Download Logs from Umbrella Log Management in AWS S3

Contents

Introduction

This document describes how to download logs from the Umbrella Log Management in AWS S3.

Overview

Once you set up and test that Log Management in the Amazon S3 is working correctly, you might wish to begin automatically downloading and storing the logs within your network infrastructure, either for retention or consumption (or both).

In order to do this, we have outlined an approach using s3tools from http://s3tools.org. s3tools uses the s3cmd command line utility for Linux or OS/X. There are other tools that can accomplish a similar function for Windows users:

- For a command line tool, you can download a small command line executable here.
- If you prefer a graphical interface, check out S3 Browser (https://s3browser.com/), although we are not covering how to use it because the graphical interface is not scriptable to automate the process. This article gives you steps to setup both command line tools. You can use the information in stage 1 to configure the s3browser application if you prefer.

Start by downloading the tool for the operating system you intend to use. For now, we are just covering s3cmd for OS/X and Linux, although the steps to access your bucket and download the data are effectively the same for Windows.

Grab the installer from s3tools here.

The intaller does not require you to install the program to run the command line, so simply extract the

package you have downloaded.

Stage 1: Configuring Your Security Credentials in AWS

Step 1

1. Add an access key to your Amazon Web Services account to enable remote access to your local tool and the ability to upload, download and modify files in S3. Log into AWS and click your account name in the upper-right hand corner. In the drop-down, choose **Security Credentials**.
2. A prompt instructs you to use Amazon Best Practices and create an AWS Identity and Access Management (IAM) user. In essence, an IAM user ensures that the account that s3cmd uses to access your bucket is not the primary account (for example, your account) for your entire S3 configuration. By creating individual IAM users for people accessing your account, you can give each IAM user a unique set of security credentials. You can also grant different permissions to each IAM user. If necessary, you can change or revoke an IAM user's permissions at any time.
   For more information on IAM users and AWS best practice, read here.

Step 2

1. Click **Get Started with IAM Users** to create an IAM user with access your S3 bucket. Navigate to a screen where you can create an IAM User.
2. Click **Create New Users** and fill out the fields.
3. After creating the user account, you have only one opportunity to grab two critical pieces of information containing your Amazon User Security Credentials. **We highly suggest you download these using the button in the lower right to back them up. They are not available after this stage in the setup.** Ensure you make a note of both your Access Key ID and Secret Access Key, as we need them in a later step.

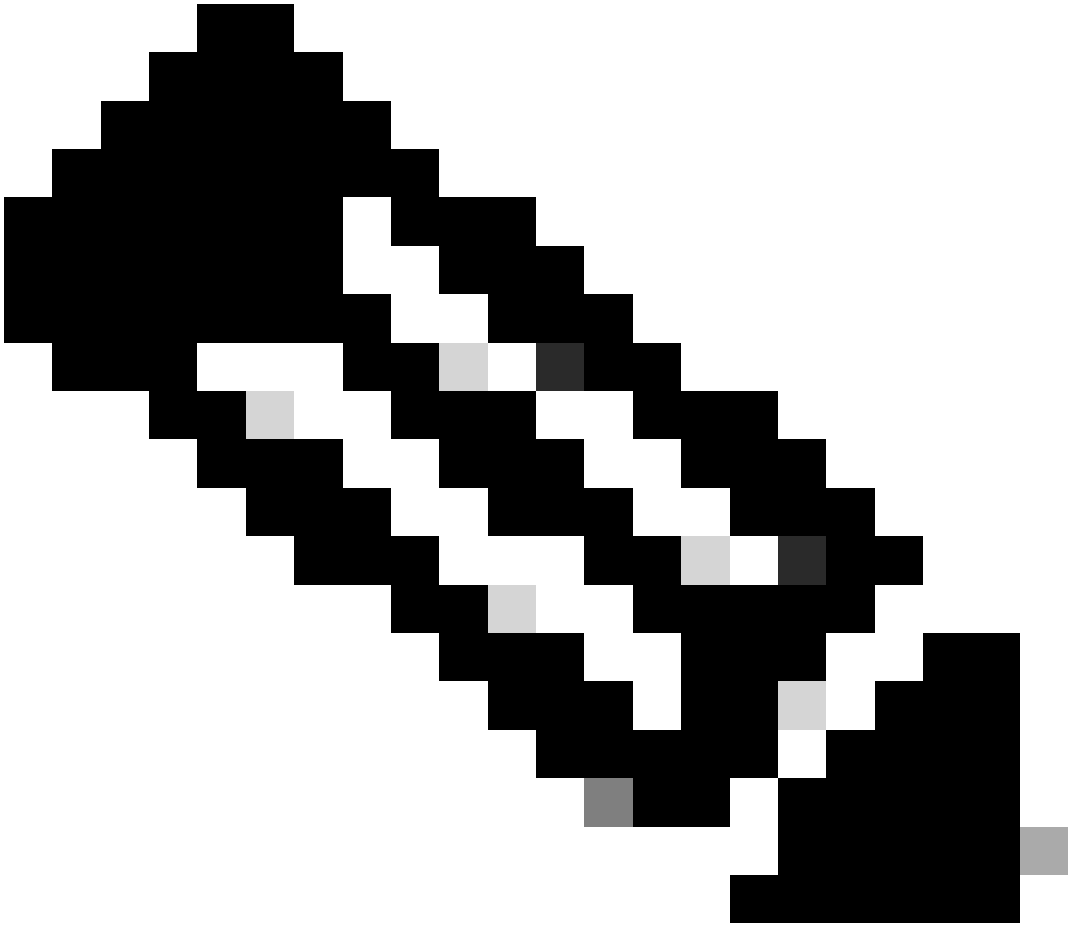

Note: The user account cannot contain spaces.

Step 3

1. Next, you want to add a policy for your IAM user so they have access to your S3 bucket. Click the user you have just created and then scroll down through the users' properties until you see the Attach Policy button.
2. Click **Attach Policy**, then enter 's3' in the policy type filter. This should show two results "AmazonS3FullAccess" and "AmazonS3ReadOnlyAccess".
3. Select **AmazonS3FullAccess** and then click **Attach Policy**.

Stage 2: Configuring a Tool to Download DNS Logs From the Bucket

s3cmd for MacOS and Linux

1. Go to the path you have extracted the s3cmd in the previous stage and from Terminal, type:

```
./s3cmd --configure
```

This should bring you to a prompt requesting that you provide your security credentials:

Enter new values or accept defaults in brackets with Enter.

Refer to user manual for detailed description of all options.

Access key and Secret key are your identifiers for Amazon S3. Leave them empty for using the env variables.

Access Key [YOUR ACCESS KEY]:

Secret Key [YOUR SECRET KEY]:

2. Next, you be asked a series of questions regarding how you would like to configure access to your bucket. In this case, we are not setting up an encryption password (GPG), and we are not using HTTPS or a proxy server. If your network or preferences differ, fill in the required fields:

Default Region [US]:

Encryption password is used to protect your files from reading by unauthorized persons while in transfer to S3

Encryption password:

Path to GPG program [None]:

When using secure HTTPS protocol all communication with Amazon S3 servers is protected from 3rd party eavesdropping. This method is

slower than plain HTTP, and can only be proxied with Python 2.7 or newer

Use HTTPS protocol [No]:

On some networks all internet access must go through a HTTP proxy.

Try setting it here if you can't connect to S3 directly

HTTP Proxy server name:

After entering any network-specific settings or any encryption, you have a chance to review:

New settings:

Access Key: YOUR KEY

Secret Key: YOUR SECRET KEY

Default Region: US

Encryption password:

Path to GPG program: None

Use HTTPS protocol: False

HTTP Proxy server name:

HTTP Proxy server port: 0

Lastly, you are asked to test and if successful, save the settings:

Test access with supplied credentials? [Y/n] y

Please wait, attempting to list all buckets...

Success. Your access key and secret key worked fine ��

Now verifying that encryption works...

Not configured. Never mind.

Save settings? [y/N]

Windows Command Line executable (s3.exe)

After downloading the tool (https://s3.codeplex.com/releases/view/47595), copy the .exe to your preferred working folder and from the command prompt type this, substituting your access key and secret:

<#root>

```
s3 auth [<YOUR ACCESS KEY> <YOUR ACCESS SECRET>]
```

For more information on authentication syntax, read here.

Stage 3: Testing the Download of Files From Your Bucket

Step 1: Test the download

s3cmd for OS/X and Linux

From the terminal, run this command where "my-organization-name-log-bucket" is the name of your bucket already configured in the Log Management portion of the Umbrella dashboard. In this example, this is run from the folder that contains the s3cmd executable and the files are delivered to the same path, but these can be changed:

<#root>

```
./s3cmd sync s3://my-organization-name-log-bucket ./
```

If there is a difference between the files in your bucket and the files in the destination path on disk, the sync should download the missing or updated files. The first file retrieved should be the README file that is typically uploaded:

```
./s3cmd sync s3://my-organization-name-log-bucket ./

 s3://my-organization-name-log-bucket/README_FROM_UMBRELLA.txt ->
<fdopen>  [1 of 1]

1800 of 1800   100% in    0s    15.00 kB/s  done

 Done. Downloaded 1800 bytes in 1.0 seconds, 1800.00 B/s
```

Any log files that are present are also downloaded. It is up to you if you would like to set a cron job to schedule this function on a regular basis, but you should now be able to automatically download any new or changed log files in your bucket to a local path for long-term retention.

Windows Command Line executable (s3.exe)

From the command prompt, run this command where 'my-organization-name-log-bucket' is the name of your bucket already configured in the Log Management portion of the Umbrella dashboard. In this example, all files in the bucket (defined with the asterisk wildcard) are downloaded to the \dnslogbackups\ folder.

<#root>

```
s3 get my-organization-name-log-bucket/* c:\dnslogbackups\
```

For more information about the syntax for this command, read here.

Step 2: Automate the download

Once the syntax has been tested and works as expected, copy the instructions into a script setup a cron job (OS X / Linux) or a scheduled task (Windows) or use any other task automation tool you might have at your disposal. It is also possible using the tools to remove files from your bucket after you have downloaded them to free up space in your S3 instance. We encourage you to look at the documentation for the tool you are using to see what might work best for your data retention policy.