Understand Centralized Umbrella Log Management with Amazon's S3 Service for MSP, MSSP, and Multi-org Customers

Contents

Introduction

This document describes centralized Umbrella log management with Amazon's S3 service for MSP, MSSP, and multi-org customers.

Overview

The MSP, MSSP, and Multi-org consoles have the ability to store the DNS, URL and IP logs of your customers offline in cloud storage. The storage is in Amazon S3 and after the logs have been uploaded, they can be downloaded and kept for compliance reasons or security analysis.

This documentation helps you understand this feature, set it up in both your Umbrella dashboard and your Amazon S3 console, and run through several options for configuration, including the duration of time you would like the logs to be kept in S3.

Umbrella for MSP, MSSP, and Multi-Org all have the ability to upload the traffic activity logs from the child organizations of the console and store those logs in the cloud. Amazon’s AWS S3 (**S**imple **S**torage **S**ervice) is the service that archives logs and is sometimes referred to as it isoffline storageit is or it islog retention.it is

Archiving logs can be useful for several reasons, depending on your need. For some people, the exported and archived logs can be imported into data analysis or security forensic tools, such as SIEMs. For others, an archive of activity logs can be useful for data forensics in case of a security incident, or human resources records.

AWS S3 stores logs in a compressed (gzip) archive in CSV format. Because logs are uploaded every ten minutes, thereit iss a minimum ten-minute delay between network traffic coming from your network, logged by Umbrella, and then made available to download from S3.

The orgID number from the Console

Each customer organization uploads their logs individually, using the orgID number from the Console to map each customer to a folder. The feature can also be enabled or disabled on a per-customer / per-organization basis.

Two Types of Umbrella Log Management

Log Management is performed by uploading logs to what is called a it isbucketit is (essentially a folder within AWSit iss S3 environment). There are two ways to host a bucket for your Umbrella logs:

- Administered, managed and paid for by you, the company administrator.
- Administered, managed and paid for by Cisco Umbrella.

There are pros and cons to having Cisco manage your S3 bucket.

The pros of Cisco managing your bucket:

- Extremely easy to setup. It only takes a couple of minutes and afterward itit iss extremely easy to manage.
- Cisco bucket-managment is included in your license cost with Umbrella, effectively making the service free. While it is inxpensive to have your own bucket, the overhead cost of managing another bill can be prohibitive.

The pros of managing an S3 instance yourself:

- There is no limitation on how long data can be stored offline. Cisco limits offline storage to 30 days at the maximum.
- You can add anything to your bucket, including log files from Umbrella, so the bucket can be used by other applications as well.
- You can get support directly from Amazon for advanced configuration assistance, such as automation or help with command line.

For most customers the cost of maintaining a bucket is very inexpensive, but can prove to be hassle.

Getting Started

The Log Management feature can be found in the Console under **Settings > Log Management** (you can hit the dropdown arrow).

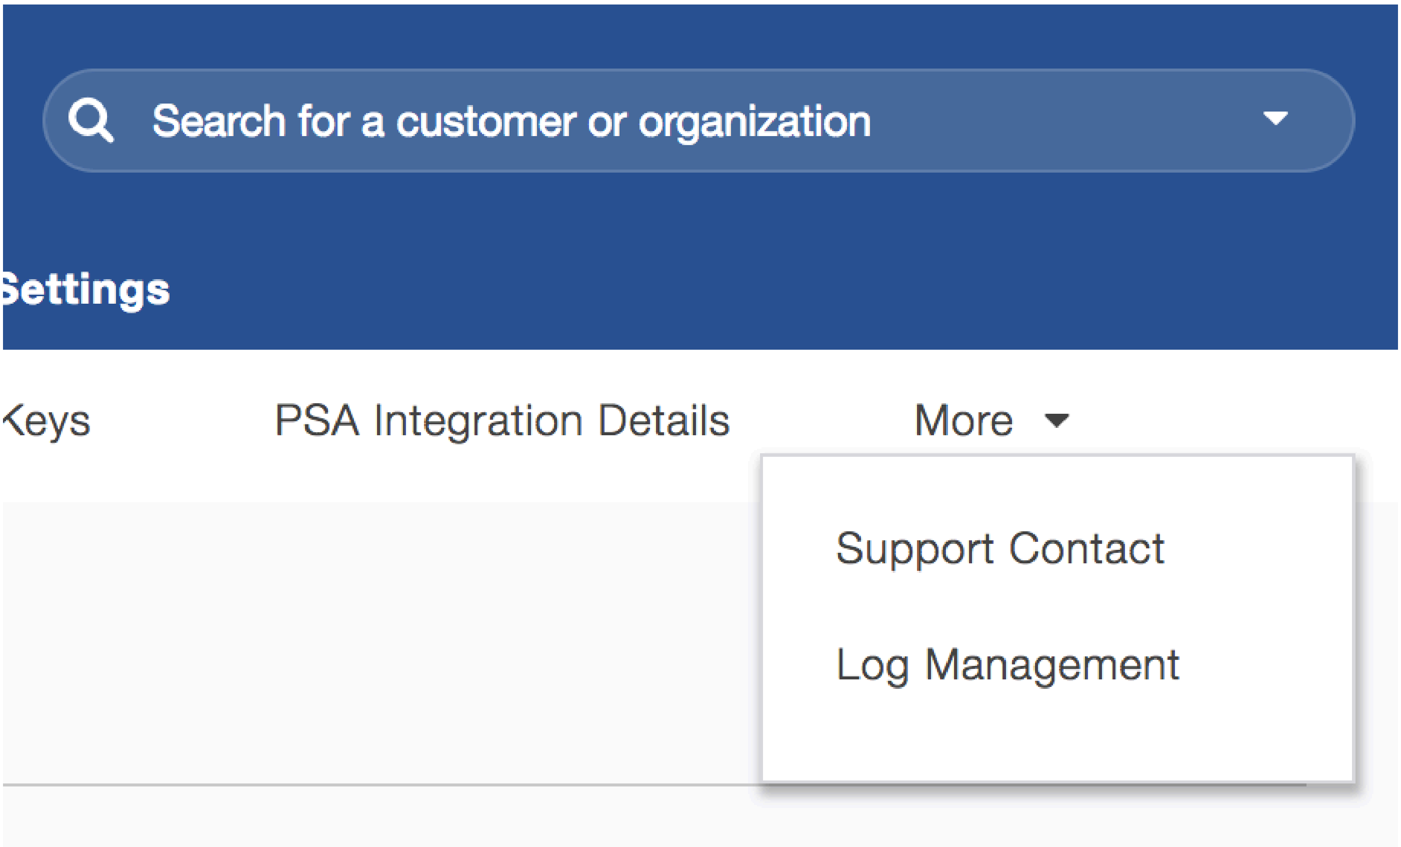

115012963103

Configuring a Self-managed S3 Bucket

Prerequisites

In order to archive logs, you must meet these requirements:

- Full administrative access to the Cisco Umbrella MSP, MSSP, or Multi-org Console.
- A login to Amazon AWS service (https://aws.amazon.com/console/). If you donit ist have an account, Amazon provides free signup for S3. However, they require a credit card in case your usage exceeds free plan usage.
- A bucket configured in Amazon S3 for log storage. See the next section for instructions on configuring and setting up the Amazon S3 bucket.

Setting up Your Amazon S3 Bucket

1. Start by signing into the AWS Console, and selecting "**S3**" from the list of options under Storage.

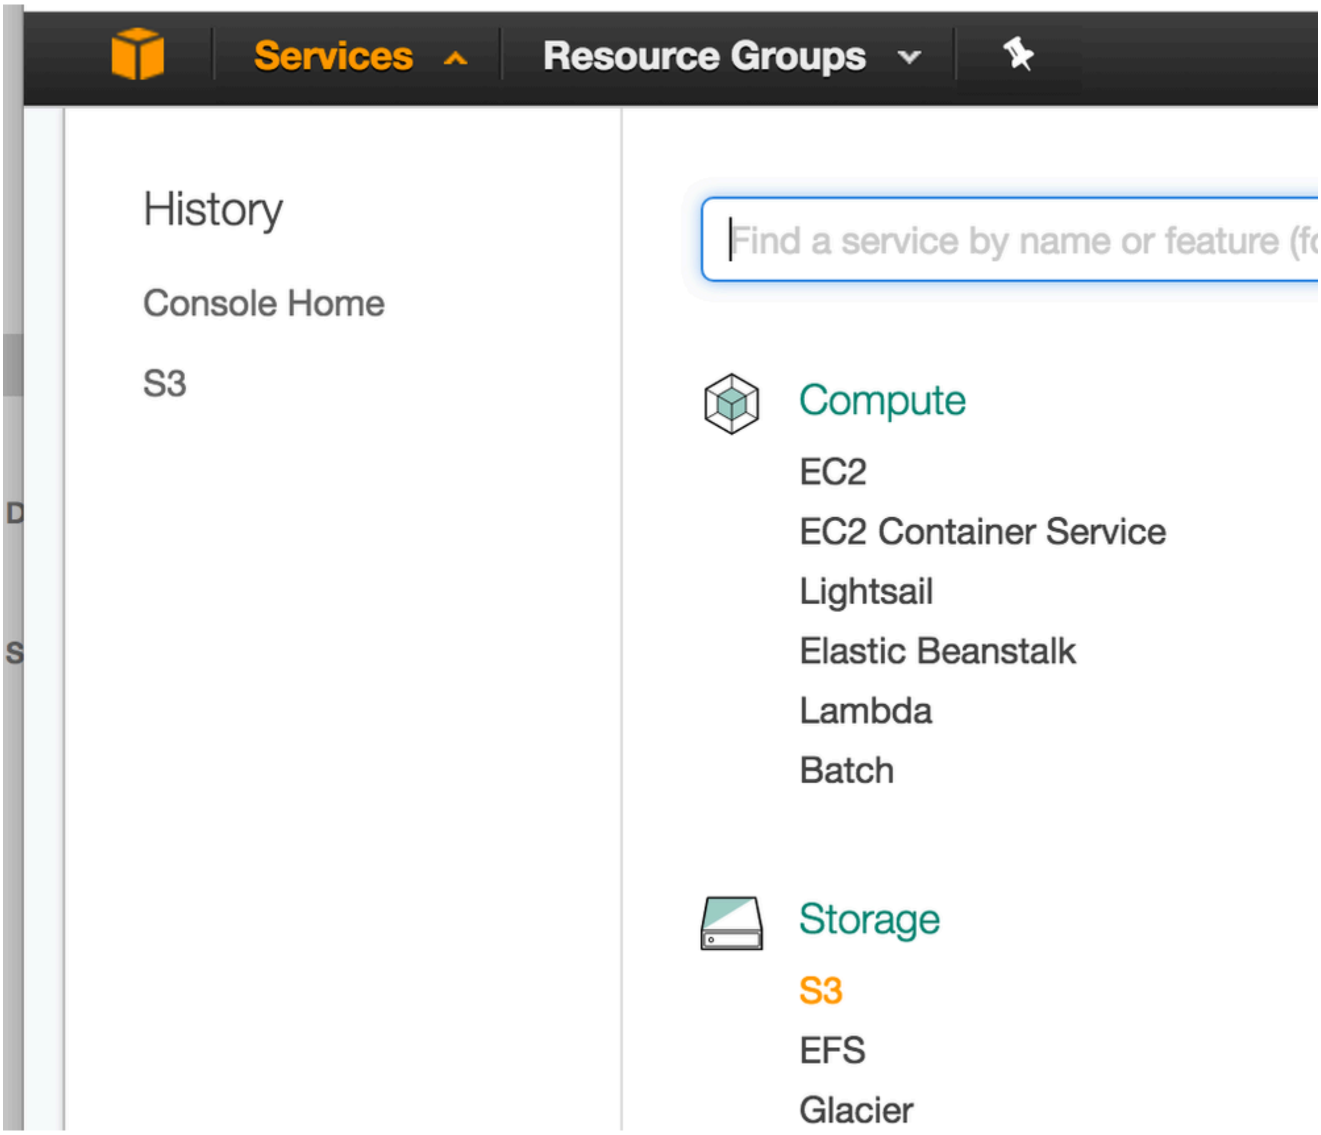

115012842106

2. You see an introduction screen welcoming you to the Amazon Simple Storage System

3. Next, if you do not already have a bucket, you want to create one. Click **Create Bucket**

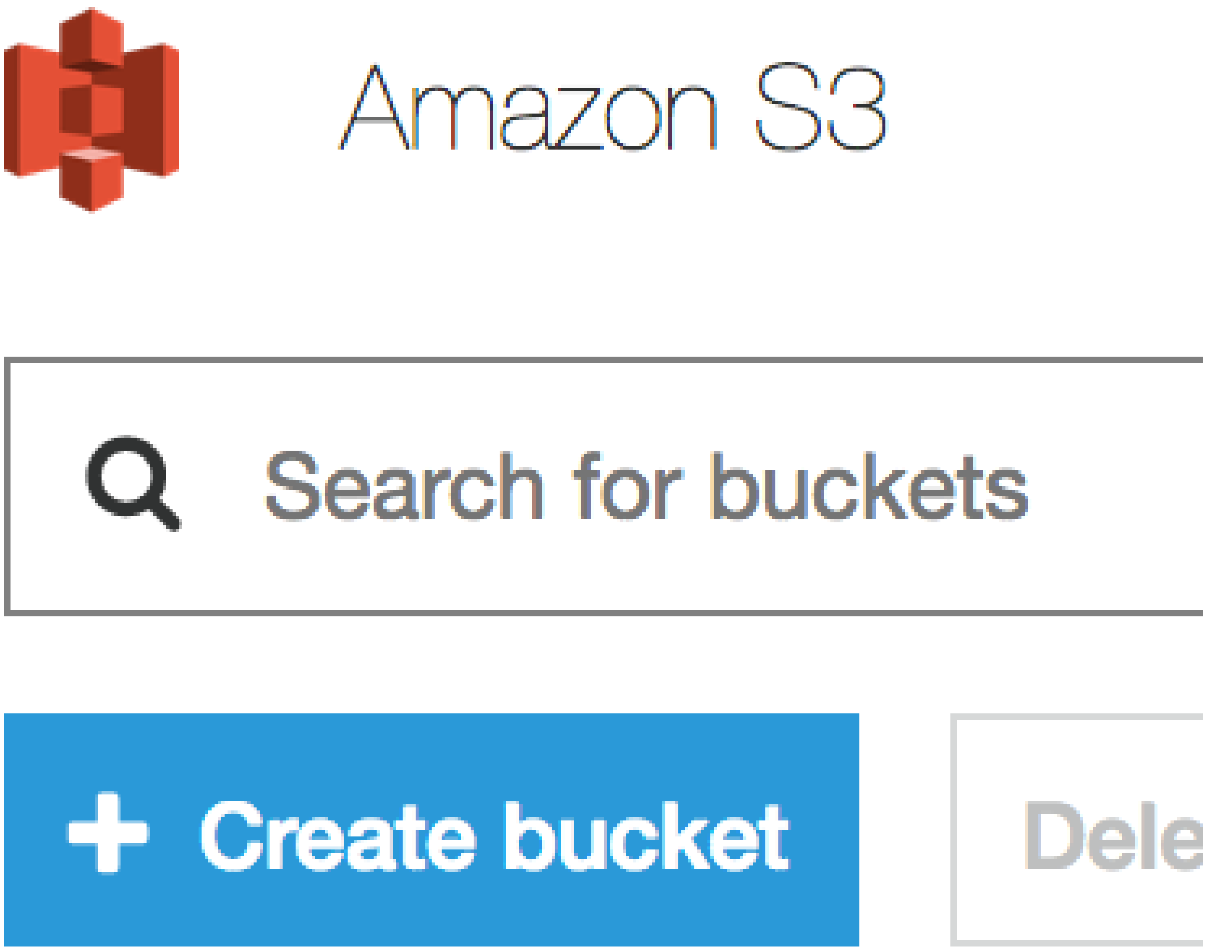

115012842326

.

4. Start by entering a **Bucket Name**
   The bucket name must be universally unique—not just to your AWS or your Umbrella, but to all of Amazon AWS. Using something personal, such as *"my-organization-name-log-bucket"* can help you bypass the requirement for universally unique bucket name. The bucket name must only use lowercase letters and cannot contain spaces or periods, and must comply with DNS naming conventions. For more information on name restrictions, read here**.** For more information on bucket creation, including naming, read here.

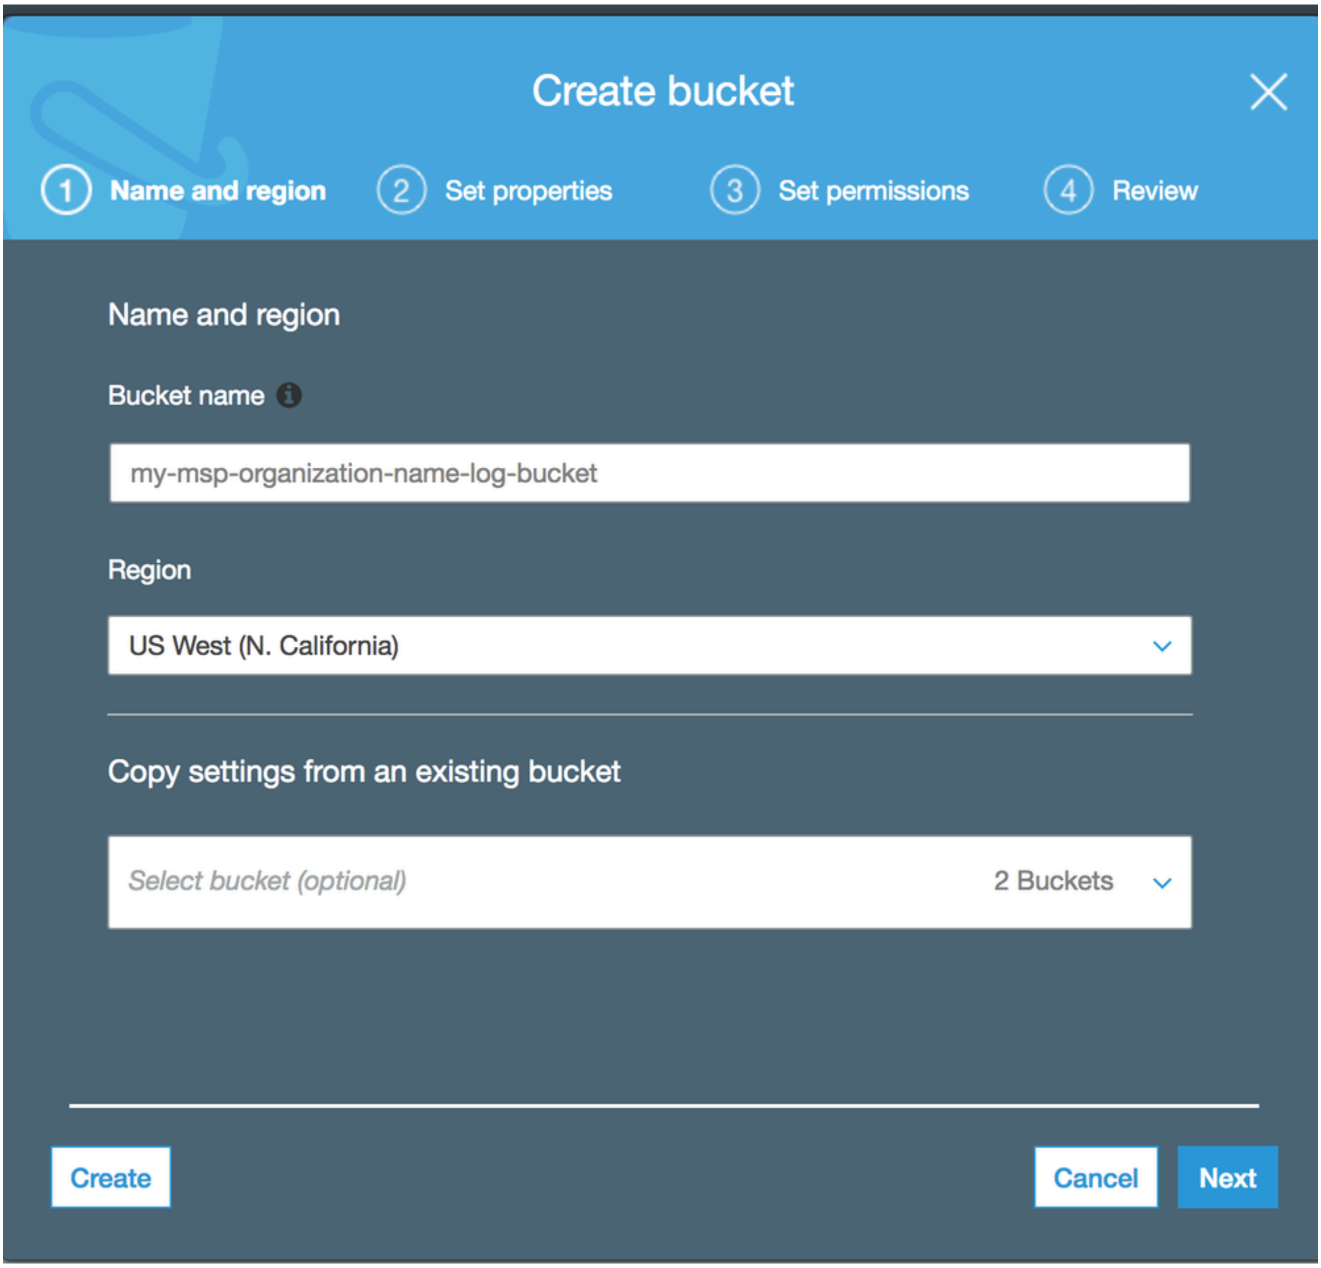

Create bucket

Name and region

Bucket name

my-msp-organization-name-log-bucket

Region

US West (N. California)

Copy settings from an existing bucket

Select bucket (optional) 2 Buckets

115013010503

5. Select whichever region works best for your location and click **Create**. Do not copy the settings from another bucket
6. In the "Set properties" step, just click **Next**. These can be adjusted later
7. In the "Set permissions" step, just click **Next**. We are going to revisit the permissions later to set up the bucket for uploading
8. Finalize the review process and click **Create bucket**

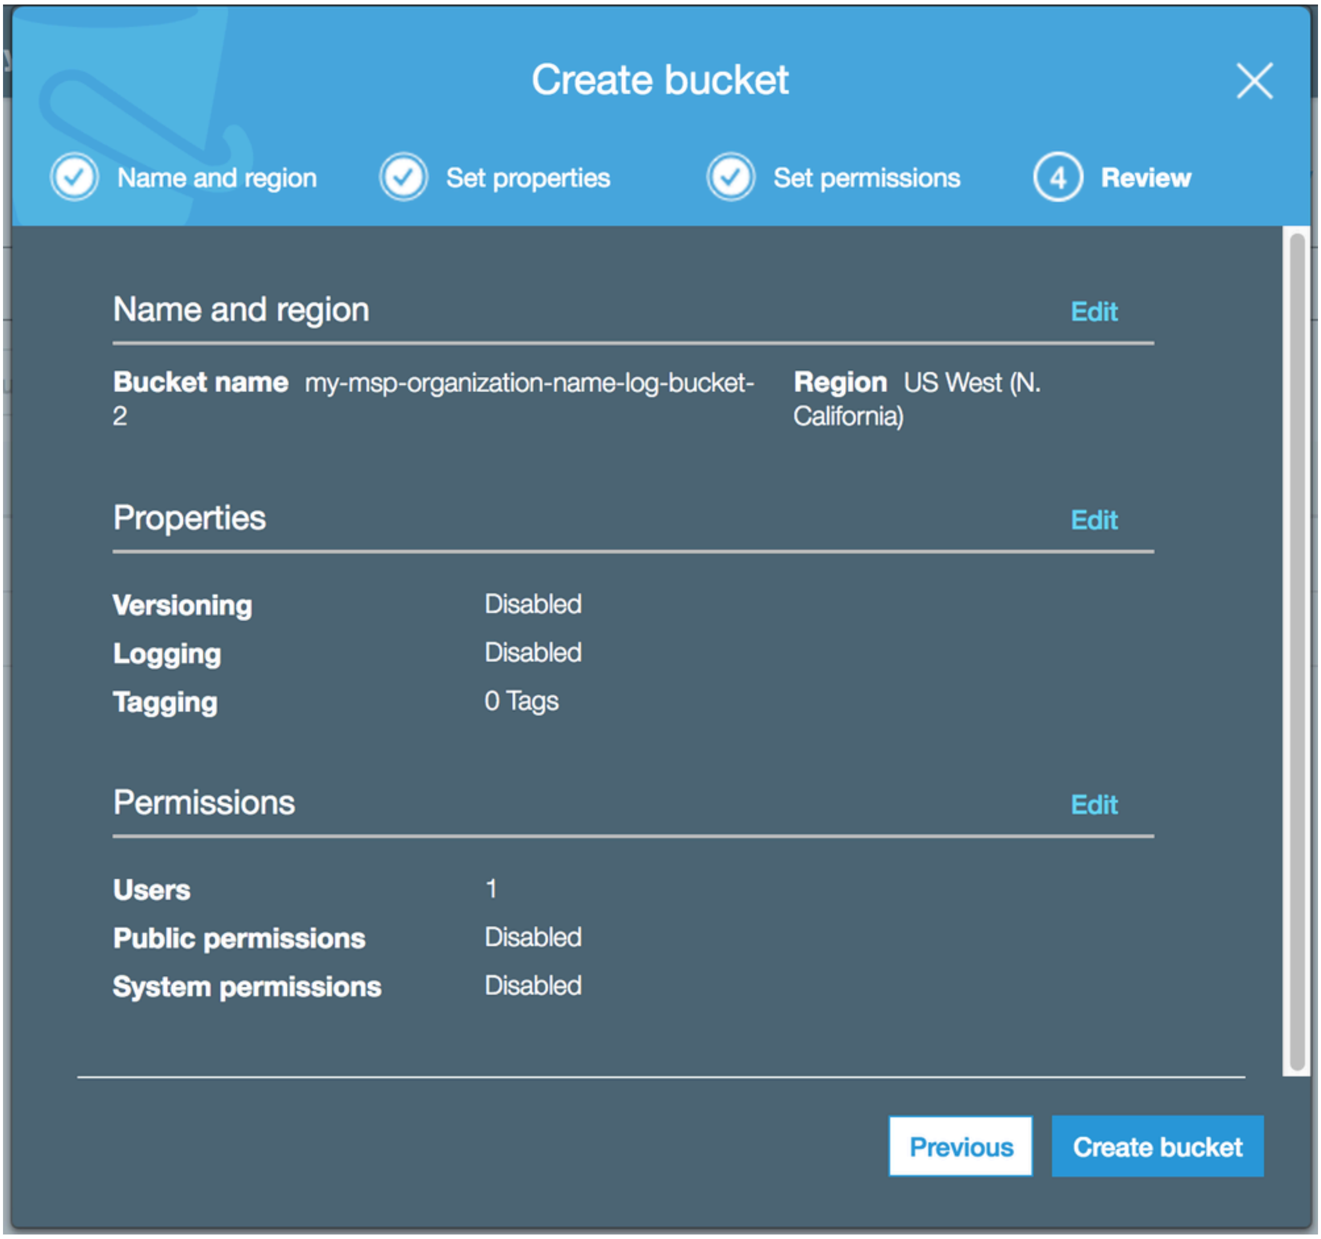

Create bucket

✓ Name and region ✓ Set properties ✓ Set permissions 4 Review

Name and region Edit

Bucket name my-msp-organization-name-log-bucket-2 **Region** US West (N. California)

Properties Edit

Versioning	Disabled
Logging	Disabled
Tagging	0 Tags

Permissions Edit

Users	1
Public permissions	Disabled
System permissions	Disabled

115012842686

9. Next, you need to configure the bucket to accept uploads from the Umbrella Service. In S3, this is referred to as a bucket policy. Click the name of your newly configured bucket and then select the **Permissions** tab at the top of the interface

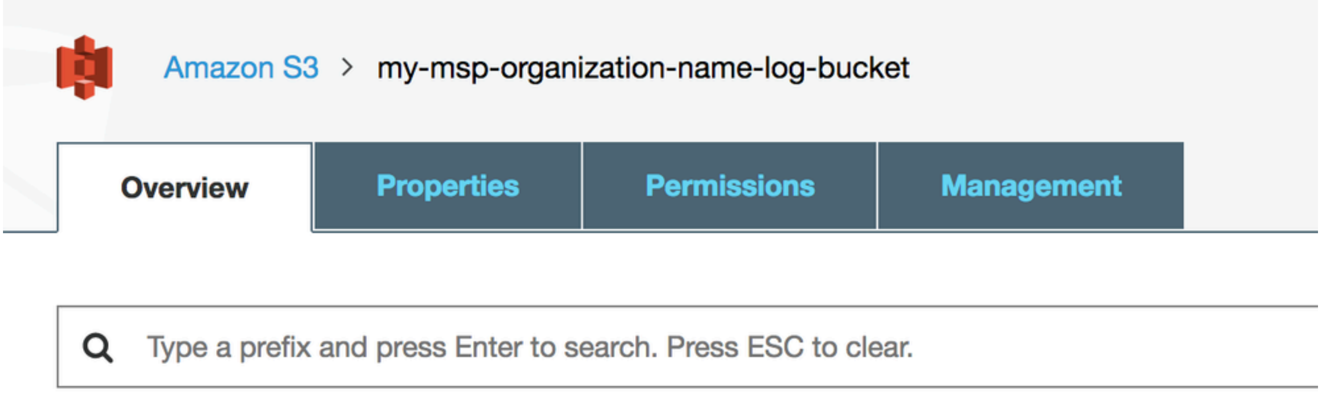

Amazon S3 > my-msp-organization-name-log-bucket

Overview
Properties
Permissions
Management
Type a prefix and press Enter to search. Press ESC to clear.

115012842906

10. Select **Bucket Policy** and then you are prompted to paste in the bucket

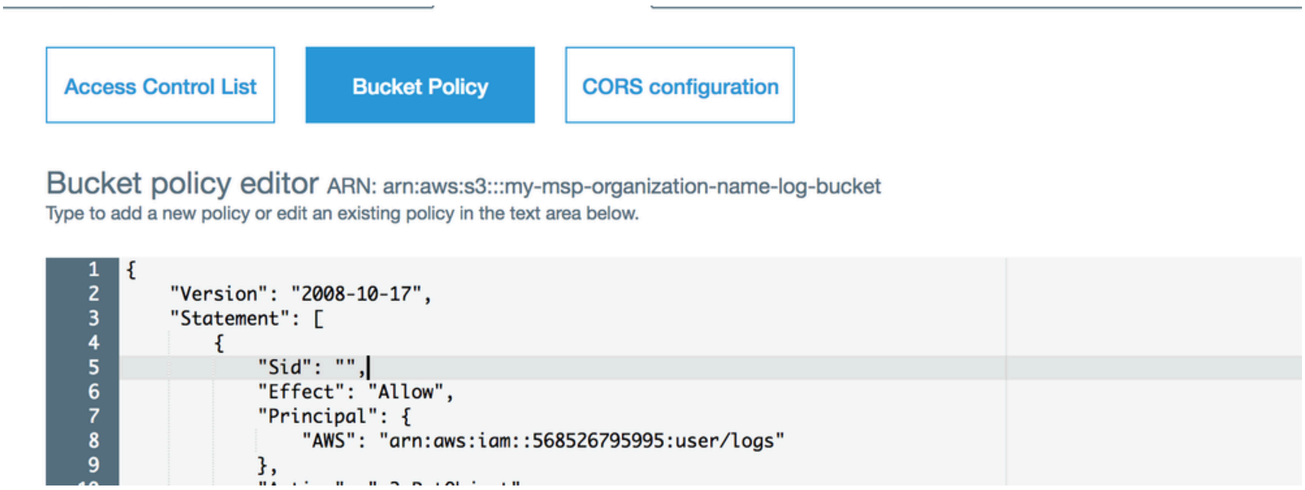

Bucket policy editor ARN: arn:aws:s3:::my-msp-organization-name-log-bucket

Type to add a new policy or edit an existing policy in the text area below.

```
{
    "Version": "2008-10-17",
    "Statement": [
        {
            "Sid": "",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::568526795995:user/logs"
            },
```

115012843006

11. Copy and paste the JSON string below, which contains the bucket policy, to a text editor or simply paste it into the window. Substitute your *exact* bucket name where **bucketname** is specified below. Failure to do this results in an error message

```
{
"Version": "2008-10-17",
"Statement": [
{
"Sid": "",
"Effect": "Allow",
"Principal": {
"AWS": "arn:aws:iam::568526795995:user/logs"
},
"Action": "s3:PutObject",
"Resource": "arn:aws:s3:::bucketname/*"
},
{
"Sid": "",
"Effect": "Deny",
"Principal": {
"AWS": "arn:aws:iam::568526795995:user/logs"
},
"Action": "s3:GetObject",
"Resource": "arn:aws:s3:::bucketname/*"
},

{
"Sid": "",
"Effect": "Allow",
"Principal":

{ "AWS": "arn:aws:iam::568526795995:user/logs" }


,
"Action": "s3:GetBucketLocation",
"Resource": "arn:aws:s3:::bucketname"
},

{
```

```
"Sid": "",
"Effect": "Allow",
"Principal": {
"AWS": "arn:aws:iam::568526795995:user/logs"
},
"Action": "s3:ListBucket",
"Resource": "arn:aws:s3:::bucketname"
}
]
}
```

12. Click **Save** to confirm this change

Verifying Your Amazon S3 Bucket

Step 1:

1. Go back to your Umbrella Console and navigate to **Settings > Log Management**
2. Click "Amazon S3" to expand the window
3. In the Bucket Name field, type or paste the *exact bucket name* you created in S3 and click **Verify** You receive a confirmation message in your dashboard indicating that the bucket was successfully verified.

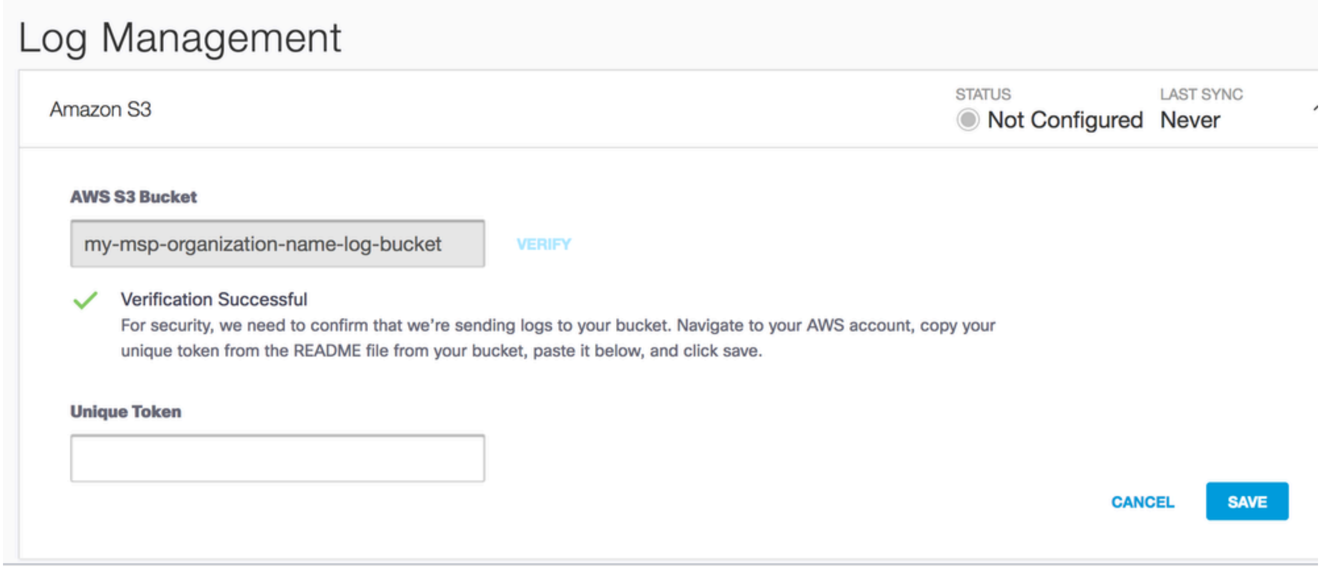

my-msp-organization-name-log-bucket

✓ Verification Successful

For security, we need to confirm that we're sending logs to your bucket. Navigate to your AWS account, copy your unique token from the README file from your bucket, paste it below, and click save.

115012847146

If you receive an error indicating that your bucket could not be verified, recheck the syntax of the bucket name and review the configuration. If problems persist, please open a case with our support department

Step 2:

As a secondary precaution to ensure the correct bucket was specified, Umbrella requests that you enter a unique activation token. The activation token can be obtained by revisiting your S3 bucket. As part of the verification process, a file named **README_FROM_UMBRELLA.txt** was uploaded from Umbrella to your Amazon S3 bucket and appears there.

1. Download the readme file by double-clicking on it and then open it in a text editor. Within the file, there is a unique token tying your S3 bucket to your Umbrella dashboard

Note: You might need to refresh your S3 bucket in the browser in order to see the README file after itit iss been uploaded.

2. Return to the Umbrella dashboard and paste the token into the field labeled "Unique token", then click **Save**. At this point, the configuration is
complete. To review your configuration, just click the Amazon S3 name in the Log Managemen section

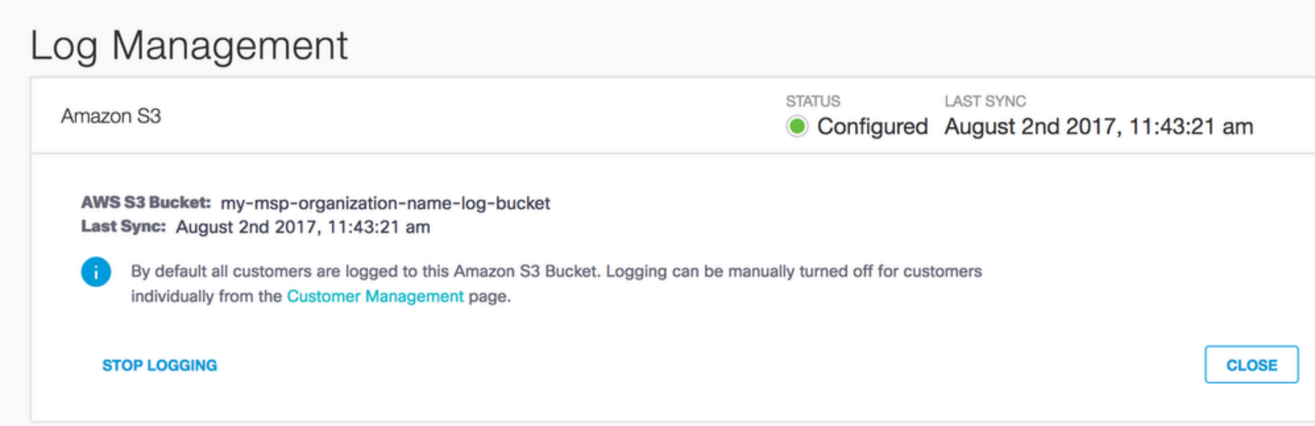

Log Management

Amazon S3

STATUS
● Configured

LAST SYNC
August 2nd 2017, 11:43:21 am

AWS S3 Bucket: my-msp-organization-name-log-bucket
Last Sync: August 2nd 2017, 11:43:21 am

By default all customers are logged to this Amazon S3 Bucket. Logging can be manually turned off for customers individually from the Customer Management page.

115012848126

Managing the Log Lifecycle

When you are using S3, you can manage the lifecycle of the data within the bucket to extend the duration of time you would like to retain logs for. Depending on the reason you are using the external log management, the duration could be very short or very long. For instance, you can simply download the logs from the S3 bucket after 24 hours and store them offline, or retain the logs indefinitely in the cloud. By default, Amazon stores the data in a bucket indefinitely but unlimited storage does raise the cost of maintaining the bucket. For more information on S3 lifecycles, read here.

To configure the lifecycle of your bucket:

1. Select **Management** and then click **Lifecycle**

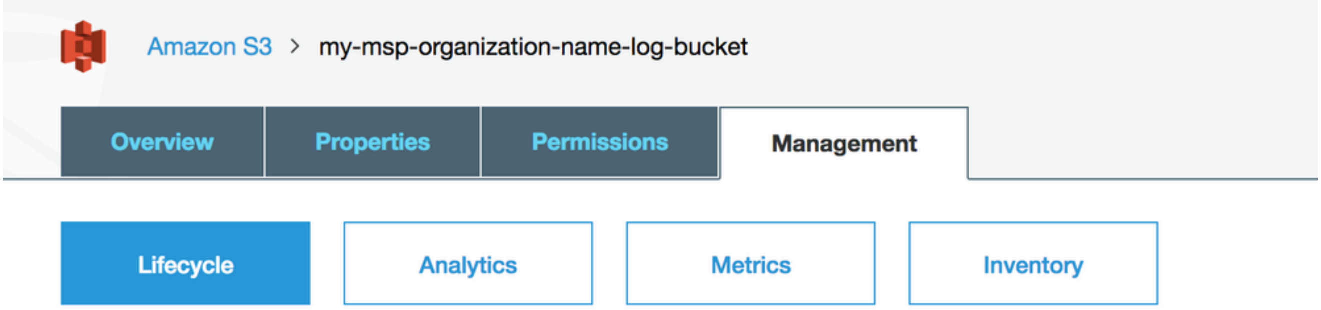

115012848246

2. Click **Add a Rule**, then Apply the Rule to the whole bucket (or a subfolder if you have configured it as such).
3. Select an Action on Objects, such as Delete or Archive, then select the time period and whether you would like to use Glacier storage to help reduce your Amazon costs. (Glacier is it iscoldit is off-line storage, which while slower to access, is less expensive.)

4. If you prefer to manage logs with another method (such as your internal backup solution) you can simply download the logs from S3 and preserve them in another way, then set your retention time to a few days.

Configuring a Cisco-Managed S3 Bucket

Navigate to **Settings > Log Management** in your Umbrella dashboard.

There are two options:

- Use your company-managed Amazon S3 bucket
- Use a Cisco-managed Amazon S3 bucket

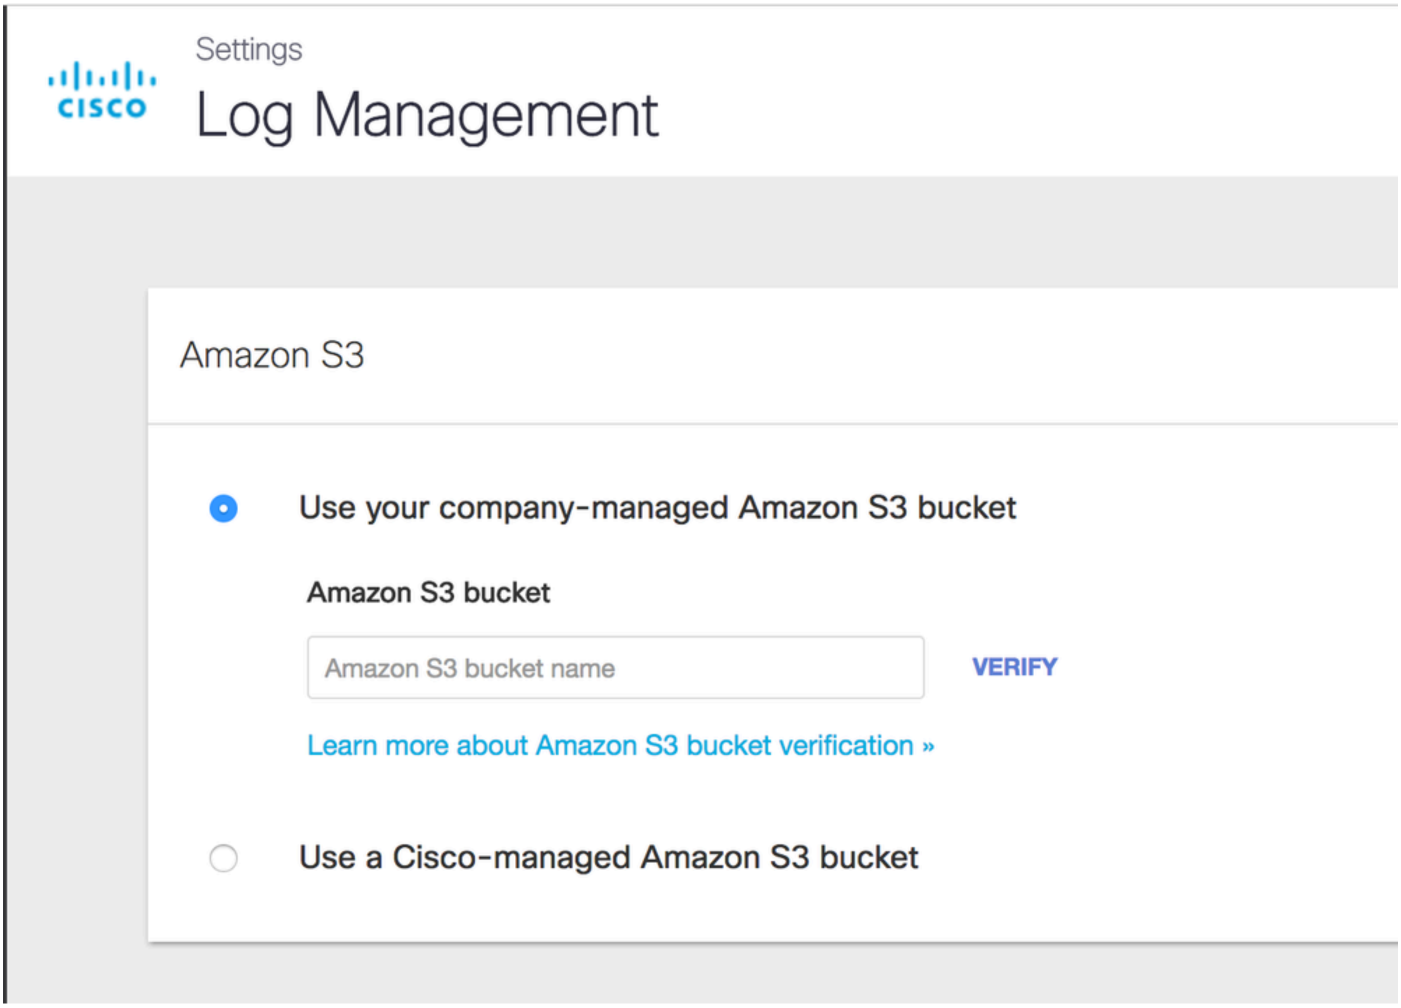

Settings

Log Management

Amazon S3

● Use your company-managed Amazon S3 bucket

Amazon S3 bucket

Amazon S3 bucket name

Learn more about Amazon S3 bucket verification »

○ Use a Cisco-managed Amazon S3 bucket

25231151138964

Pick "Use a Cisco-managed Amazon S3 bucket" and you are given two new options: "**Select a Region**" and "**Select a Retention Duration**".

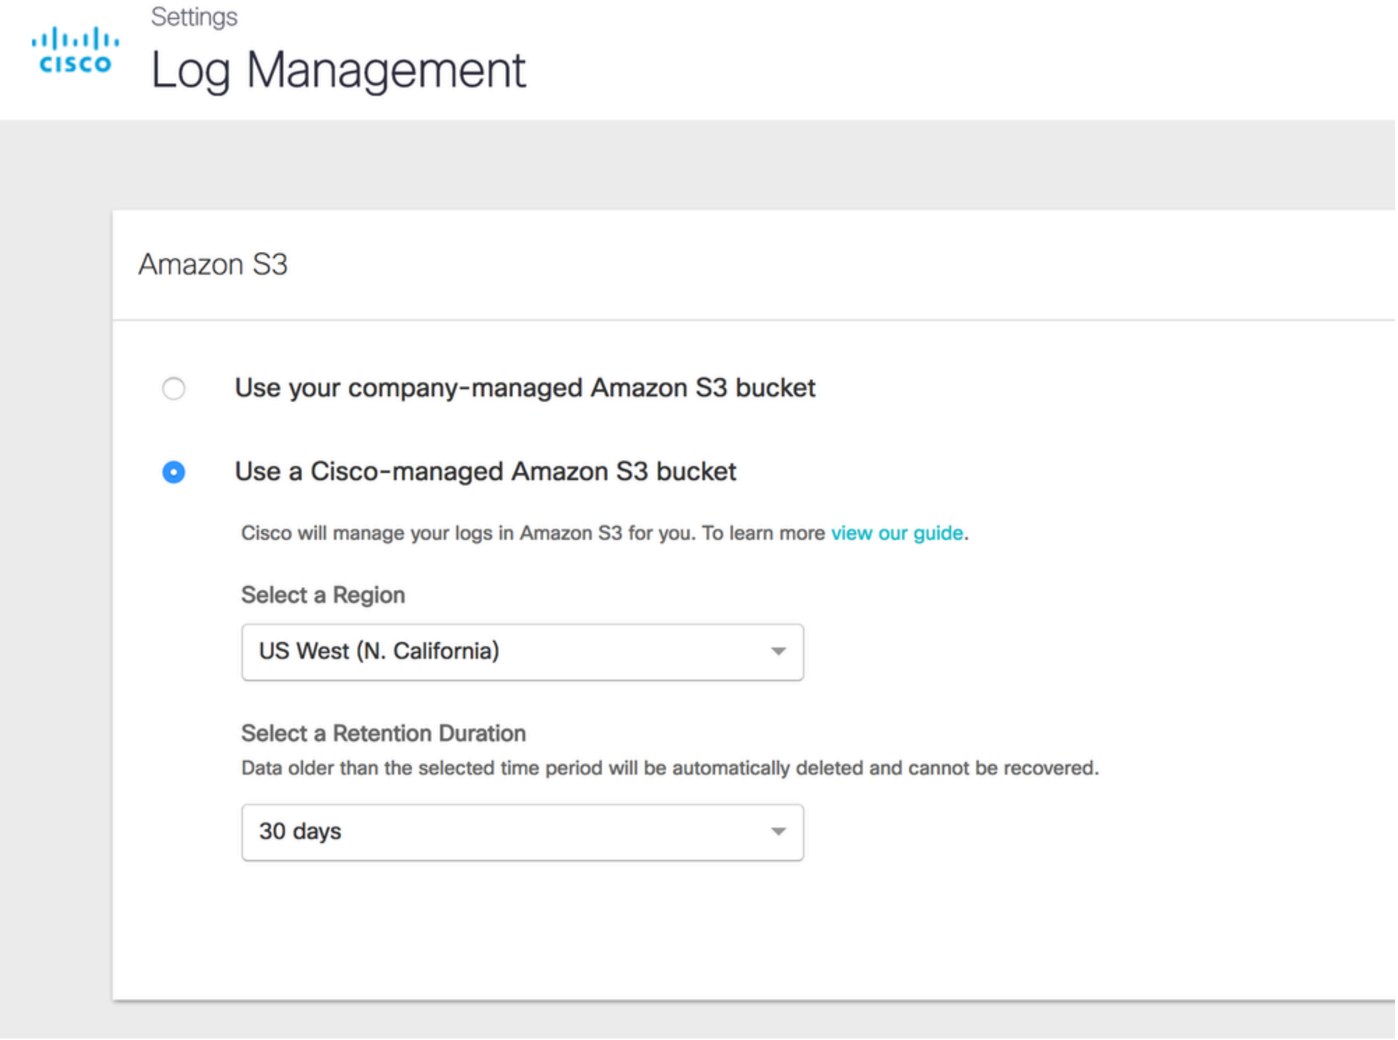

Settings

Log Management

Amazon S3

○ Use your company-managed Amazon S3 bucket

● Use a Cisco-managed Amazon S3 bucket

Cisco will manage your logs in Amazon S3 for you. To learn more view our guide.

Select a Region

US West (N. California)

Select a Retention Duration

Data older than the selected time period will be automatically deleted and cannot be recovered.

30 days

25231151158036

Select a Region

Regional endpoints are important to minimize latency when downloading logs to your servers. The listed regions match those available in Amazon S3, but not all regions are available. For instance, China is not listed.

Pick the region thatit iss closest to you from the dropdown. If you wish to change your region in the future, you need to delete your current settings and start over.

Select a Retention Duration

The retention duration is simply 7, 14 or 30 days. After the selected time period, all data is purged and cannot be retrieved no matter what. We recommend a smaller time period if your ingestion cycle is regular. The retention duration can be changed at a later time.

After you make your selections, click **Next** and you are asked to confirm your region and duration

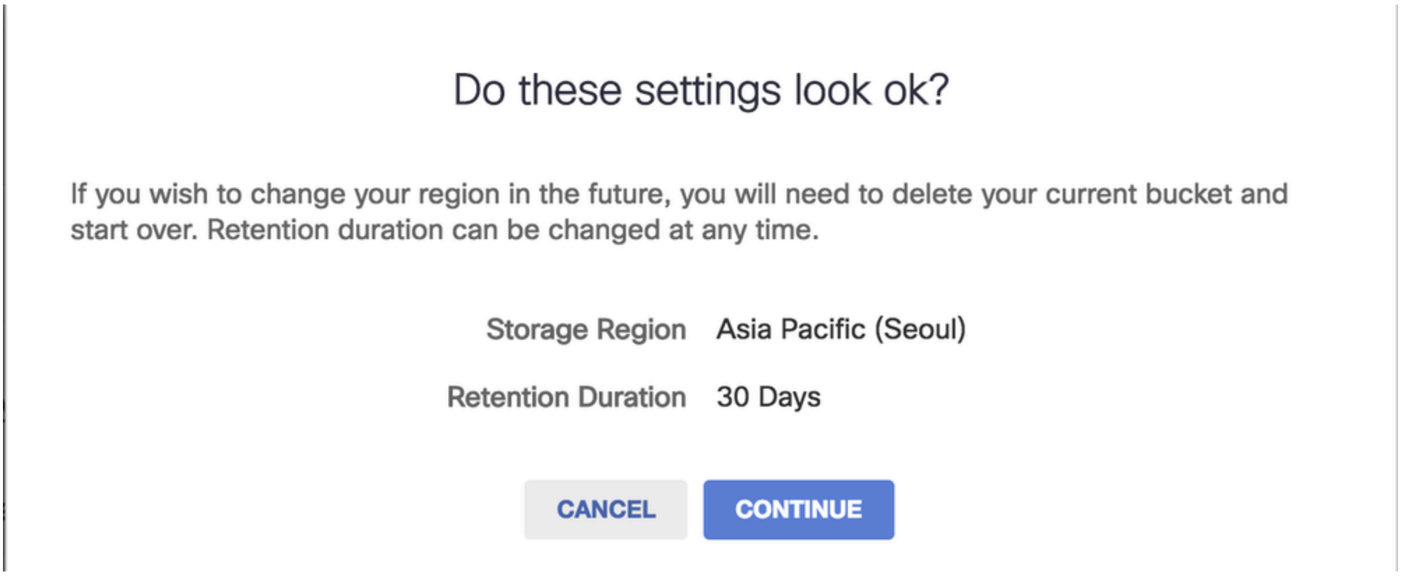

Do these settings look ok?

If you wish to change your region in the future, you will need to delete your current bucket and start over. Retention duration can be changed at any time.

Storage Region Asia Pacific (Seoul)

Retention Duration 30 Days

25231181211796

Once you agree to continue, you get an activation notification.

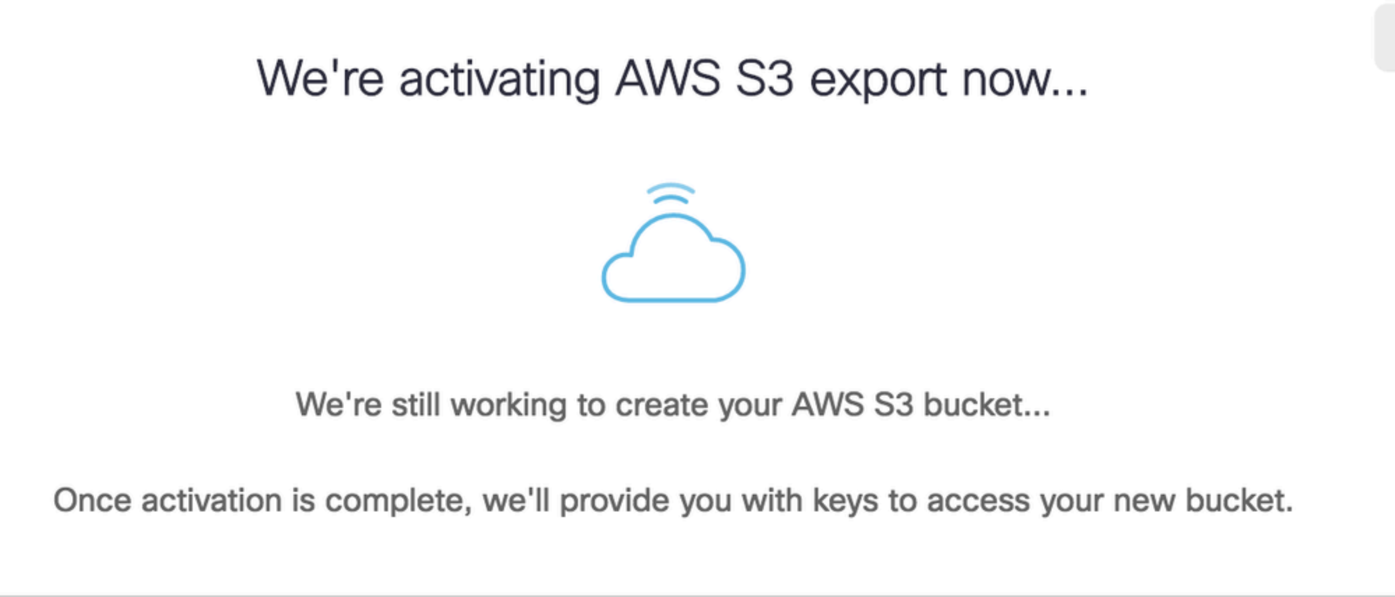

We're activating AWS S3 export now...

We're still working to create your AWS S3 bucket...

Once activation is complete, we'll provide you with keys to access your new bucket.

25231181218708

You then recieve an access key and s secret key. You *must accept* (click "Got it!") because this the *only* time you get to see either of the keys. The access and secret keys are required in order to access your bucket and download your logs.

Lastly, you see the summary screen showing the configuration and most importantly, your bucket name.

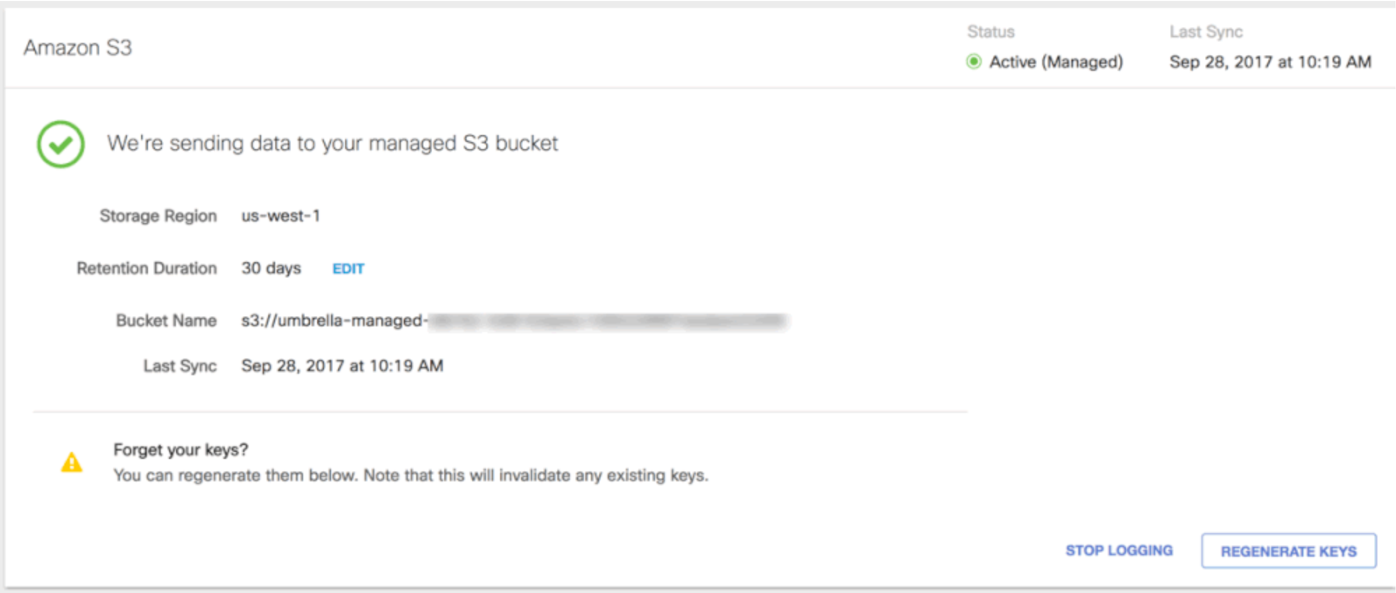

Amazon S3

Status
● Active (Managed)

Last Sync
Sep 28, 2017 at 10:19 AM

✓ We're sending data to your managed S3 bucket

Storage Region us-west-1

Retention Duration 30 days EDIT

Bucket Name s3://umbrella-managed-[redacted]

Last Sync Sep 28, 2017 at 10:19 AM

⚠ Forget your keys?
You can regenerate them below. Note that this will invalidate any existing keys.

25231181228180

You can turn logging on or off at your convenience.

Note: Cisco continues to purge logs based on the selected retention duration, even if logging has been turned off.

Post Configuration Options

Log Upload Failures

In the case of a failure to upload logs from Cisco Umbrella to your S3 bucket, there is a grace period of four hours during which the service retries every 20 minutes. After four hours, a case is opened with our Support team, who begin an investigation as to the cause of the issue and proactively reach out to you to let you know about the problem.

Checking Uploaded Logs and Format

Logs are uploaded in ten-minute intervals from the Umbrella log queue to the S3 buckets. After completing configuration, the first log uploads to your S3 bucket within two hours, although the process is usually immediate or close to immediate. However, uploading anything requires newly-generated log data to exist, so if you are trying this on a test environment, ensure network data is being logged in the Activity Search.

To verify if everything is working, the Last Sync time in the Umbrella dashboard updates and logs begin to appear in your S3 bucket.

Within your bucket, each customer or organization are labelled with their org ID, so the folder structure is:

```
Amazon S3/<bucket-name>/<orgID>/<subfolder>
```

<bucket-name> is your bucket name, **<orgID>** is your organizationit iss ID, and **<subfolder>** are either dnslogs, proxylogs, or iplogs, depending on the types of logs within.

For MSP and MSSP customers, the orgID matches the one in Customer Settings under each customer detail in the deployment parameters section. Multi-org customers can gather the orgID by logging into each individual sub-org and noting the orgID in the browser url: (https://dashboard.umbrella.com/o/#####/).

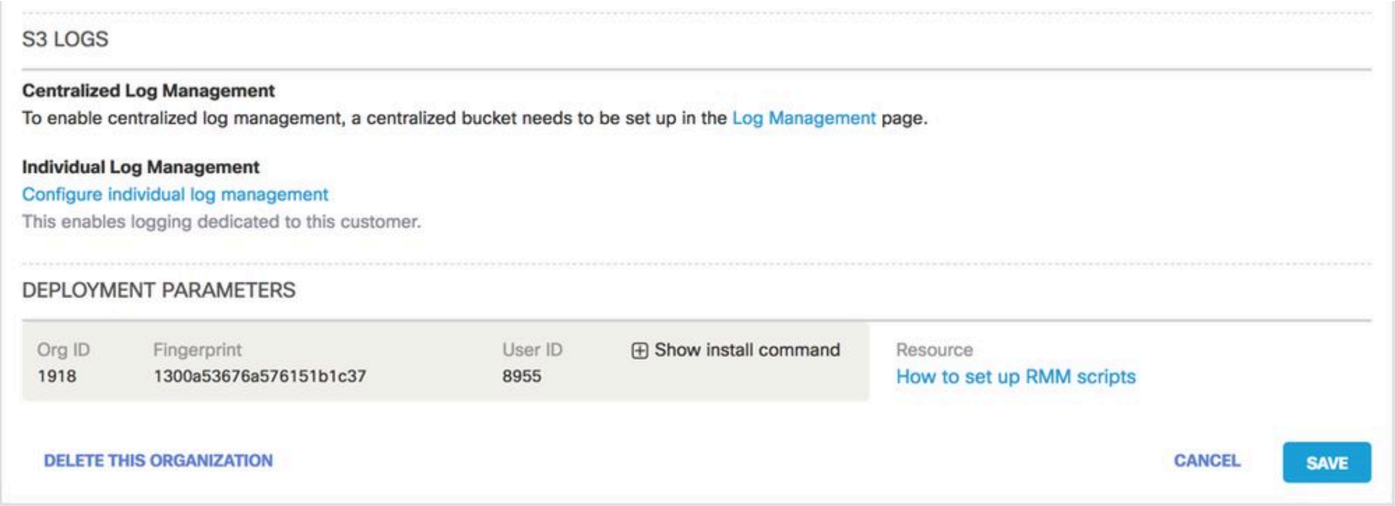

S3 LOGS

Centralized Log Management

To enable centralized log management, a centralized bucket needs to be set up in the Log Management page.

Individual Log Management

Configure individual log management

This enables logging dedicated to this customer.

DEPLOYMENT PARAMETERS

Org ID	Fingerprint	User ID	⊞ Show install command	Resource
1918	1300a53676a576151b1c37	8955		How to set up RMM scripts

Currently, the log format version for the MSP, MSSP, and Multi-org customers is version 1.1. The logs appear in a GZIP format and are uploaded to S3 buckets in the appropriate subfolder with this naming format:

```
<subfolder>/<YYYY>-<MM>-<DD>/<YYYY>-<MM>-<DD>-<hh>-<mm>-<xxxx>.csv.gz
```

<subfolder> is either dnslogs, proxylogs, or iplogs, depending on the types of logs within. **<xxxx>** is a random string of four alphanumeric characters, which prevents duplicate file names from being overwritten.

For example:

```
dnslogs/2019-01-01/2019-01-01-00-00-e4e1.csv.gz
```

If you do not see logs in your bucket within 10 minutes, please contact support outlining the steps you have taken thus far.

Once logs do appear, we recommend reviewing the data by unzipping the contents of the first few log uploads that are received to ensure the data is viewable in a text editor (or even Microsoft Excel, often the default for .CSV). For information on which each field in the log represents read here.

If a log upload from Cisco Umbrella to your S3 bucket fails, there is a four-hour grace period in which the service retries every 20 minutes. After four hours, a case opens within our Support team, who begins an investigation as to the cause of the issue and proactively reach out to you to let you know about the problem.

Enable Logging on a Per-customer Basis

Out of the box, this feature is enabled for all customers unless otherwise specified. The feature can be turned off for individual customers, which is helpful if you have different service levels for customers who do have the feature. This is under each customerit iss settings in the Console. The screenshot in the previous section shows the toggle to turn it off.

It is also possible to create IAM users in Amazon and assign those IAM users to individual orgit iss subfolders of the bucket. By doing so, you can allow an end user access to their logs, but *only* their logs.

Downloading Logs, Understanding the Format and Splunk / QRadar Integration

In order to download the logs for retention or consumption, there are a few approaches to downloading the DNS logs from S3. Weit isve created an article outlining a few approaches to this problem here.

You might also have a few questions about the log format and how it differs slightly from the logs that are displayed in the Umbrella dashboard. For more information about the exported log format, read this article.

Lastly, one of the primary uses for exporting DNS logs is integration with SIEM tools. Although configuration for a SIEM when dealing with logs like this can often come down to an administratorit iss personal preferences, we have some guidance for the most popular SIEMs.

For more information on setting up the Splunk plug-in for Amazon AWS S3 and Umbrella, read here.

For information about configuring IBM QRadar to pull logs from Amazon S3 and digest them, read here.

How Large Are S3 Logs?

The size of your S3 logs depends on the number of events that occur, which is dependent on the volume of your DNS traffic.

You can find the log format for the S3 Logging here.

The example entry is 220 bytes, but the size of each log line varies based on a number of items (length of domain name, number of categories, etc). Assuming each log line is 220 bytes, a million requests would be 220 MB.

To get an estimate of how many DNS queries is seen each day:

1. In the Umbrella dashboard, navigate to **Reporting > Activity Search**.
2. Under **Filters**, run a report for the last 24 hours and then click the **Export CSV** icon.
3. Open the downloaded .csv file. The number of rows (minus one for the header) is the number of DNS queries per day; multiply that by 220 bytes to get the estimate for one day.

In terms of cost, although it is variable, we find that even our most voluminous customers spend only a few dollars a month on the service. One cost is tied to storage time and another is tied to data download from S3 to your environment. Check with Amazon for more details.

As with any of our features, weit isd love to know what you think, especially around SIEM integrations or any additional questions that arenit ist covered in this documentation. If you have any feedback, please let us know!