Meanwhile, Cost and Usage Reports and Cost Explorer show data transfer charges as a usage type for each service.

I want to backfill data because I changed the settings of my report

Open a support case to request a backfill of your cost data. In your support case, make sure to provide the report name and the billing period that you want backfilled. For more information on opening a case, see [Contacting customer support \(p. 93\)](#).

Note that you can't get a backfill of cost data for the following scenarios:

- You can't get a backfill for cost data from before the date that you created the account.
- If you use Amazon Organizations and the structure of your organization changed, such as which account is designated the management account, then you can't get a backfill of data with the previous organization structure.
- If you use Amazon Organizations and you change organizations, then you can't get a backfill of data from prior to joining your current organization.

My report file folder in Amazon S3 is in an unnamed folder

Any **/** character in the **Report path prefix** of your report generates an unnamed folder in your Amazon S3 bucket. To remove the unnamed folder in your next report update, edit your report settings and remove the **/** character from the **Report path prefix**. For instructions, see [Editing your Cost and Usage Reports configuration \(p. 11\)](#).

I can't select the option to include resource IDs on my report

When you create your report, you can select the option to **Include resource ID**. If you create your report with **Report versioning** set to **Overwrite existing report**, then you can't modify your **Include resource ID** selection after you create your report. To include resource IDs, you must create a new report and select the **Include resource ID** option.

My Cost and Usage Report queries for Amazon Athena don't work on Amazon Redshift, or my Amazon Redshift queries don't work on Amazon Athena

Amazon Athena and Amazon Redshift databases format Cost and Usage Report columns differently. Amazon Athena adds an underscore between words in the column name (line_item_normalized_usage_amount). Amazon Redshift adds an underscore between the column type and the attribute (lineitem_normalizedusageamount). Make sure to modify your queries to match the column name format in Amazon Athena or Amazon Redshift.

The columns included in my report have changed from a previous month

The columns that Amazon includes in your report depend on your Amazon usage. Every report includes columns with the **identity/**, **bill/**, and **lineitem/** prefixes:

- identity/LineItemId
- identity/TimeInterval
- bill/InvoiceId
- bill/BillingEntity
- bill/BillType
- bill/PayerAccountId
- bill/BillingPeriodStartDate
- bill/BillingPeriodEndDate
- lineitem/UsageAccountId
- lineitem/LineItemType
- lineitem/UsageStartDate
- lineitem/UsageEndDate
- lineitem/ProductCode
- lineitem/UsageType
- lineitem/Operation
- lineitem/AvailabilityZone
- lineitem/ResourceId

- lineItem/UsageAmount
- lineItem/NormalizationFactor
- lineItem/NormalizedUsageAmount
- lineItem/CurrencyCode
- lineItem/UnblendedRate
- lineItem/UnblendedCost
- lineItem/BlendedRate
- lineItem/BlendedCost
- lineItem/LineItemDescription
- lineItem/TaxType
- lineItem/LegalEntity

All other columns are included only if your monthly Amazon usage generates data to populate those columns. For example, your report includes **savingsPlan/** columns only if you used Savings Plans during that month.

The queries or tables based on my report don't work because the columns in my report have changed

The columns that Amazon includes in your report depend on your Amazon usage for the month. Because the columns included in your report can change, it's a best practice to reference column names instead of column numbers in any custom queries or tables based on your report.

I need help querying my report

For detailed information about querying your Cost and Usage Report, see [CUR Query Library Help](#) in the Amazon Well-Architected Labs website.

I can't find the billing data for my Amazon EC2 Dedicated Host

In the **ResourceID** column, look for the Dedicated Host ID rather than the instance ID. Because Dedicated Hosts are metered by Dedicated Host running hours, your report shows Dedicated Host usage by metered hours associated with the host ID.

I don't understand the billing data for my Amazon EC2 Elastic IP addresses

Amazon EC2 Elastic IP addresses are metered in aggregate. This means that each line item in your report doesn't correspond with an individual Elastic IP address. Each line item represents the total number

of chargeable hours. You can have one Elastic IP address assigned to a running instance at no charge. You're charged per hour on a pro-rata basis for each additional Elastic IP address that you assign to the instance. Additionally, Amazon charges an hourly fee for unassigned Elastic IP addresses.

I use consolidated billing and I don't understand the difference between unblended and blended rates or costs

With consolidated billing for Amazon Organizations, unblended and blended rates or costs can help you understand how much an account's usage would cost for a standalone account versus a linked account in an organization. Some services offer pricing tiers that can lower unit costs as usage increases. Because Amazon aggregates all usage for a service in an organization, individual accounts might access lower-priced tiers sooner when their usage is aggregated in an organization's monthly usage.

Unblended rates are the rates associated with an individual account's service usage. For a line item, the unblended cost is usage multiplied by the unblended rate. The unblended cost would be the cost of the account's usage if it were a standalone account. Blended rates are the rates associated with total usage in an organization averaged across accounts. For a line item, the blended cost is usage multiplied by the blended rate. The blended cost is the cost attributed to the account's usage as a linked account in an organization.

For more information and examples of calculating unblended and blended costs, see [Understanding Consolidated Bills](#) in the *Amazon Billing User Guide*

Some line items in my report have a blended rate or blended cost of 0

Amazon EC2 line items with a Reserved Instance discount have an blended rate of zero. For these line items, the **LineItemType** is **Discounted Usage**.

The blended cost is the usage multiplied by the blended rate. If the value for blended rate or usage is zero, then the blended cost is also zero.

I don't understand how All Upfront Reserved Instances are amortized in my report

Because All Upfront Reserved Instances are paid in full upfront, the amortized costs are reflected in your report as the upfront payment divided over the associated time period (one year or three years).

reservation/AmortizedUpfrontCostForUsage and **reservation/EffectiveCost** are the same rate for All Upfront Reserved Instances. This is because both columns are an equal division of the upfront payment for the Reserved Instance over the total hours of its term.

It's expected that your report has **RIFee** line items populated for All Upfront Reserved Instances, even though the **RIFee** is \$0.00. These line items represent the recurring hourly costs for the month, and they have additional usage data in other columns. All Reserved Instances generate **RIFee** line items.

Contacting customer support

Use the following resources to get help with your Amazon Billing and Cost Management and Amazon Cost and Usage Reports questions.

- [Amazon Knowledge Center](#): This is the quickest way to find answers to questions about your Amazon CUR or bill. We recommend you start here.
- **Account and Billing Support**: If you're an Amazon account owner, you have access to account and billing support free of charge. Only personalized technical support requires a support plan. For more information, visit the Amazon Web Services Support web [site](#).
- **Opening a Support Case**: You can contact Amazon Web Services Support and open a support case for your inquiry. This is the most direct method for communicating with Amazon Web Services Support. Amazon Web Services Support does not publish a direct phone number to reach representatives, but instead will call you through the following procedure.

Note

To open an Amazon Web Services Support case and specify *Regarding: Account and Billing Support*, you must either be signed into Amazon as the root account owner, or have IAM permissions to open a support case. For more information, see [Accessing Amazon Web Services Support](#) in the Amazon Web Services Support User Guide.

To contact Amazon Web Services Support

1. Sign in and navigate to the [Amazon Web Services Support Center](#). If prompted, enter the email address and password for your account.
2. Choose **Open a new case**.
3. On the **Open a new case** page, select **Account and Billing Support** and fill in the required fields on the form.
4. Choose your response preference:
 - **Web**: You'll receive an email response from an Amazon Web Services Support representative.
 - **Phone**: You'll receive a telephone call from an Amazon Web Services Support representative.

Note

Instant messaging support is not available for billing inquiries. If you've closed your Amazon account, you can still sign in to contact Customer Support and view past bills.

Security in Amazon Cost and Usage Reports

Cloud security at Amazon is the highest priority. As an Amazon customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Amazon Cost and Usage Reports is a feature in the Amazon Billing and Cost Management console. For details about security considerations, see [Security in Amazon Billing and Cost Management](#) in the *Amazon Billing User Guide*.

For more information on access control and IAM permissions to use Amazon CUR, see [Overview of Managing Access Permissions](#).

Quotas and restrictions

The following table describes the current quotas and restrictions within Amazon Cost and Usage Reports.

Cost and Usage Reports

Number of Cost and Usage Reports	10 per account
Rates	Cost and Usage Reports are free of charge, but standard Amazon S3 rates apply.
Number of Free Tier Amazon Cost and Usage Reports	10

Document history for Amazon Cost and Usage Reports User Guide

The following table describes the documentation for this release of Amazon Cost and Usage Reports.

Latest documentation update: February 10, 2020

update-history-change	update-history-description	update-history-date
New definitions added to product data dictionary	Added 42 new definitions to the Product details page.	February 23, 2022
New security chapter	Added a new security chapter that provides information about various security controls. Former "Controlling Access" chapter contents have been migrated here.	February 10, 2020
Initial launch (p. 96)	Migrated and reorganized all Amazon CUR content from Billing and Cost Management guide.	January 21, 2020

Amazon glossary

For the latest Amazon terminology, see the [Amazon glossary](#) in the *Amazon General Reference*.