

# zadara

## **Zadara Cost and Usage (CAU) Explorer**

**Release 25.07**

**Zadara**  
**Jan 21, 2026**



## CONTENTS

<b>1</b>	<b>Introduction to Zadara Cost and Usage (CAU) Explorer</b>	<b>1</b>
1.1	Early Preview Period . . . . .	1
1.2	Data flow . . . . .	2
1.3	Data Retention period . . . . .	2
<b>2</b>	<b>First steps</b>	<b>3</b>
2.1	Enabling Zadara Cost and Usage (CAU) Explorer . . . . .	3
2.2	CAU Explorer access eligibility . . . . .	3
2.3	How usage is detected . . . . .	5
<b>3</b>	<b>Main view</b>	<b>7</b>
3.1	CAU Explorer views . . . . .	7
3.2	Key functions . . . . .	7
<b>4</b>	<b>Exploring Costs and Usage</b>	<b>13</b>
4.1	Charts . . . . .	13
4.2	Tables . . . . .	14



---

## CHAPTER

# ONE

---

## INTRODUCTION TO ZADARA COST AND USAGE (CAU) EXPLORER

Zadara Cost and Usage (CAU) Explorer is a centralized analytics tool that enables organization billing administrators to analyze their usage and costs across Zadara services.

The CAU Explorer offers both monthly and daily cost visualizations in graphical and tabular formats, with grouping, filtering, drill-downs, and CSV export functionality.

---

### **Important**

All cost figures, usage estimates, and projections displayed within the Cost Explorer do not include applicable taxes or VAT. For official and complete financial details-including tax and VAT-please refer exclusively to the invoices issued to your account, which are considered the definitive source for billing and compliance purposes.

---

### 1.1 Early Preview Period

When accessing the CAU Explorer during the Early Preview period, the CAU Explorer displays an Early Preview notice, informing you of potential limitations and issues.

By proceeding, you acknowledge and accept these conditions:

---

### **Important**

#### **Early Preview**

Please note that the information presented in this service is for informational purposes only and is not a replacement for official invoices provided by Zadara.

The Cost and Usage Explorer service is currently in Early Preview. This means the service is still under active development and may undergo significant changes.

While we strive to provide accurate and reliable data, features and functionality may be limited, and occasional disruptions or inaccuracies may occur.

For billing and payment purposes, always refer to your official invoice.

By using the service during this preview phase, you acknowledge and accept these limitations.

---

#### 1.1.1 CAU Explorer Availability

---

### **Note**

CAU Explorer is available to users with a valid Zadara account in the Global Provisioning Portal at <https://manage.zadarastorage.com>.

---

During the Early Preview period, the CAU Explorer is available for selected customers and on-demand.

---

## 1.2 Data flow

The system tracks cost and usage activity for provisioned resources for all clouds, and pushes that information to a centralized location.

### Note

The CAU Explorer's cost and usage data is processed by a scheduled job on a daily basis, and not as a live stream. As a result, refreshed data viewing can be expected to be delayed by up to 24 hours.

By default, the job that processes periodic costs and usage and populates the CAU database is scheduled to run at 14:00 daily. The job can run for several minutes.

---

### 1.2.1 Estimation



If the daily scheduled process fails to populate the CAU database due to a network outage or another issue, the system calculates an estimate based on historical usage data.

Each week, the system attempts to replace estimated entries with actual usage data.

In the table, affected periods are shown with different shading and a circled question mark icon in the column heading.

When you hover over the shaded column in the table or the corresponding bar in the chart, a tooltip appears with the **Estimation** label.

## 1.3 Data Retention period

Customers can review their historic usage in the CAU Explorer:

- Monthly data is retained for 62 months.
- Daily data is retained for 92 days.

---

## CHAPTER

# TWO

---

## FIRST STEPS

This guide walks you through the first steps of using Zadara Cost and Usage (CAU) Explorer.

### 2.1 Enabling Zadara Cost and Usage (CAU) Explorer

The CAU Explorer must be enabled for your environment before you can use it. Contact Zadara Support to enable the CAU Explorer.

To verify that the CAU Explorer is enabled:

1. Sign on to your Global Provisioning Portal account at <https://manage.zadarastorage.com>.

---

 **Note**

If you do not have an account in the Global Provisioning Portal, go to <https://manage.zadarastorage.com> and select **Sign Up** to create one.

---

2. If the CAU Explorer is enabled, its icon appears in the icon collection on the Provisioning Portal's top navigation bar.



Clicking the icon opens the CAU Explorer, displaying the default view.

### 2.2 CAU Explorer access eligibility

CAU Explorer is available to users with a valid Zadara account in the Global Provisioning Portal at <https://manage.zadarastorage.com>.

Only users who have the **Organization Admin** or **Organization Billing** role can see their organization's costs and usage.

---

 **Note**

For any issues accessing the CAU Explorer, contact Zadara Support for assistance.

---

The CAU Explorer in the Global Provisioning Portal enables any user with a valid Zadara account to view their costs and usage, regardless of whether their account is associated with an on-prem implementation of Zadara services.

The CAU Explorer is a service that operates at the organization level. When enabled, it can be accessed by multiple personas in the organization.

**Organization Admins** who want to grant CAU Explorer permissions to other members of the organization should proceed with [Inviting Users to CAU Explorer](#).

---

 **Important**

By default, the organization's administrator will have access to the CAU service, if it is enabled for that organization.

For organizations that manage sub-organizations, the administrator of each sub-organization will also have access to the CAU service, provided it is enabled for their respective sub-organization.

**Restrict access**

In sub-organizations, if CAU access is restricted, users with the *Member* role will **not** have access to the CAU service, even if the service is enabled.

---

### 2.2.1 Inviting Users to CAU Explorer

Users with the **Organization Admin** role can invite others to use CAU Explorer.

1. In the Global Provisioning Portal account, select **Management** from the options bar at the top right.
2. In the **Management** screen, select the **Users** tab.

A list of users and their roles appears.

3. At the top right of the **Users** tab header, click **Invite User**.

The **Invite User** dialog opens.

1. In the **Invite User** dialog, enter the user's details:

• **Username:**

Enter a sign-on username for the user.

Minimum length: 3 characters

Maximum length: 75 characters

Alphanumeric characters and the following are permitted in the **Username**:

- Hyphen: -
- Underscore: \_
- At-symbol: @
- Plus-symbol: +

• **Email address:**

Enter the user's email address.

An email invitation will be sent to this address. It includes:

- A link to the Provisioning Portal sign-on screen, that activates the account.
- The sign-on username and a temporary password.

• **Organization name:**

From the dropdown, select the organization that the user can access for viewing billing data.

- **Organization role:**

From the dropdown, select one of the following roles:

- **Organization Admin:**

Full admin privileges in the Provisioning Portal.

- **Organization Billing:**

A role that can access the CAU Explorer, and can also manage service inventory items and provision new Zadara services for the account.

This role grants access to CAU Explorer in addition to permissions to manage service inventory items and provision Zadara services.

---

 **Note**

**Organization Member** role users cannot access the CAU Explorer.

---

2. Click **Invite User**.

The user receives an email invitation to access the Provisioning Portal as a member of the specified role.

On first sign-on, users are prompted to change their password.

---

 **Note**

Our recommended practice is to configure the optional second-factor authentication (2FA) during the first sign-on.

---

## 2.3 How usage is detected

Both public and private Zadara clouds track and report usage activity for provisioned resources to a centralized service, which is later securely visualized in Zadara's Global Provisioning Portal.



---

## CHAPTER

# THREE

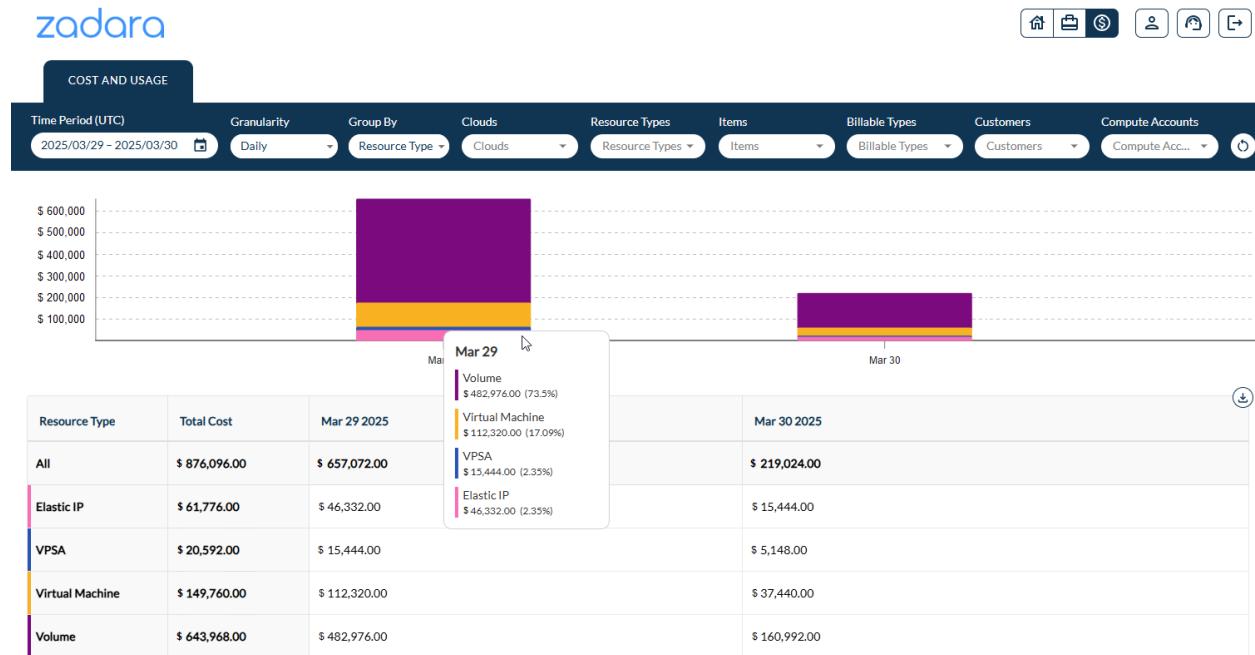
---

## MAIN VIEW

Access CAU Explorer by clicking its icon on the Provisioning Portal's top navigation bar.

### 3.1 CAU Explorer views

The CAU Explorer **Main** screen provides an overview of costs in a bar chart and table, with grouping and filtering.

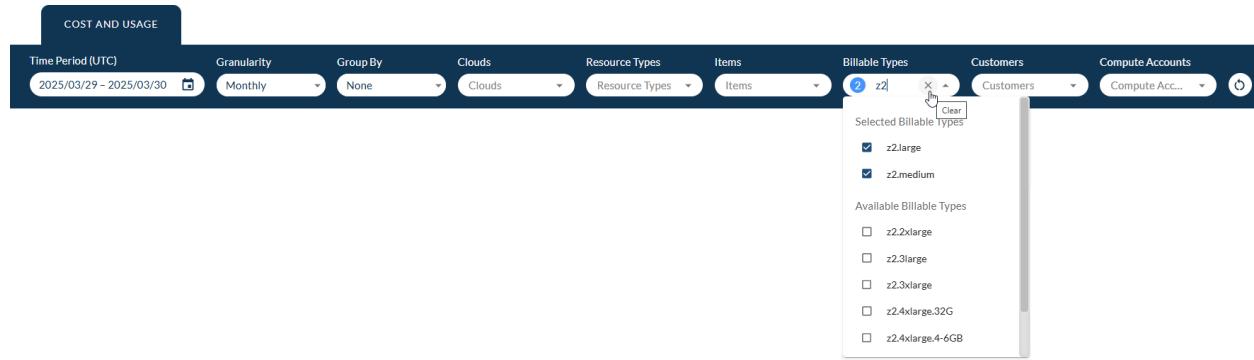


### 3.2 Key functions

The CAU Explorer's dashboard can focus from full consolidated cost information, drilling down to detailed scopes with filtering and grouping functions, and the ability to download the data in the view, and reset the view to default settings.

#### 3.2.1 Filters and groupings

The main view filters bar provides filters and grouping options, to refine and narrow the scope of data in the view.



On the main view filters bar:

- Unfiltered options remain in gray text.
- Selected filters appear in black.
  - **Time Period (UTC)** shows the chosen date range.
  - Single-choice options (**Granularity**, **Group By**) show the selected grouping.
  - Multiple-choice options show the count of selected values that the filter matches and displays.
- **Search** for matching parameter values, by entering a case-insensitive string.
- Hover over a multiple-choice filter to reveal the “x” button for clearing the selection.

You can determine the scope and summary group breakdowns, by applying your selection of the following filters and groupings:

- **Time Period (UTC)**

The CAU Explorer dashboard allows the option to define a specific date range, or select a predefined date range.

The available predefined date range buttons vary, depending on the selected **Granularity** level (**Daily** or **Monthly**):

Date Range Button	Description	Granularity Level	
		Daily	Monthly
<b>Last Week</b>	Data from the same day of the week, one week ago, until now.	✓	
<b>Last Month</b>	Data from the same day of the month, last month, until now.	✓	
<b>Last 3 Months</b>	Data from the same day of month, three months ago, until now.	✓	✓
<b>Last 6 Months</b> (current default)	Data from the same day of month, six months ago, until now.	✓	✓
<b>Last Year</b>	Data from the same day of month, one year ago, until now.	✓	
<b>Last 3 Years</b>	Data from the same day of month, three years ago, until now.	✓	
<b>Last 5 Years</b>	Data from the same day of month, five years ago, until now.	✓	
<b>Default</b>	The default date range.	✓	✓

 **Important**

The data in the date range display is determined by cost and usage event dates aligned to the UTC time zone, and not by any specific site's or cloud's geographic location.

- **Granularity**

Granularity determines the available predefined date range buttons for **Time Period**, and the summary totals level.

Select one of:

- **Monthly** (default) totals, broken down to **Group By** element subtotals (if selected).
- **Daily** totals, broken down to **Group By** elements subtotals (if selected).

- **Group By**

View the table summary totals, broken down according to a selected grouping:

- **None** (default):

Displays a total per **Granularity** unit (month or day) for the defined date range and other filters, and a grand total for all.

- **Resource Type**:

Displays a total per selected **Resource Type** for each **Granularity** unit (month or day) for the defined date range and other filters, a summary total per **Resource Type** and a grand total for all.

- **Cloud**

Displays a total per selected **Cloud** for each **Granularity** unit (month or day) for the defined date range and other filters, a summary total per **Resource Type** and a grand total for all.

- **Clouds**

A cloud refers to a specific deployment of Zadara's infrastructure services, such as compute, storage, and networking, at a specified physical or virtual location.

Each Zadara cloud is uniquely identified and managed as an independent service region or zone.

You can select or deselect clouds from the account's clouds list, to control the scope of billing data displayed.

By default, billing data from all clouds in the account is presented in the view.

- **Resource types**

From the list of billed resource types, there's the option to select or deselect resource types, such as Elastic IP, Virtual Machine, Volume and VPSA, from the scope of the presented billing information.

By default, billing data from all resource types is presented in the view.

- **Items**

An item is an instance of a resource type component associated with a cloud, for example, an HDD or SSD of a specified capacity, a VPSA Engine, a VM instance, and so on.

From a list of items, there's the option to select or deselect Items such as specific instances of resource types, from the scope of the presented billing information.

By default, billing data from all items is presented in the view.

- **Billable types**

From a list of billable types, there's the option to select or deselect billable types, such as data reduction services, VPSA engine and disk types, and compute instance types, from the scope of the presented billing information.

By default, billing data from all billable types is presented in the view.

- **Customers**

A customer is a user's Provisioning Portal username.

From a list of customers, there's the option to select or deselect customers from the scope of the presented billing information.

By default, billing data from all customers is presented in the view.

- **Compute Accounts**

A compute account is a tenant-scoped identity and container, used to isolate and manage virtualized compute resources within a Zadara cloud environment.

From a list of compute accounts, there's the option to select or deselect compute accounts from the scope of the presented billing information.

By default, billing data from all compute accounts is presented in the view.

- **Reset to Default**

The **Reset to Default** button is on the right of the filter and grouping selection options bar.

Click **Reset to Default** to remove the custom period, granularity, grouping and filter selections, and restore the default view.

### 3.2.2 Estimation



If the daily scheduled process fails to populate the CAU database due to a network outage or another issue, the system calculates an estimate based on historical usage data.

Each week, the system attempts to replace estimated entries with actual usage data.

In the table, affected periods are shown with different shading and a circled question mark icon in the column heading.

When you hover over the shaded column in the table or the corresponding bar in the chart, a tooltip appears with the **Estimation** label.

### 3.2.3 Downloads

#### Download main view to CSV (Explorer UI)

The main view **Download CSV** button is located at the top right.

Click **Download CSV** to export the current displayed data, and save it as a file in .csv format. The scope of the saved data is according to the selected time period, granularity, grouping and filters.

The saved .csv file has the same structure and values as the table in the lower section of the CAU Explorer.

#### Download main view to JSON (CLI)

From the command line, you can use the Provisioning Portal's [GET Organizations API](#) to download CAU Explorer data, and optionally apply a custom time period, granularity, grouping and filters.

See the [Zadara Provisioning Portal API](#) documentation's Overview section for information about the **Endpoint** base URL and retrieving an **Authentication** token.

Usage:

```
curl -X 'GET' \
'<Zadara Provisioning Portal URL>/api/v2/organizations/<Organization UUID>/cost_explorer.
˓→.json' \
-H 'accept: application/json' \
-H 'X-Token: <token>'
```

You can optionally extend the URL with a query that applies a custom time period, granularity, grouping and filters.

For example, applying the following filters:

- **Granularity** as monthly:

```
?granularity=monthly
```

- **Time period** from 2024-03-01 to 2024-06-30 in yyyy-mm-dd format:

```
&from=2024%2F03%2F01&to=2024%2F06%2F30
```

- **Group By** cloud:

```
&group_by=cloud
```

Extending the URL to apply the above filters:

```
curl -X 'GET' \
'https://manage.zadarastorage.com/api/v2/organizations/1a2b..36.char.UUID.string..x9z/
˓→cost_explorer.json?granularity=monthly&from=2024%2F03%2F01&to=2024%2F06%2F30&group_
˓→by=cloud \
-H 'accept: application/json' \
-H 'X-Token: A1b2..token.string..yz'
```

The request above returns data in a JSON structure similar to the following response:

```
{
  "data": {
    "2024-03-31": {
      "total": 14508.00,
      "cloud-01": 899.00,
      "cloud-02": 1457.00,
      "cloud-03": 1209.00,
      "cloud-04": 1674.00,
      "cloud-05": 1209.00,
      "cloud-06": 1395.00,
      "cloud-07": 1519.00,
      "cloud-08": 1426.00,
      "cloud-09": 1860.00,
      "cloud-10": 1860.00
    },
    "2024-04-30": {
      "total": 14040.00,
      "cloud-01": 870.00,
```

(continues on next page)

(continued from previous page)

```
"cloud-02": 1410.00,
"cloud-03": 1170.00,
"cloud-04": 1620.00,
"cloud-05": 1170.00,
"cloud-06": 1350.00,
"cloud-07": 1470.00,
"cloud-08": 1380.00,
"cloud-09": 1800.00,
"cloud-10": 1800.00
},
"2024-05-31": {
"total": 14508.00,
"cloud-01": 899.00,
"cloud-02": 1457.00,
"cloud-03": 1209.00,
"cloud-04": 1674.00,
"cloud-05": 1209.00,
"cloud-06": 1395.00,
"cloud-07": 1519.00,
"cloud-08": 1426.00,
"cloud-09": 1860.00,
"cloud-10": 1860.00
},
"2024-06-30": {
"total": 14040.00,
"cloud-01": 870.00,
"cloud-02": 1410.00,
"cloud-03": 1170.00,
"cloud-04": 1620.00,
"cloud-05": 1170.00,
"cloud-06": 1350.00,
"cloud-07": 1470.00,
"cloud-08": 1380.00,
"cloud-09": 1800.00,
"cloud-10": 1800.00
}
},
"totals": {
"cloud-01": 3538.00,
"cloud-02": 5734.00,
"cloud-03": 4758.00,
"cloud-04": 6588.00,
"cloud-05": 4758.00,
"cloud-06": 5490.00,
"cloud-07": 5978.00,
"cloud-08": 5612.00,
"cloud-09": 7320.00,
"cloud-10": 7320.00,
"total": 57096.00
},
"currency": "USD"
}
```

---

## CHAPTER

# FOUR

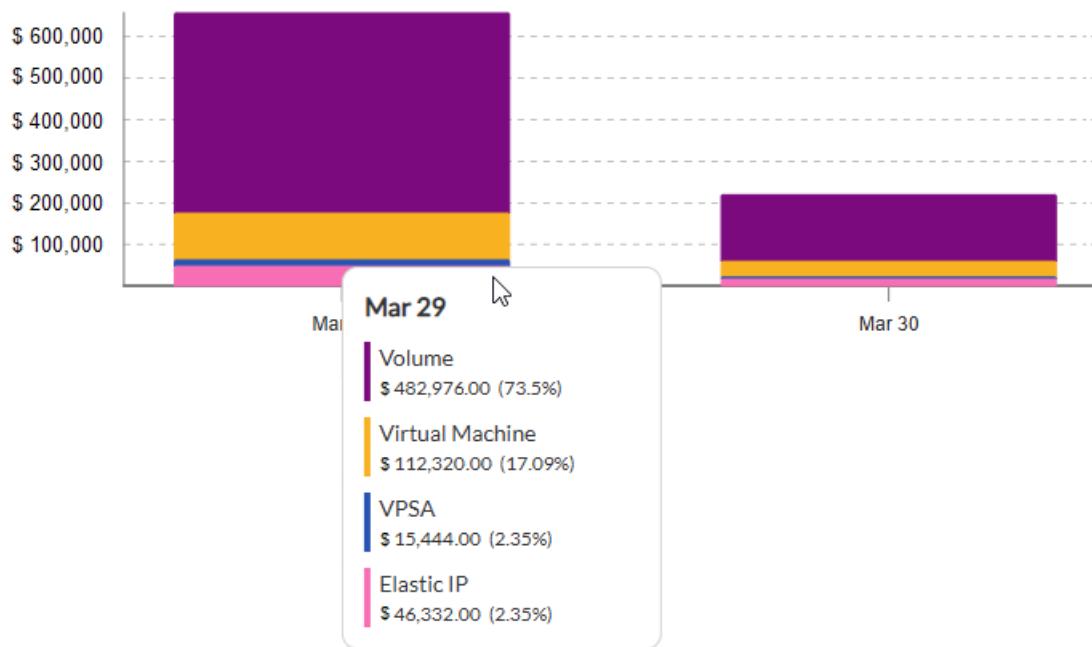
---

## EXPLORING COSTS AND USAGE

This section will walk you through the process of exploring costs and usage using CAU Explorer.

### 4.1 Charts

Charts are displayed in the upper section of the CAU Explorer, and provide visual representations of costs and usage patterns for the selected dates, and filters.



By exploring chart data, a visual review of a chart might rapidly indicate an anomaly or periodic usage trend.

For example, displaying daily data grouped by resource types might show that there are certain days each week that a bar on the chart is comparatively higher than its neighboring bars, showing a peak in cost and usage for a certain resource type on a particular day of the week. This pattern might reflect a known weekly usage-intensive activity.

#### 4.1.1 Structure

The CAU Explorer charts have a standard structure:

- X-axis: **Time** in summary units of days or months, according to the selected **Granularity** level.
- Y-axis: **Cost** or **Usage** scale in units of the configured currency.
- Data points: Hovering the mouse over a **Time** unit bar in the chart displays a pop-up detailing the individual **Cost** and **Usage** data points for each **Group By** item (if selected) of that time unit.

#### 4.1.2 Legends

Legends provide an interpretation of elements in a chart.

- The Y-axis is marked with currency unit steps indicating costs and usage levels, providing a visual comparison across time units.
- The X-axis is marked with the time unit's month or day beneath each bar in the chart.
- A chart with a **Group By** selection displays the proportion of each group's summary cost and usage as a separate color in stacked bars.
- A "data points" pop-up indicates the color used for each group item, names the group and displays its cost and usage in monetary value and percentage of the time unit's total.

#### 4.1.3 Currency

All costs and usages are displayed in the user's invoice currency, as configured in Zadara's billing system. In the event that costs in the CAU Explorer are not presented in the expected currency, please contact Zadara Support.

### 4.2 Tables

Tables are displayed in the lower section of the CAU Explorer.

Resource Type	Total Cost	Mar 29 2025	Mar 30 2025
All	<b>\$ 876,096.00</b>	<b>\$ 657,072.00</b>	<b>\$ 219,024.00</b>
Elastic IP	<b>\$ 61,776.00</b>	\$ 46,332.00	\$ 15,444.00
VPSA	<b>\$ 20,592.00</b>	\$ 15,444.00	\$ 5,148.00
Virtual Machine	<b>\$ 149,760.00</b>	\$ 112,320.00	\$ 37,440.00
Volume	<b>\$ 643,968.00</b>	\$ 482,976.00	\$ 160,992.00

Tables show the same information as the charts, presented in a tabular monetary layout. Subtotals of filtered costs and usage are listed per time unit of the selected period, along with a grand total. If **Group By** is selected, a further subtotal breakdown is also displayed per group in each time unit, along with a grand total for all time units of the group.

Similar to the exploration of chart data, a visual review of a table might rapidly indicate an anomaly or periodic usage trend.

Using the same example, displaying daily data grouped by resource types might show that there are certain days each week that the cost and usage amounts are comparatively higher than neighboring columns, showing a peak in cost and usage for a certain resource type on a particular day of the week. This pattern might reflect a known weekly usage-intensive activity.

### 4.2.1 Data structure

The tables in CAU follow a standard data structure:

- Rows: Represent individual cost or usage categories, for example, compute costs or storage costs.
- Columns: Represent time periods, for example, daily or monthly costs.

### 4.2.2 Layout

The table layout is designed for clarity and ease of use.

#### Main view table layout

- Each row represents a specific category, and each column represents a specific time period.
- Below the heading of each column, the topmost row displays a total that is the sum of all rows below it in the same column.
- If the **Group By** option is not selected, there is a single data row below the headings, and the leftmost column displays the sum of usages displayed in all other columns for the selected period.
- If a **Group By** option is selected, the leftmost column lists a member of the selected Group per row, and the second column displays the sum of that row's Group's usage for the selected period.

#### Drill-down view

Use the drill-down view to explore the breakdown of costs that accumulate to a total or subtotal in the main view:

1. Click a cell in the top-level *Main view table layout*.
2. Click a bar in the *Charts*.

#### ✓ Note

If a **Group By** option is applied, both the main table and the chart reflect the grouped values for the selected time period:

- The table displays one row per group value.
- The chart shows stacked bars that illustrate the proportions of each group.

Click a table cell or a segment of a stacked bar to drill down to reveal the detailed breakdown for the selected group and time period.

You can also explore the associated usage information, indicating how the costs were derived.

COST AND USAGE										
Details of <b>February 1st 2025 to February 10th 2025   Elastic IP</b>		Clouds	Resource Types	Items	Billable Types	Customers				
Report Date ↑	Cloud	Customer	Resource Type	Item	Compute Account Name					
□ After	□ Before	🔍 Search	🔍 Search	🔍 Search	🔍 Search	🔍 Search				
Rows per page: 50	1-50 of 1660	<	1	2	3	4	5	...	34	>

The drill-down view's caption heading indicates the scope of the view, using the following convention:

**Details of**

< Date range > | < Group By >

Details of  
<date range selection> | <Group By selection>

 **Note**

Reverting back to top level main view resets the drill-down view's custom settings to default.

---

Additional features in the detailed drill-down view further assist cost and usage exploration:

- *Estimation*
- *Setting the number of rows to display*
- *Adding and removing columns from the view*
- *Column order*
- *Sorting the table data*
- *Searching and filtering column values*
- *Download detailed drill-down view*

**Estimation**



If the daily scheduled process fails to populate the CAU database due to a network outage or another issue, the system calculates an estimate based on historical usage data.

Each week, the system attempts to replace estimated entries with actual usage data.

In the table, affected periods are shown with different shading and a circled question mark icon in the column heading.

When you hover over the shaded column in the table or the corresponding bar in the chart, a tooltip appears with the **Estimation** label.

## Setting the number of rows to display

Under the filters in the top filters bar, open the **Rows per page** dropdown to select the number of rows to display.

- Default: 50
- Minimum: 5
- Maximum: 100

## Adding and removing columns from the view

### Note

All columns are listed and described in *Drill-down detail table columns*.

Click the **Columns** button to the right of the top filters bar. Depending on the selected **Granularity** level (**Daily** or **Monthly**) the dropdown that opens displays:

- **Selected Columns**

Default columns are auto-selected for the drill-down detail table.

Most default columns are relevant for both **Granularity** levels (**Daily** or **Monthly**).

The default **Quantity**, **Hours** and **Amount** columns are specific to the selected **Granularity** level:

Column Name	Both	Monthly Only	Daily Only
Report Date	✓		
Cloud	✓		
Customer	✓		
Resource Type	✓		
Item	✓		
Compute Account Name	✓		
Unit Price	✓		
Discount	✓		
Billing UUID	✓		
Quantity MTD		✓	
Hours MTD		✓	
Net Amount MTD		✓	
Daily Quantity			✓
Daily Hours			✓
Daily Gross Amount			✓
Daily Net Amount			✓

Uncheck columns to remove them from the display.

- **Available Columns**

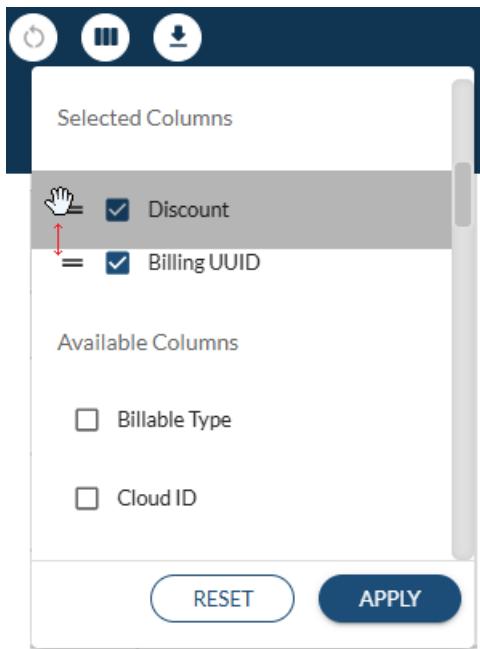
By default, many of the available columns are not displayed in the detail table.

To add a column to the display, mark it, and then in the **Selected Column** section, drag it up or down to the preferred location.

## Column order

There are two ways to rearrange the column display sequence:

- **Columns** filter:



1. Click the **Columns** button to the right of the top filters bar.
2. Drag one selected column at a time, up or down.

- **Detailed table:**

In the detailed table view, click and hold the column header, and drag the column left or right.

## Sorting the table data

To sort the table data according to the ascending or descending values of a selected column:

In the detailed table view, click a column header to display the table data sorted according to the sequence of the data in the selected column.

An arrow appears next to the column heading, indicating ascending order.

Click the heading again to reverse to descending order.

Click the heading a third time to cancel sorting by that column.

## Searching and filtering column values

Report Date	Cloud	Customer	Resource Type		Item	Compute Account Name	Unit Price
<input checked="" type="checkbox"/> After <input type="checkbox"/> Before	<input type="text"/> Search <input type="button" value="▼"/>	<input type="text"/> Search <input type="button" value="▼"/>	<input type="text"/> Search <input type="button" value="▲"/>	<input type="text"/> Search <input type="button" value="▼"/>	<input type="text"/> Search <input type="button" value="▼"/>	<input type="text"/> Search <input type="button" value="▼"/>	7
2025-02-01	Cloud1	Customer71	Elastic IP	Elastic IP	Item71	Account Name71	71.0
2025-02-01	Cloud1	Customer171	Virtual Mach	Virtual Machine ZIOS	Item171	Account Name171	171.0

- **Date** and **timestamp** type columns support filtering for entries **After** or **Before** a specified date and time, or between the two.
- **String** and **UUID** type columns support filtering, by entering a partial match in the **Search** field below the column heading.
- Numeric type columns are regarded as strings for the purpose of searching and filtering.

## Download detailed drill-down view

The CAU Explorer supports downloading a detailed drill-down view as a .csv file from the UI, or as either JSON or CSV output using the CLI.

Both the .csv and JSON outputs contain all of the *Drill-down detail table columns*.

This includes hidden columns that were not selected for display in the UI.

### ✓ Note

The maximum downloadable limit for export is 10,000 rows.

If downloading more than 10,000 rows is attempted:

- **UI:** An alert appears, advising narrowing down the scope.  
The number of rows in the current scope appears in the top filters bar below the **Download** button, after **Rows per page**.
- **CLI:** The response returns an error, advising narrowing down the scope.

### • Download detailed drill-down view to CSV (UI)

The drill-down view **Download CSV** button is to the right of the top filters bar.

Click **Download CSV** to export the detailed drill-down view, and save it as a file in .csv format. The scope of the saved data is according to the selected time period, granularity, grouping and filters.

The downloaded .csv file has the following naming convention:

CAU\_<organization>\_<from yyyy-mm-dd>-<to yyyy-mm-dd>\_<daily/monthly>\_<detail/  
→aggregated>.csv

- **Download detailed drill-down view to JSON or CSV (CLI)**

From the command line, you can use the Provisioning Portal's [GET Organization API](#) to download CAU Explorer detailed drill-down data in either JSON (default) or CSV format, and apply a custom time period, granularity, sorting and other optional filters.

Usage:

```
curl -X GET "https://<Zadara Provisioning Portal URL>/api/v2/organizations/
-><Organization UUID>/cost_explorer/details.json" \
-H "accept: application/json" \
-H "X-Token: <token>" \
--get --data-urlencode "granularity=<monthly/daily>" \
--data-urlencode "from=<start date>" \
--data-urlencode "to=<end date>" \
--data-urlencode "<optional parameter1>" \
--data-urlencode "<optional parameter2>"
```

## Parameters

The [GET Organization API](#) supports the following parameters:

Parameter	Type	Description	Mandatory
granularity	string	Granularity level Valid values: - monthly (default) - daily	✓
from	date	Report start date in yyyy-mm-dd format	✓
to	date	Report end date in yyyy-mm-dd format	✓
cloud	array of strings	Filter by cloud ID	
resource_type	array of strings	Filter by resource type	
item	array of strings	Filter by item	
billable_type	array of strings	Filter by billable type	
user	array of strings	Filter by user	
compute_account_name	array of strings	Filter by compute account	
sort_by	string	Name of column to sort by	
sort_direction	string	Order order of sorted values Valid values: - asc (default) - desc	
download_csv	string	Download file in JSON or CSV format Valid values: - false: JSON (default) - true: CSV	

Example request :

```
curl -X GET "manage.zadarastorage.com/api/v2/organizations/1a2b..36.char.UUID.
  -s string..x9z/cost_explorer/details.json" \
  -H "accept: application/json" \
  -H "X-Token: A1b2..token.string..yz" \
  --get --data-urlencode "granularity=monthly" \
  --data-urlencode "from=2025-03-01" \
  --data-urlencode "to=2025-05-12" \
  --data-urlencode "download_csv=false"
  --data-urlencode "billable_type[] = elastic_ip" \
  --data-urlencode "billable_type[] = volume"
```

Response example, displaying 2 of the total\_count of rows returned:

 **Note**

When metadata is returned in the response with the value "estimation": true, it indicates a temporary calculated *Estimation* in place of unpopulated values for the row.

```
{
  "status": "success",
  "total_count": 936,
  "data": [
    {
      "id": 3309950,
      "report_date": "2025-03-30",
      "org_id": "myOrg1",
      "parent_org": "myParentOrg1",
      "cloud": "myCloud1",
      "cloud_id": "4da21e44-66d4-4b1c-b8c3-abcdefg1234",
      "customer": "",
      "resource_type": "Elastic IP",
      "resource_id": "c5a62293-1896-4778-9e34-abcdefg1234",
      "resource_name": "172.221.32.11",
      "billable_type": "elastic_ip",
      "item": "Elastic IP - No Charge",
      "item_description": "Elastic IP - No Charge",
      "resource_created_at": "2024-11-08T15:53:33.000Z",
      "resource_deleted_at": null,
      "resource_status": "active",
      "compute_account_id": "123493a10ee44da98549c5a65def5678",
      "compute_account_name": "myAccount1",
      "compute_project_id": "1234454c6e294b5bacd5eb96f2155678",
      "compute_project_name": "myProj1",
      "compute_user_id": "",
      "compute_user_name": "",
      "vm_cpus": null,
      "vm_ram": null,
      "volume_disk_size": null,
      "vm_attached_to_uuid": "919f62b1-c189-4aa1-b421-abcdefg1234",
      "service_type": "",
      "compute_instance_type": "",
      "compute_instance_type_aws_alias": ""
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

"unit_price": "0.0",
"discount": "0.0",
"quantity_mtd": 906,
"quantity_est": 672,
"hours_mtd": 906,
"hours_est": 672,
"gross_amount_mtd": "0.0",
"gross_amount_est": "0.0",
"net_amount_mtd": "13573000.0",
"net_amount_est": "0.0",
"image_source": "",
"billing_uuid": null,
"created_at": "2025-03-31T15:49:02.119Z",
"updated_at": "2025-03-31T15:49:02.119Z",
"metadata": {
    "estimation": true
},
{
    "id": 3309951,
    "report_date": "2025-03-30",
    "org_id": "myOrg1",
    "parent_org": "myParentOrg1",
    "cloud": "myCloud1",
    "cloud_id": "4da21e44-66d4-4b1c-b8c3-abcdefg1234",
    "customer": "",
    "resource_type": "Elastic IP",
    "resource_id": "a465b79b-c982-4858-99b8-abcdefg1234",
    "resource_name": "172.221.32.110",
    "billable_type": "elastic_ip",
    "item": "Elastic IP - No Charge",
    "item_description": "Elastic IP - No Charge",
    "resource_created_at": "2023-09-09T13:38:40.000Z",
    "resource_deleted_at": null,
    "resource_status": "active",
    "compute_account_id": "12340e3ddb843979d8c782cb3c05a3c",
    "compute_account_name": "myAccount99",
    "compute_project_id": "1234c1cee6fe4eb890c73a4fc51f5678",
    "compute_project_name": "myProj99",
    "compute_user_id": "",
    "compute_user_name": "",
    "vm_cpus": null,
    "vm_ram": null,
    "volume_disk_size": null,
    "vm_attached_to_uuid": "",
    "service_type": "",
    "compute_instance_type": "",
    "compute_instance_type_aws_alias": "",
    "unit_price": "0.0",
    "discount": "0.0",
    "quantity_mtd": 906,
    "quantity_est": 672,
}

```

(continues on next page)

(continued from previous page)

```

"hours_mtd": 906,
"hours_est": 672,
"gross_amount_mtd": "0.0",
"gross_amount_est": "0.0",
"net_amount_mtd": "13574000.0",
"net_amount_est": "0.0",
"image_source": "",
"billing_uuid": null,
"created_at": "2025-03-31T15:49:02.119Z",
"updated_at": "2025-03-31T15:49:02.119Z",
"metadata": {
}
}
]
}

```

#### Drill-down detail table columns

Both daily and monthly reports include the following columns.

The default columns of the detailed drill-down view are indicated.

Column Name	Data Type	Default	Description
Report Date	date	✓	Date when the usage data was recorded.
Organization ID	UUID		Unique identifier for the organization.
Parent Organization	UUID		Identifier for the parent organization, if applicable.
Cloud	string	✓	Name of the cloud where the resource is hosted.
Cloud ID	UUID		Unique identifier for the cloud.
Customer	string	✓	Provisioning Portal username associated with the resource.
Resource Type	string	✓	Resource family name (e.g. Object Storage, VPSA, VM, Volume)
Resource ID	UUID		Unique identifier for the resource.
Resource Name	string		Resource display name.
Billable Type	string		Resource's billing code.
Item	string	✓	Specific item or service being billed.
Item Description	string		Description of the billed item or service.
Resource Created At	timestamp		UTC timestamp when the resource was created.
Resource Deleted At	timestamp		UTC timestamp when the resource was deleted, if applicable.
Resource Status	string		Current status of the resource (e.g., active, deleted).
Compute Account ID	UUID		Unique identifier for the zCompute account.
Compute Account Name	string	✓	zCompute account display name.
Compute Project ID	UUID		Unique identifier for the zCompute project.
Compute Project Name	string		zCompute project display name.
Compute User ID	UUID		Unique identifier for the zCompute user.
Compute User Name	string		zCompute username.
VM CPUs	integer		Number of virtual CPUs allocated to the VM.
VM RAM	integer		Amount of RAM allocated to the VM, in GiB.
Volume Disk Size	integer		zCompute EBS Disk Size (GB).
VM Attached To UUID	UUID		UUID of the VM to which the volume is attached.
Service Type	string		Type of service provided (e.g., compute, storage).

continues on next page

Table 1 – continued from previous page

Column Name	Data Type	Default	Description
Compute Instance Type	string		zCompute instance type.
Compute Instance Type AWS Alias	string		AWS-equivalent instance type alias.
Unit Price	decimal	✓	Price per unit of the billed item.
Discount	decimal	✓	Discount percentage applied to the unit price, if any.
Quantity MTD	integer		Quantity of units used month-to-date.
Quantity EST	integer		Estimated quantity to be used until the end of the month.
Hours MTD	integer		Number of hours since the beginning of the month.
Hours EST	integer		Estimated number of hours in the current billing cycle.
Gross Amount MTD	decimal		Gross amount month-to-date (Price * Qty Units * Hours).
Gross Amount EST	decimal		Estimated gross bill amount (Price * Est Qty * Est Hours).
Net Amount MTD	decimal		Net month-to-date billable amount after discounts.
Net Amount EST	decimal		Estimated net amount to be billed after discounts.
Image Source	UUID		UUID of zCompute image source used for the VM.
Billing UUID	UUID	✓	Unique identifier for the billing record.

When daily granularity is selected, daily reports also include the following columns in the drill-down view:

Column Name	Data Type	De-fault	Description
Daily Quantity	integer	✓	Calculated: Yesterday's Qty MTD subtracted from Today's Qty MTD.
Daily Hours	integer	✓	Calculated: Yesterday's Hours MTD subtracted from Today's Hours MTD.
Daily Gross Amount	decimal	✓	Calculated: Yesterday's Gross Amount MTD subtracted from Today's Gross Amount MTD.
Daily Net Amount	decimal	✓	Calculated: Yesterday's Net Amount MTD subtracted from Today's Net Amount MTD.