



PUBLIC
2025-02-04

Application Autoscaler

Content

- 1 What Is Application Autoscaler? 3**
- 2 Initial Setup. 5**
- 3 Development. 7**
 - 3.1 Working with the Application Autoscaler Dashboard. 7
 - Access the Application Autoscaler Dashboard. 7
 - Use the Application Autoscaler Dashboard. 8
 - 3.2 Defining a Scaling Policy. 9
 - Dynamic Scaling Policy. 9
 - Schedule-Based Scaling Policy. 20
 - 3.3 Defining a Custom Metric. 34
- 4 Administration. 39**
 - 4.1 Preventing the Application Autoscaler from Scaling your Application. 39
- 5 Accessibility Features in Application Autoscaler. 41**



1 What Is Application Autoscaler?

Automatically scale your applications to meet their dynamic resource needs.

Application Autoscaler lets you automatically increase or decrease the number of your application instances based on the policies you have defined.



→ Tip

The English version of this guide is open for contributions and feedback using GitHub. This allows you to get in contact with responsible authors of SAP Help Portal pages and the development team to discuss documentation-related issues. To contribute to this guide, or to provide feedback, choose the corresponding option on SAP Help Portal:

- [Feedback](#) > [Create issue](#) : Provide feedback about a documentation page. This option opens an issue on GitHub.
- [Feedback](#) > [Edit page](#) : Contribute to a documentation page. This option opens a pull request on GitHub.

You need a GitHub account to use these options.

More information:

- [Contribution Guidelines](#)
- [Introduction Video](#) 
- [Introduction Blog Post](#) 

Features

Scale your applications automatically

Automatically increase or decrease the number of application instances to meet the changing demands of application resources.

Make your scaling resource-specific

Scale your applications on the basis of any standard metric type, such as memory consumed, CPU, response time, or throughput.

Define custom metrics

Scale your applications by defining custom metrics that match their requirements.

Manage your costs effectively

Only consume and pay for the compute resources, which your applications require.

Environment

This service runs in the Cloud Foundry environment.

Regional Availability

For information on the regional availability of Application Autoscaler, see [SAP Discovery Center](#).

For more information, see [Regions](#).

Documentation Feedback Survey

[Give Feedback](#)

Help us improve our documentation. Please take this two-minute survey and let us know what you think.

2 Initial Setup

Create an instance of the Application Autoscaler service and bind it to your application.

Prerequisites

Note

If you are using this service as part of SAP Build Code, follow the [SAP Build Code Initial Setup](#) instructions instead.

- You have downloaded and set up the Cloud Foundry Command Line Interface (cf CLI). See [Download and Install the Cloud Foundry Command Line Interface](#).
- You are logged on to your Cloud Foundry space. See [Log On to the Cloud Foundry Environment Using the Cloud Foundry Command Line Interface](#).
- You are a Space Developer in the chosen Cloud Foundry space.
- You have deployed the application you want to scale in the Cloud Foundry environment. See [Deploy Business Applications in the Cloud Foundry Environment](#).

Procedure

1. Check if the Application Autoscaler service is listed in the Service Marketplace using the following command:

```
cf marketplace
```

2. Create an instance of the service using the following command:

```
cf create-service autoscaler <service plan name> <instance name>
```

Sample Code

```
cf create-service autoscaler standard myservice
```

3. Bind the service instance to your application using the following command. You can bind as many applications as you want to an Application Autoscaler service instance.

```
cf bind-service <application name> <instance name> -c <file name>.json
```

The JSON file contains the policy, which is needed to initiate the scaling of an application. For more information about how to define a policy, see [Defining a Scaling Policy \[page 9\]](#).

ⓘ Note

If you are building a multitarget application (MTA), you can leverage the MTA deployment descriptor to provide the scaling policy. See [Service Binding Parameters](#).

4. Check if the instance is successfully bound to the application using the following command:

```
cf service <instance name>
```

The *Bound apps* field displays all applications bound to the instance.

3 Development

Use the Application Autoscaler in your development.

Using the Application Autoscaler in your development comprises the following steps:

- [Working with the Application Autoscaler Dashboard \[page 7\]](#)
Access and use the Application Autoscaler dashboard.
- [Defining a Scaling Policy \[page 9\]](#)
Define a policy to scale your application instances either dynamically or based on schedules.
- [Defining a Custom Metric \[page 34\]](#)
Define your own metrics to scale applications based on your requirements.

3.1 Working with the Application Autoscaler Dashboard

Access and use the Application Autoscaler dashboard.

3.1.1 Access the Application Autoscaler Dashboard

Access the Application Autoscaler dashboard, which lets you create or update a scaling policy, and view the scaling details.

Prerequisites

- You have created an Application Autoscaler service instance and bound it to your application. See [Initial Setup \[page 5\]](#).
- You are a Space Developer in the chosen Cloud Foundry space.

Procedure

1. In the SAP BTP cockpit, navigate to the **Organization** > **Space** containing the Application Autoscaler.
2. From the navigation pane, choose [Service Marketplace](#).
3. Choose [Application Autoscaler](#).
4. In the navigation pane, choose [Instances](#).

The list of created instances appears.

5. From the [Actions](#) column, choose (Open Dashboard).
6. Provide the platform credentials to access the dashboard.

The [Overview](#) page appears.

7. In the navigation pane, choose [Referencing Apps](#).

The list of applications bound to the service instance appears. In the [Actions](#) column, the [Manage Policy](#) and [Scaling History](#) options are available.

3.1.2 Use the Application Autoscaler Dashboard

Manage your scaling policy and view the scaling history.

The Application Autoscaler dashboard comprises two features, which let you perform different actions: [Manage Policy](#) and [Scaling History](#).

Manage Policy

A scaling policy defines the scaling requirements of an application. [Manage Policy](#) lets you perform the following tasks:

- [Create a Policy](#)
Define scaling requirements if you haven't used configuration parameters with the bind request during the binding of the service with an application.
- [View a Policy](#)
Preview your policy, especially if you've defined it during the binding process.
- [Edit a Policy](#)
Change an existing policy, for example to:
 - Add more scaling rules
 - Add schedule-based scaling along with dynamic scaling
 - Update scaling parameters like cool down and threshold adjustment based on their behavior as observed in the [Scaling History](#) feature.

For information about scaling policy, see [Dynamic Scaling Policy \[page 9\]](#).

Scaling History

This feature lists all scaling events that were triggered by the Application Autoscaler service instance after it was bound to an application. The scaling history helps you to troubleshoot and resolve issues if the scaling doesn't match the defined policy.

❖ Example

If the scaling failed due to quota restrictions, an appropriate message is displayed in the scaling history. Accordingly, you can request additional quota.

3.2 Defining a Scaling Policy

Define a policy to scale your application instances either dynamically or based on schedules.

To initiate the scaling of an application, you need to define a policy. A policy is a JSON file containing an array of rules or a single rule for scaling. This JSON file is passed while binding the service with the application to be scaled. For more information, see [Initial Setup \[page 5\]](#).

You can choose between the following types of scaling policies:

- [Dynamic Scaling Policy \[page 9\]](#)
Scale your application instances based on memory/CPU usage, response time, throughput, or custom metrics.
- [Schedule-Based Scaling Policy \[page 20\]](#)
Scale your application instances based on schedules.

Related Information

[Parameters for a Dynamic Scaling Policy \[page 12\]](#)

[Parameters for a Schedule-Based Scaling Policy \[page 23\]](#)

[Time Zones for a Schedule-Based Policy \[page 25\]](#)

3.2.1 Dynamic Scaling Policy

Scale your application instances based on memory or CPU usage, response time, throughput, or custom metrics.

Note

Both built-in and custom metrics are collected from each application instance individually. Then, their average for all application instances is determined. Based on these average metrics, scaling rules are evaluated once per app.

Metric Types: memory consumed and memory utilized

These scaling rules define that if the memory consumption either exceeds or falls below the set threshold, the application is respectively scaled up or down by an instance. If an array of scaling rules is defined, the Application Autoscaler runs periodic checks of the scaling policy to verify if any of the specified conditions matches. New instances are generated based on specified rules.

Example

A scaling policy comprises two rules:

- Rule 1: If the memory exceeds 500 MB, a single instance of an application is created.
- Rule 2: If the memory exceeds 800 MB, two application instances are created.

When defining an array of rules, its sequence must match the descending order of the threshold values, so that the higher value appears first in the sequence. In this example, rule 2 must be defined before rule 1 to ensure that the Application Autoscaler checks all conditions.

The following table provides examples for dynamic scaling policies based on the metric types **memory consumed** and **memory utilized**.

Metric Type memory consumed

↔ Sample Code

```
{
  "instance_min_count": 1,
  "instance_max_count": 5,
  "scaling_rules": [
    {
      "metric_type":
"memoryused",
      "threshold": 90,
      "operator": ">=",
      "adjustment": "+1"
    },
    {
      "metric_type":
"memoryused",
      "threshold": 30,
      "operator": "<",
      "adjustment": "-1"
    }
  ]
}
```

Metric Type memory utilized

↔ Sample Code

```
{
  "instance_min_count": 1,
  "instance_max_count": 5,
  "scaling_rules": [
    {
      "metric_type":
"memoryutil",
      "threshold": 90,
      "operator": ">=",
      "adjustment": "+1"
    },
    {
      "metric_type":
"memoryutil",
      "threshold": 30,
      "operator": "<",
      "adjustment": "-1"
    }
  ]
}
```

Metric Types: throughput and responsetime

The following samples are simple scaling policies for the metric types **throughput** and **responsetime**. The scaling rule determines if the throughput or response time exceeds or falls below the set threshold.

↔ Sample Code

```

{
  "instance_min_count": 1,
  "instance_max_count": 5,
  "scaling_rules": [
    {
      "metric_type":
"throughput",
      "threshold": 100, //
the throughput value in Requests/
second
      "operator": ">",
      "adjustment": "+1"
    },
    {
      "metric_type":
"throughput",
      "threshold": 100,
      "operator": "<=",
      "adjustment": "-1"
    }
  ]
}

```

↔ Sample Code

```

{
  "instance_min_count": 1,
  "instance_max_count": 5,
  "scaling_rules": [
    {
      "metric_type":
"responsetime",
      "threshold": 900, //
the response time value in
milliseconds
      "operator": ">",
      "adjustment": "+1"
    },
    {
      "metric_type":
"responsetime",
      "threshold": 900,
      "operator": "<=",
      "adjustment": "-1"
    }
  ]
}

```

Advanced Usage

For advanced usage, you can provide optional values along with the mandatory ones. The following sample displays such a scaling policy:

↔ Sample Code

```

{
  "instance_min_count": 1,
  "instance_max_count": 5,
  "scaling_rules": [
    {
      "metric_type": "memoryused",
      "breach_duration_secs": 600,
      "threshold": 90,
      "operator": ">=",
      "cool_down_secs": 300,
      "adjustment": "+1"
    },
    {
      "metric_type": "memoryused",
      "breach_duration_secs": 600,
      "threshold": 30,
      "operator": "<",
      "cool_down_secs": 300,
      "adjustment": "-1"
    }
  ]
}

```

For more information about the parameters used, see [Parameters for a Dynamic Scaling Policy \[page 12\]](#).

Related Information

[Parameters for a Dynamic Scaling Policy \[page 12\]](#)

3.2.1.1 Parameters for a Dynamic Scaling Policy

Get to know the parameters used for dynamic scaling.

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
<code>instance_min_count</code>	The minimum number of application instances that are always running.	Yes	number	minimum = 1; maximum = no upper limit	none	1
<code>instance_max_count</code>	The maximum number of application instances that can be provisioned as part of application scaling.	Yes	number	minimum = <code>instance_min_count</code> ; maximum = no upper limit	none	5
<code>scaling_rules</code>	A rule comprises a group of key values set to automatically trigger the scaling activity. You can have one or more rules. At least one rule out of a possible array size of two is validated.	Yes	array	none	none	none

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
<code>metric_type</code>	<p>The metric type can be one of the following options:</p> <ul style="list-style-type: none"> memory usage in mebibytes (<code>memoryused</code>) memory utilization as a percentage of the memory quota (<code>memoryutil</code>) CPU as a percentage of virtual CPUs used (<code>cpu</code>) 	Yes	string	none	none	<p>One of the following options:</p> <ul style="list-style-type: none"> <code>memoryused</code> <code>memoryutil</code> <code>cpu</code> <code>cpuutil</code> <code>disk</code> <code>diskutil</code> <code>throughput</code> <code>response</code> <code>etime</code>

→ Tip

Instead, we recommend to use the CPU entitlement utilization metric (`cpuutil`), which eliminates

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
	<p>the need to calculate the CPU entitlement for your application.</p> <ul style="list-style-type: none"> • CPU entitlement utilization as percentage (cpuutil) • disk usage in mebibytes (disk) • disk utilization as a percentage of the disk quota (diskutil) • throughput in requests per second (throughput) • response time in milliseconds (responseime) 					

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
breach_duration_secs	<p>The duration to analyze the collected metrics data points.</p> <p>The duration is considered from the current time to the past, for example: the last 600 seconds.</p> <p>The service performs this analysis to check if the data points are consistently above or below the set threshold and makes the decision to scale. For a more accurate decision, set a larger duration, as in this case, more data points are analyzed.</p> <p>In situations that cause a rapid increase of application load in short intervals, set a small duration. This enables the application to scale out before the maximum memory threshold is exceeded and prevents the application from crashing.</p>	No	number	minimum = 60 seconds; maximum = 3600 seconds	300 seconds	300 seconds

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
threshold	<p>The threshold is compared to one of the following options:</p> <ul style="list-style-type: none"> memory usage in mebibytes (metric type <code>memoryused</code>) memory utilization as a percentage of the memory quota (metric type <code>memoryutil</code>) CPU usage as a percentage of virtual CPUs used (metric type <code>cpu</code>) CPU entitlement utilization as percentage (metric type <code>cpuutil</code>) disk usage in mebibytes (metric type <code>disk</code>) disk utilization as a percentage 	Yes	integer	<ul style="list-style-type: none"> <code>memoryused</code>: minimum = 1; maximum = no upper limit <code>memoryutil</code>: minimum = 1; maximum = 100 <code>cpu</code>: minimum = 1; maximum = 400 <code>cpuutil</code>: minimum = 1; maximum = 100 <code>disk</code>: minimum = 1; maximum = 10240 <code>diskutil</code>: minimum = 1; maximum = 100 <code>throughput</code>: minimum = 1; maximum = no upper limit <code>response time</code>: minimum = 1; maximum = no upper limit 	none	30

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
	<p>of the disk quota (metric type <code>diskutil</code>)</p> <ul style="list-style-type: none"> throughput in requests per second (metric type <code>throughput</code>) response time in milliseconds (metric type <code>response_time</code>) 					
<code>operator</code>	<p>The operator is used in combination with the threshold value to compare the current metric value. The expression is as follows:</p> <pre>[MetricUsage value] [operator] [threshold].</pre>	Yes	string	<p>< > <= >=</p>	none	<code>64 > 30</code>

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
cool_down_secs	The minimum duration between two successive scaling triggers. After the first scaling trigger, the second trigger occurs after the specified interval. This duration enables the application to be stable enough before it starts the next trigger. For production environments, a longer duration is recommended.	No	number	minimum = 60 seconds; maximum = 3600 seconds	300 seconds	300

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
adjustment	<p>The number of application instances to be scaled as part of the scaling activity. The format is + or – followed by one of the following:</p> <ul style="list-style-type: none"> the number of required instances the number of required instances followed by a percentage sign (%) <p>The number is then divided by 100 and multiplied with the number of application instances currently running to determine the final adjustment amount.</p>	Yes	number	none	none	<ul style="list-style-type: none"> +1 to scale up by an instance; –2 to scale down by two instances +100% to double the number of instances; –50% to halve the number of instances

Related Information

[Dynamic Scaling Policy \[page 9\]](#)

3.2.2 Schedule-Based Scaling Policy

Scale your application instances based on schedules.

You can either define recurring schedules or specific date schedules for scaling. If you want to define multiple schedules, make sure that they don't overlap.

⚠ Caution

Overlapping schedules don't trigger the scaling.

In recurring schedules, the date is optional. However, you need to provide a date for a specific date schedule. If you set the parameters `instance_min_count` and `instance_max_count` within a schedule, they override the global values set for these parameters.

Recurring Schedules

The following example shows a schedule-based scaling policy:

📄 Sample Code

```
{
  "instance_min_count": 1,
  "instance_max_count": 5,
  "schedules": {
    "timezone": "Asia/Shanghai",
    "recurring_schedule": [
      {
        "start_time": "10:00",
        "end_time": "18:00",
        "days_of_week": [
          1,
          2,
          3
        ],
        "instance_min_count": 1,
        "instance_max_count": 10,
        "initial_min_instance_count": 5
      }
    ]
  }
}
```

The policy in the example defines a recurring schedule, which is set for the Asia/Shanghai time zone. The schedule triggers application instances during the specified time frame on Sunday, Monday, and Tuesday. The days of the week are specified through the numbers 1 to 7 starting with 1 for Monday.

The policy defines the minimum and maximum number of application instances during the time frame as well as the number of application instances that should be available at the beginning of the recurring schedule.

For more information about the time zones for schedule-based policies, see [Time Zones for a Schedule-Based Policy \[page 25\]](#).

Schedules with Start Day and Day of the Month, and Specific Date Schedules

The following examples show schedule-based scaling with start date and day of the month, as well as a specific date schedule.

Schedule with Start Date and Day of the Month

Sample Code

```
"schedules": {
  "timezone": "Asia/Shanghai",
  "recurring_schedule": [
    {
      "start_date": "2016-06-27",
      "end_date": "2016-07-23",
      "start_time": "11:00",
      "end_time": "19:30",
      "days_of_month": [
        5,
        15,
        25
      ],
      "instance_min_count": 3,
      "instance_max_count": 10,
    }
  ],
  "initial_min_instance_count": 5
}
```

Specific Date Schedule

Sample Code

```
{
  "instance_min_count": 1,
  "instance_max_count": 5,
  "schedules": {
    "timezone": "Asia/Shanghai",
    "specific_date": [
      {
        "start_date_time": "2015-06-02T10:00",
        "end_date_time": "2015-06-15T13:59",
        "instance_min_count": 1,
        "instance_max_count": 4,
        "initial_min_instance_count": 2
      },
      {
        "start_date_time": "2015-01-04T20:00",
        "end_date_time": "2015-02-19T23:15",
        "instance_min_count": 2,
        "instance_max_count": 5,
        "initial_min_instance_count": 3
      }
    ]
  }
}
```

Advanced Usage

For advanced usage, you can provide optional values along with the mandatory ones. The following sample displays such a scaling policy:

Sample Code

```
{
  "instance_min_count": 1,
  "instance_max_count": 5,
  "schedules": {
```

```

"timezone": "Asia/Shanghai",
"recurring_schedule": [
  {
    "start_time": "10:00",
    "end_time": "18:00",
    "days_of_week": [
      1,
      2,
      3
    ],
    "instance_min_count": 1,
    "instance_max_count": 10,
    "initial_min_instance_count": 5
  },
  {
    "start_date": "2016-06-27",
    "end_date": "2016-07-23",
    "start_time": "11:00",
    "end_time": "19:30",
    "days_of_month": [
      5,
      15,
      25
    ],
    "instance_min_count": 3,
    "instance_max_count": 10,
    "initial_min_instance_count": 5
  }
]
}

```

For more information about the parameters used, see [Parameters for a Schedule-Based Scaling Policy \[page 23\]](#).

Related Information

[Parameters for a Schedule-Based Scaling Policy \[page 23\]](#)

[Time Zones for a Schedule-Based Policy \[page 25\]](#)

3.2.2.1 Parameters for a Schedule-Based Scaling Policy

Get to know the parameters used for recurring schedules or specific date schedules.

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
<code>schedules</code>	A schedule lets you configure scaling rules for specific days or on a recurring basis. Schedules guard against expected high surges or low activity periods.	Yes	none	none	none	none
<code>timezone</code>	A valid time zone to run the schedule. For more information, see Time Zones for a Schedule-Based Policy [page 25] .	Yes	string	none	none	Asia/Sanghai
<code>recurring_schedule</code>	Triggers the scaling rule recursively during the specified intervals	Yes	none	none	none	none
<code>start_time</code>	Start time of a recurring schedule in 24-hr format (HH:MM).	Yes	string	minimum = 00:00; maximum = 23:59	none	10:00
<code>end_time</code>	End time of a recurring schedule in 24-hr format (HH:MM).	Yes	string	minimum = 00:00; maximum = 23:59	none	21:00

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
start_date	Start date of the schedule in YYYY-MM-DD format.	No	string	minimum = current date; maximum = no upper limit	none	2020-06-27
end_date	End date of the schedule in YYYY-MM-DD format.	No	string	minimum = set by start_date; maximum = no upper limit	none	2020-07-23
days_of_week	Triggers the scaling on weekdays ranging from 1 (Monday) to 7 (Sunday). The rule is executed during the weekdays specified within the array.	Yes	array	minimum = 1; maximum = 7	none	[1 , 3 , 5]
days_of_month	Triggers the scaling on days of the month ranging from 1 to 31. The rule is executed during the specified days of the month.	Yes	array	minimum = 1; maximum = 31	none	[1 , 11 , 24 , 30]
instance_min_count	Minimum number of instances during the recurrence period.	Yes	number	minimum = 1; maximum = no upper limit	none	1
instance_max_count	Maximum number of instances during the recurrence period.	Yes	number	minimum = set by instance_min_count; maximum = no upper limit	none	5

Parameter	Description	Mandatory	Data Type	Value Range	Default Value	Example
<code>initial_minimum_instance_count</code>	Minimum number of instances to scale up at the beginning of the recurrence period.	Yes	number	minimum = set by <code>initial_minimum_instance_count</code> ; maximum = <code>initial_maximum_instance_count</code>	none	5
<code>specific_date</code>	Triggers the scaling rule during the specified date intervals.	Yes	none	none	none	none
<code>start_date_time</code>	Start date and time of the schedule in YYYY-MM-DDTHH:MM format	Yes	string	minimum = current date and time; maximum = no upper limit	none	2015-06-02T10:00
<code>end_date_time</code>	End date and time of the schedule in YYYY-MM-DDTHH:MM format	Yes	string	minimum = set by <code>start_date_time</code> ; maximum = no upper limit	none	2015-09-02T10:00

Related Information

[Schedule-Based Scaling Policy \[page 20\]](#)

[Time Zones for a Schedule-Based Policy \[page 25\]](#)

3.2.2.2 Time Zones for a Schedule-Based Policy

Get to know the time zones supported for recurring schedules or specific date schedules.

The following code block lists all time zones supported by the Application Autoscaler.

```
{
  "Etc/GMT+12",
  "Etc/GMT+11",
  "Pacific/Midway",
  "Pacific/Niue",
```

```
"Pacific/Pago_Pago",
"Pacific/Samoa",
"US/Samoa",
"Etc/GMT+10",
"HST",
"Pacific/Honolulu",
"Pacific/Johnston",
"Pacific/Rarotonga",
"Pacific/Tahiti",
"US/Hawaii",
"Pacific/Marquesas",
"America/Adak",
"America/Atka",
"Etc/GMT+9",
"Pacific/Gambier",
"US/Aleutian",
"America/Anchorage",
"America/Juneau",
"America/Metlakatla",
"America/Nome",
"America/Sitka",
"America/Yakutat",
"Etc/GMT+8",
"Pacific/Pitcairn",
"US/Alaska",
"America/Creston",
"America/Dawson",
"America/Dawson_Creek",
"America/Ensenada",
"America/Hermosillo",
"America/Los_Angeles",
"America/Phoenix",
"America/Santa_Isabel",
"America/Tijuana",
"America/Vancouver",
"America/Whitehorse",
"Canada/Pacific",
"Canada/Yukon",
"Etc/GMT+7",
"MST",
"Mexico/BajaNorte",
"PST8PDT",
"US/Arizona",
"US/Pacific",
"US/Pacific-New",
"America/Belize",
"America/Boise",
"America/Cambridge_Bay",
"America/Chihuahua",
"America/Costa_Rica",
"America/Denver",
"America/Edmonton",
"America/El_Salvador",
"America/Guatemala",
"America/Inuvik",
"America/Managua",
"America/Mazatlan",
"America/Ojinaga",
"America/Regina",
"America/Shiprock",
"America/Swift_Current",
"America/Tegucigalpa",
"America/Yellowknife",
"Canada/East-Saskatchewan",
"Canada/Mountain",
"Canada/Saskatchewan",
"Etc/GMT+6",
"MST7MDT",
```

"Mexico/BajaSur",
"Navajo",
"Pacific/Galapagos",
"US/Mountain",
"America/Atikokan",
"America/Bahia_Banderas",
"America/Bogota",
"America/Cancun",
"America/Cayman",
"America/Chicago",
"America/Coral_Harbour",
"America/Eirunepe",
"America/Guayaquil",
"America/Indiana/Knox",
"America/Indiana/Tell_City",
"America/Jamaica",
"America/Knox_IN",
"America/Lima",
"America/Matamoros",
"America/Menominee",
"America/Merida",
"America/Mexico_City",
"America/Monterrey",
"America/North_Dakota/Beulah",
"America/North_Dakota/Center",
"America/North_Dakota/New_Salem",
"America/Panama",
"America/Porto_Acre",
"America/Rainy_River",
"America/Rankin_Inlet",
"America/Resolute",
"America/Rio_Branco",
"America/Winnipeg",
"Brazil/Acre",
"CST6CDT",
"Canada/Central",
"Chile/EasterIsland",
"EST",
"Etc/GMT+5",
"Jamaica",
"Mexico/General",
"Pacific/Easter",
"US/Central",
"US/Indiana-Starke",
"America/Caracas",
"America/Anguilla",
"America/Antigua",
"America/Aruba",
"America/Asuncion",
"America/Barbados",
"America/Blanc-Sablon",
"America/Boa_Vista",
"America/Campo_Grande",
"America/Cuiaba",
"America/Curacao",
"America/Detroit",
"America/Dominica",
"America/Fort_Wayne",
"America/Grand_Turk",
"America/Grenada",
"America/Guadeloupe",
"America/Guyana",
"America/Havana",
"America/Indiana/Indianapolis",
"America/Indiana/Marengo",
"America/Indiana/Petersburg",
"America/Indiana/Vevay",
"America/Indiana/Vincennes",

```

"America/Indiana/Winamac",
"America/Indianapolis",
"America/Iqaluit ",
"America/Kentucky/Louisville ",
"America/Kentucky/Monticello",
"America/Kralendijk",
"America/La_Paz",
"America/Louisville ",
"America/Lower_Princes",
"America/Manaus",
"America/Marigot",
"America/Martinique",
"America/Montreal",
"America/Montserrat",
"America/Nassau",
"America/New_York",
"America/Nipigon",
"America/Pangnirtung ",
"America/Port-au-Prince ",
"America/Port_of_Spain",
"America/Porto_Velho",
"America/Puerto_Rico ",
"America/Santo_Domingo ",
"America/St_Barthelemy",
"America/St_Kitts",
"America/St_Lucia",
"America/St_Thomas",
"America/St_Vincent",
"America/Thunder_Bay",
"America/Toronto",
"America/Tortola",
"America/Virgin",
"Brazil/West",
"Canada/Eastern",
"Cuba",
"EST5EDT",
"Etc/GMT+4",
"US/East-Indiana",
"US/Eastern",
"US/Michigan",
"America/Araguaina ",
"America/Argentina/Buenos_Aires ",
"America/Argentina/Catamarca ",
"America/Argentina/ComodRivadavia ",
"America/Argentina/Cordoba ",
"America/Argentina/Jujuy ",
"America/Argentina/La_Rioja ",
"America/Argentina/Mendoza ",
"America/Argentina/Rio_Gallegos ",
"America/Argentina/Salta ",
"America/Argentina/San_Juan ",
"America/Argentina/San_Luis ",
"America/Argentina/Tucuman ",
"America/Argentina/Ushuaia",
"America/Bahia",
"America/Belem",
"America/Buenos_Aires",
"America/Catamarca",
"America/Cayenne",
"America/Cordoba",
"America/Fortaleza",
"America/Glace_Bay",
"America/Goose_Bay",
"America/Halifax",
"America/Jujuy",
"America/Maceio",
"America/Mendoza",
"America/Moncton",

```

```
"America/Montevideo",
"America/Paramaribo",
"America/Recife",
"America/Rosario",
"America/Santarem",
"America/Santiago",
"America/Sao_Paulo",
"America/Thule",
"Antarctica/Palmer",
"Antarctica/Rothera",
"Atlantic/Bermuda",
"Atlantic/Stanley",
"Brazil/East",
"Canada/Atlantic",
"Chile/Continental",
"Etc/GMT+3",
"America/St_Johns",
"Canada/Newfoundland",
"America/Godthab",
"America/Miquelon",
"America/Noronha ",
"Atlantic/South_Georgia",
"Brazil/DeNoronha",
"Etc/GMT+2",
"Atlantic/Cape_Verde",
"Etc/GMT+1",
"Africa/Abidjan",
"Africa/Accra",
"Africa/Bamako",
"Africa/Banjul",
"Africa/Bissau",
"Africa/Conakry",
"Africa/Dakar",
"Africa/Freetown",
"Africa/Lome",
"Africa/Monrovia",
"Africa/Nouakchott",
"Africa/Ouagadougou",
"Africa/Sao_Tome",
"Africa/Timbuktu",
"America/Danmarkshavn",
"America/Scoresbysund",
"Atlantic/Azores",
"Atlantic/Reykjavik",
"Atlantic/St_Helena",
"Etc/GMT",
"Etc/GMT+0",
"Etc/GMT-0",
"Etc/GMT0",
"Etc/Greenwich",
"Etc/UCT",
"Etc/UTC",
"Etc/Universal",
"Etc/Zulu",
"GMT",
"GMT+0",
"GMT-0",
"GMT0",
"Greenwich",
"Iceland",
"UCT",
"UTC",
"Universal",
"Zulu",
"Africa/Algiers",
"Africa/Bangui",
"Africa/Brazzaville",
"Africa/Casablanca",
```

```
"Africa/Douala",
"Africa/El_Aaiun",
"Africa/Kinshasa",
"Africa/Lagos",
"Africa/Libreville",
"Africa/Luanda",
"Africa/Malabo",
"Africa/Ndjamena",
"Africa/Niamey",
"Africa/Porto-Novo",
"Africa/Tunis",
"Africa/Windhoek",
"Atlantic/Canary",
"Atlantic/Faeroe",
"Atlantic/Faroe",
"Atlantic/Madeira",
"Eire",
"Etc/GMT-1",
"Europe/Belfast",
"Europe/Dublin",
"Europe/Guernsey",
"Europe/Isle_of_Man",
"Europe/Jersey",
"Europe/Lisbon",
"Europe/London",
"GB",
"GB-Eire",
"Portugal",
"WET",
"Africa/Blantyre",
"Africa/Bujumbura",
"Africa/Cairo",
"Africa/Ceuta",
"Africa/Gaborone",
"Africa/Harare",
"Africa/Johannesburg",
"Africa/Kigali",
"Africa/Lubumbashi",
"Africa/Lusaka",
"Africa/Maputo",
"Africa/Maseru",
"Africa/Mbabane",
"Africa/Tripoli",
"Antarctica/Troll",
"Arctic/Longyearbyen",
"Atlantic/Jan_Mayen",
"CET",
"Egypt",
"Etc/GMT-2",
"Europe/Amsterdam",
"Europe/Andorra",
"Europe/Belgrade",
"Europe/Berlin",
"Europe/Bratislava",
"Europe/Brussels",
"Europe/Budapest",
"Europe/Busingen",
"Europe/Copenhagen",
"Europe/Gibraltar",
"Europe/Kaliningrad",
"Europe/Ljubljana",
"Europe/Luxembourg",
"Europe/Madrid",
"Europe/Malta",
"Europe/Monaco",
"Europe/Oslo",
"Europe/Paris",
"Europe/Podgorica",
```

```
"Europe/Prague",
"Europe/Rome",
"Europe/San_Marino",
"Europe/Sarajevo",
"Europe/Skopje",
"Europe/Stockholm",
"Europe/Tirane",
"Europe/Vaduz",
"Europe/Vatican",
"Europe/Vienna",
"Europe/Warsaw",
"Europe/Zagreb",
"Europe/Zurich",
"Libya",
"MET",
"Poland",
"Africa/Addis_Ababa",
"Africa/Asmara",
"Africa/Asmera",
"Africa/Dar_es_Salaam",
"Africa/Djibouti",
"Africa/Juba",
"Africa/Kampala",
"Africa/Khartoum",
"Africa/Mogadishu",
"Africa/Nairobi",
"Antarctica/Syowa",
"Asia/Aden",
"Asia/Amman",
"Asia/Baghdad",
"Asia/Bahrain",
"Asia/Beirut",
"Asia/Damascus",
"Asia/Gaza",
"Asia/Hebron",
"Asia/Istanbul",
"Asia/Jerusalem",
"Asia/Kuwait",
"Asia/Nicosia",
"Asia/Qatar",
"Asia/Riyadh",
"Asia/Tel_Aviv",
"EET",
"Etc/GMT-3",
"Europe/Athens",
"Europe/Bucharest",
"Europe/Chisinau",
"Europe/Helsinki",
"Europe/Istanbul",
"Europe/Kiev",
"Europe/Mariehamn",
"Europe/Minsk",
"Europe/Moscow",
"Europe/Nicosia",
"Europe/Riga",
"Europe/Simferopol",
"Europe/Sofia",
"Europe/Tallinn",
"Europe/Tiraspol",
"Europe/Uzhgorod",
"Europe/Vilnius",
"Europe/Volgograd",
"Europe/Zaporozhye",
"Indian/Antananarivo",
"Indian/Comoro",
"Indian/Mayotte",
"Israel",
"Turkey",
```

```
"W-SU",
"Asia/Dubai",
"Asia/Muscat",
"Asia/Tbilisi",
"Asia/Yerevan",
"Etc/GMT-4",
"Europe/Samara",
"Indian/Mahe",
"Indian/Mauritius",
"Indian/Reunion",
"Asia/Kabul",
"Asia/Tehran",
"Iran",
"Antarctica/Mawson",
"Asia/Aqtou",
"Asia/Aqtobe",
"Asia/Ashgabat",
"Asia/Ashkhabad",
"Asia/Baku",
"Asia/Dushanbe",
"Asia/Karachi",
"Asia/Oral",
"Asia/Samarkand",
"Asia/Tashkent",
"Asia/Yekaterinburg",
"Etc/GMT-5",
"Indian/Kerguelen",
"Indian/Maldives",
"Asia/Calcutta",
"Asia/Colombo",
"Asia/Kolkata",
"Asia/Kathmandu",
"Asia/Katmandu",
"Antarctica/Vostok",
"Asia/Almaty",
"Asia/Bishkek",
"Asia/Dacca",
"Asia/Dhaka",
"Asia/Kashgar",
"Asia/Novosibirsk",
"Asia/Omsk",
"Asia/Qyzylorda",
"Asia/Thimbu",
"Asia/Thimphu",
"Asia/Urumqi",
"Etc/GMT-6",
"Indian/Chagos",
"Asia/Rangoon",
"Indian/Cocos",
"Antarctica/Davis",
"Asia/Bangkok",
"Asia/Ho_Chi_Min",
"Asia/Hovd",
"Asia/Jakarta",
"Asia/Krasnoyarsk",
"Asia/Novokuznetsk",
"Asia/Phnom_Penh",
"Asia/Pontianak",
"Asia/Saigon",
"Asia/Vientiane",
"Etc/GMT-7",
"Indian/Christmas",
"Antarctica/Casey",
"Asia/Brunei",
"Asia/Chita",
"Asia/Choibalsan",
"Asia/Chongqing",
"Asia/Chungking",
```


"Asia/Harbin",
"Asia/Hong_Kong",
"Asia/Irkutsk",
"Asia/Kuala_Lumpur",
"Asia/Kuching",
"Asia/Macao",
"Asia/Macau",
"Asia/Makassar",
"Asia/Manila",
"Asia/Shanghai",
"Asia/Singapore",
"Asia/Taipei",
"Asia/Ujung_Pandang",
"Asia/Ulaanbaatar",
"Asia/Ulan_Bator",
"Australia/Perth",
"Australia/West",
"Etc/GMT-8",
"Hongkong",
"PRC",
"ROC",
"Singapore",
"Australia/Eucla",
"Asia/Dili",
"Asia/Jayapura",
"Asia/Khandyga",
"Asia/Pyongyang",
"Asia/Seoul",
"Asia/Tokyo",
"Asia/Yakutsk",
"Etc/GMT-9",
"Japan",
"Pacific/Palau",
"ROK",
"Australia/Adelaide",
"Australia/Broken_Hill",
"Australia/Darwin",
"Australia/North",
"Australia/South",
"Australia/Yancowinna",
"Antarctica/DumontDURville",
"Asia/Magadan",
"Asia/Sakhalin",
"Asia/Ust-Nera",
"Asia/Vladivostok",
"Australia/ACT",
"Australia/Brisbane",
"Australia/Canberra",
"Australia/Currie",
"Australia/Hobart",
"Australia/Lindeman",
"Australia/Melbourne",
"Australia/NSW",
"Australia/Queensland",
"Australia/Sydney",
"Australia/Tasmania",
"Australia/Victoria",
"Etc/GMT-10",
"Pacific/Chuuk",
"Pacific/Guam",
"Pacific/Port_Moresby",
"Pacific/Saipan",
"Pacific/Truk",
"Pacific/Yap",
"Australia/LHI",
"Australia/Lord_Howe",
"Antarctica/Macquarie",
"Asia/Srednekolymsk",

```
"Etc/GMT-11",
"Pacific/Bougainville",
"Pacific/Efate",
"Pacific/Guadalcanal",
"Pacific/Kosrae",
"Pacific/Noumea",
"Pacific/Pohnpei",
"Pacific/Ponape",
"Pacific/Norfolk",
"Antarctica/McMurdo",
"Antarctica/South_Pole",
"Asia/Anadyr",
"Asia/Kamchatka",
"Etc/GMT-12",
"Kwajalein",
"NZ",
"Pacific/Auckland",
"Pacific/Fiji",
"Pacific/Funafuti",
"Pacific/Kwajalein",
"Pacific/Majuro",
"Pacific/Nauru",
"Pacific/Tarawa",
"Pacific/Wake",
"Pacific/Wallis",
"NZ-CHAT",
"Pacific/Chatham",
"Etc/GMT-13",
"Pacific/Apia",
"Pacific/Enderbury",
"Pacific/Fakaofu",
"Pacific/Tongatapu",
"Etc/GMT-14",
"Pacific/Kiritimati"
}
```

Related Information

[Schedule-Based Scaling Policy \[page 20\]](#)

[Parameters for a Schedule-Based Scaling Policy \[page 23\]](#)

3.3 Defining a Custom Metric

Define your own metrics to scale applications based on your requirements.

Context

As an alternative to the standard metrics provided by the Application Autoscaler, you can also define custom metrics such as lengths of job queues or technology-specific metrics like JVM heap usage for scaling. Custom metrics offer more flexibility, allowing you to scale applications based on your own requirements.

Applications can use custom metrics in the following scenarios:

- An application sends custom metrics to the Application Autoscaler and the same application is scaled.
- One application sends custom metrics to scale another application.

Scenario 1

The following scenario describes an application sending custom metrics to the Autoscaler service and the same application is scaled.

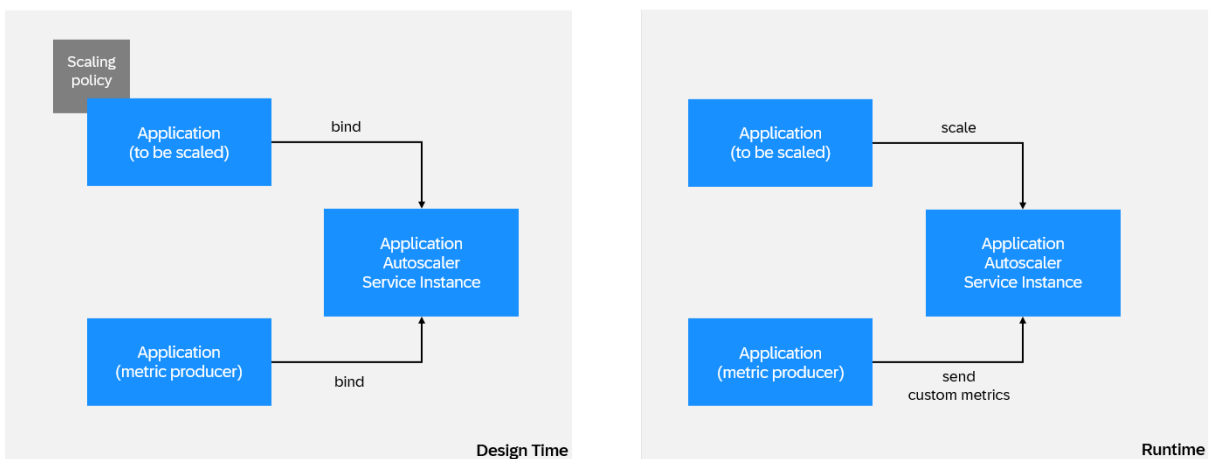
To use custom metrics, you must implicitly define the custom metric by using it in the scaling policy and the application needs to emit custom metric values using the custom metrics API at regular intervals.

The following example shows a policy with a custom metric:

```
{
  "instance_min_count":1,
  "instance_max_count":4,
  "scaling_rules":[
    {
      "metric_type":"jobqueue",
      "breach_duration_secs":60,
      "threshold":100,
      "operator":>=",
      "cool_down_secs":120,
      "adjustment":"+1"
    }
  ]
}
```

Scenario 2

The following scenario involves two different applications that are bound to the same Autoscaler service, for example, the metric producer app and an app that should be scaled. Here, the metric producer app sends custom metrics on behalf of the app to be scaled.



→ Remember

To use custom metrics, you must implicitly define the custom metric by using it in the scaling policy and the metric producer app needs to emit custom metric values using the custom metrics API at regular intervals.

The scaling app should set `metric_submission_strategy.allow-from` to `bound_app` in the configuration of the scaling policy. This scaling policy is also attached with the application to be scaled.

Note

The metric type used for custom metrics must not be any of the standard metric types.

Tip

We recommend a minimum duration of one minute between successive emissions of a custom metric.

As part of the binding process, the Application Autoscaler service instance provides necessary credentials to emit custom metrics.

The generation of these credentials is based on the credential type parameter in the service binding. The supported credential types are `binding-secret` and `x509`. If no credential type is provided, `x509` is used as the default credential type.

Procedure

1. Bind your application with the Application Autoscaler service instance using a policy that contains custom metrics scaling rule.

sample policy.json with credential Type as x509

Sample Code

```
{
  "instance_min_count":1,
  "instance_max_count":4,
  "scaling_rules":[
    {
      "metric_type":"jobqueue",
      "breach_duration_secs":60,
      "threshold":100,
      "operator":>=",
      "cool_down_secs":120,
      "adjustment":"+1"
    }
  ],
  "credential-type": "x509"
}
```

bind the service instance with the scaling policy

Sample Code

```
cf bind-service <application> <application-autoscaler-service-instance> -c
policy.json
```

Note

In the scenario in which a different app (metrics producer app) is used to send custom metrics details to the service, make sure that the metrics producer app is also bound to the service. The scaling policy

must have the `metric_submission_strategy.allow-from` set to `bound_app` in the configuration (as shown in the following code example).

Sample Code

```
{
  "configuration": {
    "custom_metrics": {
      "metric_submission_strategy": {
        "allow_from": "bound_app"
      }
    }
  },
  "instance_min_count": 1,
  "instance_max_count": 4,
  "scaling_rules": [
    {
      "metric_type": "jobqueue",
      "breach_duration_secs": 60,
      "threshold": 100,
      "operator": ">=",
      "cool_down_secs": 120,
      "adjustment": "+1"
    }
  ],
  "credential-type": "x509"
}
```

2. Read custom metrics binding credentials from the application environment.

Binding Credentials for mTLS

Sample Code

```
"custom_metrics": {
  "mtls_url": "https://autoscaler-metrics-
mtls.cf.<landscape>.hana.ondemand.com",
}
```

3. **Authenticate with X.509 Certificate:**

Your application requires a base url, certificate, and a private key for mutual TLS authentication (mTLS).

- Use the base URL provided in the `mtls_url` of the service binding.
- Use the X.509 certificate and private key from the paths provided in the environment variables `CF_INSTANCE_CERT` and `CF_INSTANCE_KEY` respectively.

Note


The environment variables `CF_INSTANCE_CERT` and `CF_INSTANCE_KEY` are not shown in the cockpit or using the `cf env` command. They're only visible inside the running application container.

Note

The X.509 certificate and private key pair are valid for 24 hours. At least 20 minutes before expiration, they're regenerated and new files replace the existing files. Make sure that your code reloads the X.509 certificate and private key pair if they're expired. See [Using Instance Identity Credentials](#) in the Cloud Foundry Documentation.

Authenticate with Basic Authentication (Deprecated, don't use):

`url`, `username`, and `password` in the application environment are provided for backwards compatibility with the HTTP basic authentication scheme and shouldn't be used going forward. If your application still needs to authenticate with basic authentication, use `credential-type binding-secret` in the service binding and plan to migrate to mTLS (`x509`) authentication until April 30, 2025.

4. Push custom metrics at an interval of your choice using APIs. For the REST API specification, see [Application Autoscaler custom metrics API](#) .

4 Administration

This topic details the tasks for the application developer/operator to configure the service and ensure proper operation.

Operations

[Preventing the Application Autoscaler from Scaling your Application \[page 39\]](#)

Related Information

[Administration and Operations in the Cloud Foundry Environment](#)

4.1 Preventing the Application Autoscaler from Scaling your Application

You want to prevent the Application Autoscaler from scaling your application without unbinding it from the Application Autoscaler service instance or deleting the scaling policy.

Prerequisites

- You have set up autoscaling (by binding your application to an Application Autoscaler service instance) and a scaling policy.
- You are a Space Developer in the space where your application is running.

Context

For example, during a deployment or maintenance of your application you want to manually control the scaling of your application (and prevent the Application Autoscaler from automatically scaling your application) without changing the setup of your application.

Procedure

Set the label `app-autoscaler.cloudfoundry.org/disable-autoscaling` on your application.

Sample Code

```
cf set-label app <application name> app-autoscaler.cloudfoundry.org/disable-  
autoscaling=<any non-empty string value>
```

Note

- You can set the value of the label `app-autoscaler.cloudfoundry.org/disable-autoscaling` to any non-empty value. We recommend to use a value that can be used to later determine the reason why it was set.
- Other applications or services might set the label as well. For example, the SAP Cloud Deployment service sets and removes the label during a deployment to be able to completely control the scaling during a deployment. For more information, see [\(Experimental\) Incremental Blue-Green Deployment Strategy](#).

5 Accessibility Features in Application Autoscaler

To optimize your experience of Application Autoscaler, SAP Business Technology Platform (SAP BTP) provides features and settings that help you use the software efficiently.

📘 Note

Application Autoscaler runs on the SAP BTP cockpit. For this reason, accessibility features for SAP BTP cockpit also apply. See the accessibility documentation for SAP BTP cockpit on SAP Help Portal at [Accessibility Features in SAP BTP Cockpit](#).

For more information on screen reader support and keyboard shortcuts, see [Accessibility for End Users](#).

Additionally, Application Autoscaler provides high contrast themes. See [Use High Contrast Themes \[page 41\]](#).

Use High Contrast Themes

Context

This feature is recommended for users with need for high contrast themes.

Steps



1. Open the [Application Autoscaler Dashboard](#).
2. In the upper right corner of the screen, choose your profile name and then [Choose Theme](#).
3. Select the high contrast theme you prefer: [High Contrast Black](#) or [High Contrast White](#).
Your selection is applied automatically.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
 - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
 - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering an SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

© 2025 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.