



# **Antivirus for Amazon S3 Deployment Guide**

Version 3.0, October 2024

## Table of Contents

<b><i>Introduction</i></b>	<b>4</b>
<b><i>Prerequisites and Requirements</i></b>	<b>4</b>
Time	4
Product License	4
AWS Account	5
AWS Identity and Access Management (IAM) Entity	5
Knowledge Requirements	5
<b><i>Architecture</i></b>	<b>7</b>
Single Region Architecture	7
Multi-Region Architecture	8
Multi-Account Architecture	9
Deployment	10
Deploy the Software	13
<b><i>Security</i></b>	<b>19</b>
IAM Roles	19
Console Access	19
Outbound Access for Console and Agents	19
API Access to the Console and Scanning Agents	20
Encrypting Amazon DynamoDB Data at Rest	20
Encrypting Amazon SNS and SQS	20
Amazon S3 Object Encryption	20
Logging/Auditing	21
<b><i>Costs</i></b>	<b>21</b>
Infrastructure Components	21
Example Infrastructure Costs	22
AWS API Calls – the other infrastructure costs	22
Additional Cost Considerations	22
<b><i>Sizing</i></b>	<b>23</b>
Questions to Ask	23
Sizing Example	23
<b><i>Testing and Deployment</i></b>	<b>24</b>

<b><i>Health Check</i></b>	<b>25</b>
Verify Up and Running	25
Verify Scanning	26
Keeping Up with the Work	27
<b><i>Backup and Recovery</i></b>	<b>28</b>
Backup	28
Instance Failure	28
Availability Zone Failure	29
Region Failure	29
General Failure Considerations	29
<b><i>Routine Maintenance</i></b>	<b>30</b>
<b><i>Emergency Maintenance</i></b>	<b>30</b>
<b><i>Support</i></b>	<b>31</b>
Troubleshooting	31
FAQ	31
Clean Up	31
Contact Us	31

## Introduction

Cloud Storage Security allows users to detect Amazon S3 files that are infected with malware and viruses. When infected files are uploaded to Amazon S3, Cloud Storage Security detects the malicious files and guards against the spread of malware.

This guide will walk you through the deployment, via AWS CloudFormation template, of the Antivirus for Amazon S3 product. The AWS CloudFormation template deploys a highly scalable, decoupled solution that runs on AWS Fargate and leverages a number of additional Amazon services (SNS, SQS, etc).

## Prerequisites and Requirements

In this section, we detail the prerequisites and requirements to run and operate our solution. This will be a straight-forward process for you.

### Time

The initial deployment will take less than 10 minutes including Marketplace signup and AWS CloudFormation run through. The initial setup of bucket protection will take less than 5 minutes.

Additional thoughts will be needed to determine which buckets should be protected (whether the entire bucket or sub-paths) and whether or not you are protecting buckets only in the deployment account or will additionally be adding other accounts to be protected (linking accounts through cross-account roles). These additional processes are easy to complete, but do need to be thought through.

### Product License

You are required to have an [AWS Marketplace](#) subscription to use the Antivirus for Amazon S3 product. You can deploy it using our [free trial](#), which allows you to use the product for a set amount of time or data consumption without incurring software charges—although you will still incur infrastructure charges. There are two options to pay for the software:

- Pay as you go (PAYG) – [Subscribe to our product](#) through the AWS Marketplace. Once your free trial has ended, the software will report the number of GBs scanned to AWS, so that AWS can charge you based on the [AWS Marketplace pricing](#).
- Bring your own license (BYOL) – Subscribe through the AWS Marketplace and apply an existing license. You can view details about your license and upload new license files by navigating to the [License Management page](#) of your Antivirus for Amazon S3 console. This option is appropriate for deployments with extremely heavy usage, abnormal access patterns, private deployments that cannot speak to AWS Marketplace and for GovCloud. Please email [sales@cloudstoragesec.com](mailto:sales@cloudstoragesec.com) to order a license directly from our sales department.

## AWS Account

You must have an AWS account to use our software. If you don't, review the Amazon instructions on [how to create an AWS account](#).

## AWS Identity and Access Management (IAM) Entity

You must also [create an IAM user or role](#). Your IAM user should have a [policy](#) that allows AWS CloudFormation actions. Do not use your root account to deploy the CloudFormation template. In addition to AWS CloudFormation actions, IAM users who create or delete stacks will also require additional permissions that depend on the stack template. This deployment requires permissions to all services listed in the following section.

## Knowledge Requirements

The following services are deployed by the CloudFormation template or created by the Console (via the ConsoleRole the CloudFormation creates) post deployment. Detailed knowledge of each of these services is not required to run and operate the solution, but you should be familiar with your AWS account and CloudFormation. Antivirus for Amazon S3 does utilize each of these services, but to simplify the experience for you we have abstracted the need to directly interact with most of these services.

Created by the AWS CloudFormation template:

- [AWS CloudFormation](#)
- [Amazon Virtual Private Cloud \(Amazon VPC\)](#)
- [Amazon VPC Security Groups](#)
- [Amazon IAM Roles](#)
- [AWS Fargate](#)
- [Amazon App Config](#)
- [Amazon CloudWatch](#)
- [Amazon DynamoDB](#)
- [Amazon Simple Notification Service](#)
- [AWS Simple Queue Service](#)
- [Amazon Cognito](#)
- [AWS Key Management Service \(AWS KMS\)](#)

The application console will be responsible for setting up additional regions based on the bucket protection you configure. Each region that contains protected buckets will get these services deployed by the console.

Created by the Console post CloudFormation template completion:

- [AWS Fargate](#)
- [Amazon App Config](#)
- [Amazon CloudWatch](#)
- [Amazon Simple Notification Service](#)
- [AWS Simple Queue Service](#)
- [AWS Key Management Service \(AWS KMS\)](#)

Additional AWS Accounts linked to the deployment will have the following resources that are created by a generated CloudFormation Template unique for each account.

- [Amazon IAM Roles](#)
- Amazon S3 Event Listener

You may need to request an increase in [permissions or additional policies](#) for your AWS account for this deployment. Our [IAM Permissions Review](#) provides more details so you can determine if your account has all the rights required to complete the deployment.

**Note:** Individuals possessing the [AWS Associate Certifications](#) should have a sufficient depth of knowledge to deploy the resources specified in this guide.

## Architecture

The nature of the architecture already lends itself to be highly available and highly scalable. The “brains” and the “brawn” have been decoupled from one another so container tasks can come and go without missing a beat.

To ensure high availability, you should always define at least two [public or private subnets](#) each in different availability zones for the Console configuration as well as the Agent configuration. This naturally creates a highly available solution as the scaling policies around both the Console and the Agent will bring the tasks back online. If one Availability Zone (AZ) goes down, they will naturally come up in the Subnet defined in the second AZ.

### Note:

The console and the agent do not require public IPs, but they do require outbound internet access to get to the [Amazon Elastic Container Registry](#) and all other AWS services listed in the previous section as well as for AV signature updates. Outbound routing can be accomplished through an [AWS Internet Gateway](#) (**not** private), an [AWS NAT Gateway](#) / [AWS Transit Gateway](#) or through [VPC Endpoints](#) + Proxy. A proxy is required because there are three services leveraged that do not have VPC Endpoints at this time: Amazon Marketplace, Amazon AppConfig and Amazon Cognito. Technically, you can use VPC Endpoints with a NAT, but most customers we’ve seen go down this road leverage a proxy.

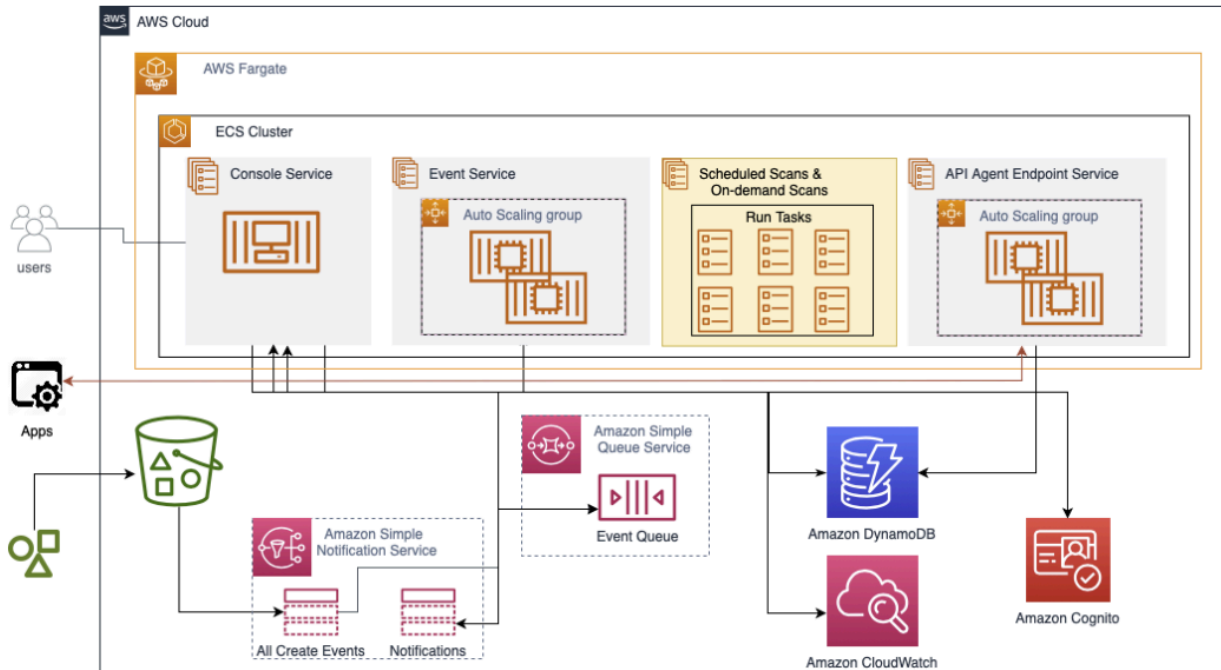
You must also ensure you have a path to the Console. If the [AWS Security Groups](#) or the subnet doesn’t allow public access, you must ensure you do have access from the network you are running.

**Note:** If you are leveraging an AWS Internet Gateway, then you **must** have public IPs assigned. AWS IGW will **not** work without public IPs.

Review the below diagrams to see how our architecture supports object flow in a single or multiple regions or multiple accounts.

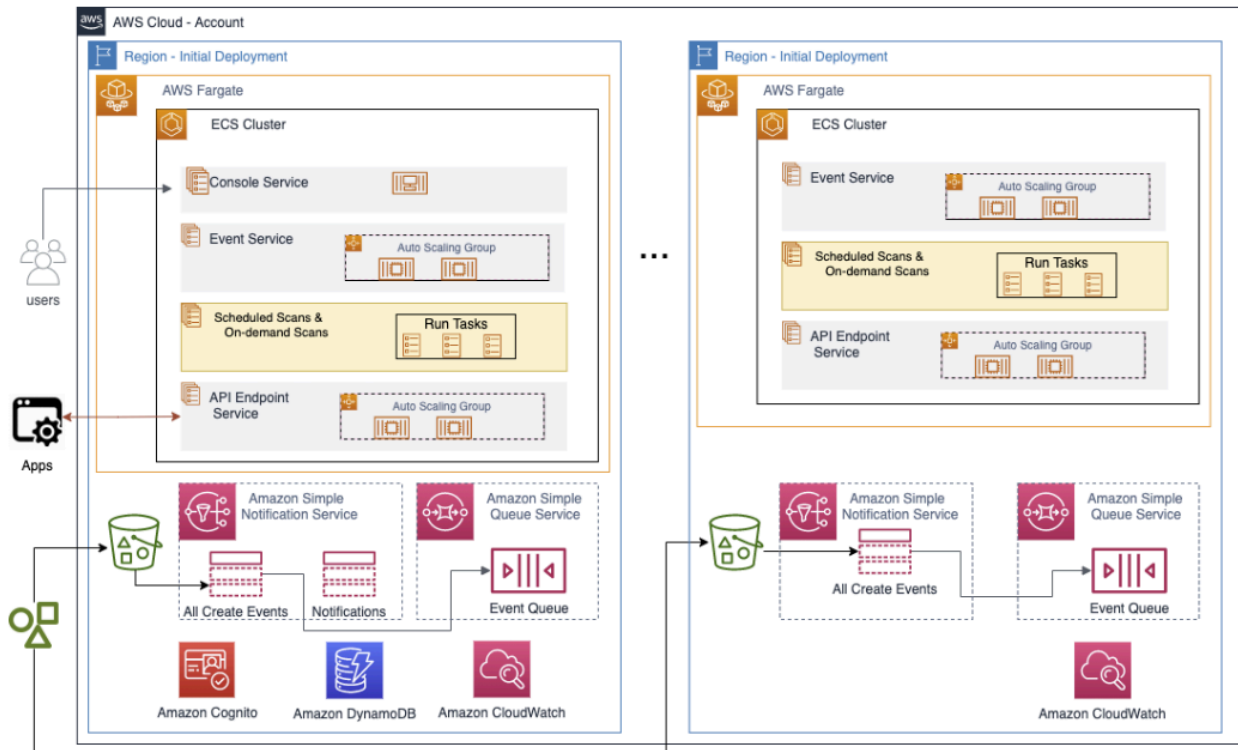
### Single Region Architecture

The initial deployment region will contain the Console Fargate Service and optionally one or all of the various object scanning services.




### Multi-Region Architecture

The initial deployment region will contain the Console Fargate Service and optionally the object scanning services. Additional regions will only contain the various scanning services that are specified based on your protection strategies.



## Multi-Account Architecture

It is possible and often desired to centrally locate your security software and leverage it for multiple accounts (related, within an Organization, etc). Antivirus for Amazon S3 has been designed to support this. You will deploy the product in the central account as you normally would. You will then [link accounts](#) and deploy a cross-account role into the linked accounts via CloudFormation to enable those accounts for cross-account scanning. The linked accounts will not have any infrastructure deployed to them.



## Deployment Options

We offer several flexible deployment options based on what you need:

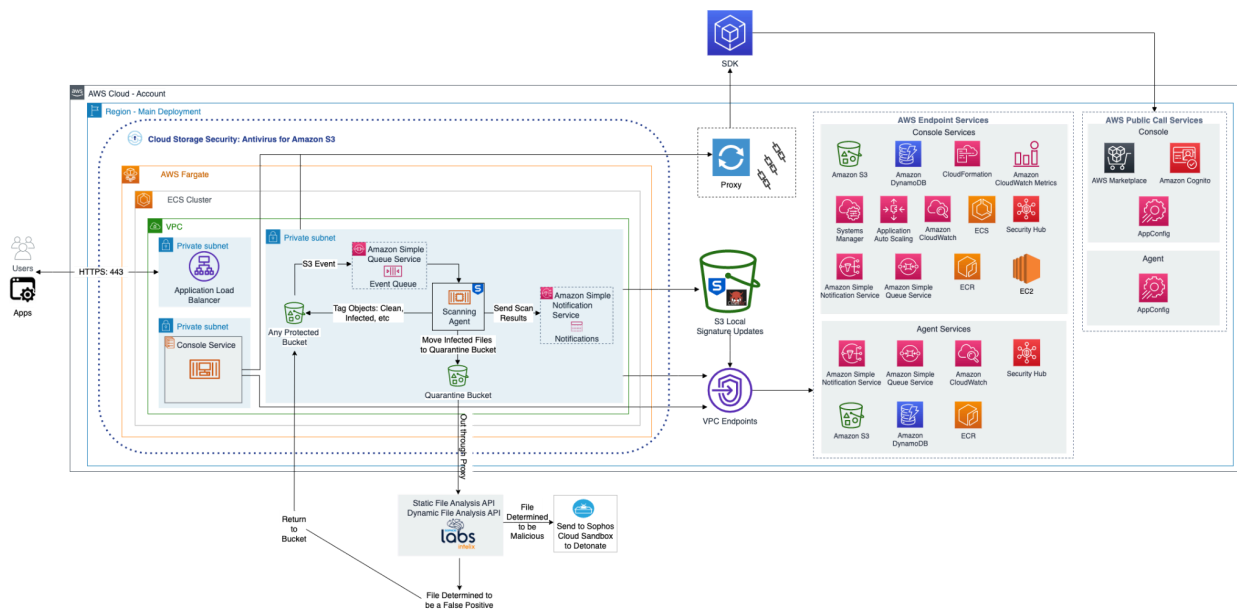
- **Standard Deployment:** requires filling out only 5 fields in the CloudFormation Template.

- Private Deployment: all our components run in private VPCs and private Subnets with no public IPs assigned at all.
- An in-between option: where you have public access (using a public Load Balancer) while still running all solution components in private VPC/Subnets.

You can mix and match as well as incorporate VPC Endpoints to keep as much traffic as possible going over the AWS backbone.

Learn more about our deployment options by viewing our [help docs](#).

The below diagram illustrates the most private deployment option available for your environment.



## Deployment

AWS CloudFormation ensures that your deployment is successful by providing you with a template that will install the necessary infrastructure components, required IAM roles, and permissions.

The CloudFormation Template will create the following resources:

Resource	Purpose
ECS Fargate Cluster with 1 Service and Task 2 Task Definitions	Runs the Antivirus for Amazon S3 Management Console
DynamoDB; AppConfig	Saves the software data
IAM Roles and Policies	Enable you to run the software and access the data
CloudWatch Logs and Metrics	Logs - track all activity for the system as well as every file touched  Metrics – manage the auto-scaling for scanning services
Cognito UserPool	Helps you manage users

Once the Management Console is running, it will create the following resources through code:

Resource	Purpose
One or two Services and Tasks (Scanning Agents – Event and/or API) to existing region cluster	Runs the scanning agents that process the objects

One ECS Fargate Cluster, Service(s) and Task(s) in each additional region you scan buckets	Run the scanning agents in new regions
SNS Topic, SQS Queue, S3 Bucket events in each additional Region	Help you keep track of the object work These will be deployed in each Region with protected buckets
CloudWatch Logs and Metrics in each additional Region	Logs - track all activity for the system as well as every file touched  Metrics – manage the auto-scaling for scanning services

This section will help you configure and run the CloudFormation template.

**Important:** When you launch the Cloud Storage Security Antivirus for Amazon S3 service from the AWS Marketplace, you will be launched in the **us-east-1** region by default.

If you want to deploy to another region, please change to one of the following supported regions before continuing:

- US East (N. Virginia)
- US East (Ohio)
- US West (Oregon)
- US West (Norther California)
- Asia Pacific (Mumbai)
- Asia Pacific (Seoul)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Asia Pacific (Singapore)
- Canada (Central)
- EU (Frankfurt)
- EU (Ireland)
- EU (London)
- EU (Paris)
- EU (Stockholm)
- Middle East (Bahrain)
- South America (Sao Paulo)
- AWS GovCloud (US-West)

**NOTE:** All components deployed, created and installed run inside of your account. We do not host any of them and we never send any of your objects/files outside of your account. All scanning is performed close to the data inside your account(s) and in-region.

Our console requires specific permissions to manage its own infrastructure and integrate with your AWS services. To provide the most secure environment, we recommend deploying it in a dedicated AWS account. This isolates the product's permissions and prevents any unintended impact on other resources.

As you'll learn below, we may send some data (IP address, account email, version numbers) to a Cloud Storage Security AWS account to assist you and your users with accessing your application. You can opt out of this if you don't want that information to be reported.

Deploying the software requires the use of multiple AWS Services. In some cases (VPC for example) you can reuse previously created entities, but you may also choose to create new entities specifically for the solution. Please evaluate your [AWS Service Quotas](#) in each region you plan to deploy a component of the product. You can also go to the [Services Quotas](#) centralized view within the AWS Console to see what your current limits are at and make changes. One item we see customers run into is bucket count. The typical limit is 100 buckets. The solution will create additional buckets for quarantining objects. *Be aware if you are close to that limit.*

Items to evaluate:

- VPCs
- Internet Gateways
- Subnets
- DynamoDB – tables per region
- CloudFormation – stack counts
- SNS – filter policy counts
- ECS Fargate

**Pre-requisites needed to deploy the solution:**

- Overall
  - Subscription to marketplace product in the deployment account (<https://aws.amazon.com/marketplace/pp/prodview-q7oc4shdnpc4w>)
- User deploying CFT has permissions to do so
  - Execute CFT, create roles and resources
- Outbound access to AWS services

- VPC and Subnets must have outbound access or VPC Endpoints to access key AWS services (like ECR, etc). Typically done through a Internet Gateway (for public VPCs) or a NAT Gateway (for private VPCs)
- Console & Event Scanners
  - VPC
    - 2 Subnets - separate AZs
    - CIDR range to support volume
  - OPTIONAL - Console Only – SSL Certificate for Load Balancer front ending console GUI  
(<https://help.cloudstoragesec.com/how-it-works/deployment-details/#private-deployment>)
- API Endpoint (if using APIs)
  - VPC (can be same as above)
    - 2 subnets - separate AZs (can be same as above)
    - CIDR range to support volume
  - SSL Certificate for Load Balancer front ending API Endpoint - cert can be same as above if wild carded, if not then a new cert is needed  
(<https://help.cloudstoragesec.com/how-it-works/object-scanning/#api-driven-scanning>)

## Deploy the Software

1. Once you have successfully subscribed to the Cloud Storage Security Antivirus for Amazon S3 service, the Quick Create Stack wizard will open in CloudFormation.
2. Under the Stack name section, type a name for your stack (a default is provided).

## Quick create stack

### Template

Template URL

<https://css-cft.s3.amazonaws.com/ConsoleCloudFormationTemplate.yaml>

Stack description

Version 6.02.002 - CloudFormation template for CloudStorageSec Console software.

### Stack name

Stack name

Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

3. Scroll down to the Parameters section. CloudFormation auto-populates many of the options, which you can leave in their default state.
4. Under Network Configuration, configure the following settings:
  - a. From the Virtual Private Cloud (VPC) ID drop-down menu, select a VPC.
  - b. From the Subnet A ID drop-down menu, select a subnet.
  - c. From the Subnet B ID drop-down menu, select a subnet.
  - d. In the Console Security Group CIDR Block field, type your network access or type **0.0.0.0/0** so that you can access the Amazon S3 Management Console from any location.

**NOTE:** We do not recommend utilizing 0.0.0.0/0 as your long term strategy. You should always lock this down to the network(s) your users will typically access from. You can modify this value in the Console Security Group via the AWS Management Console at any time after deployment.

Here is an example of what your screen may look like after you have configured these options:

**Parameters**  
Parameters are defined in your template and allow you to input custom values when you create or update a stack.

**Network Configuration**

**Virtual Private Cloud (VPC) ID**  
The VPC in which to place the public facing Console

vpc-5cf1e526 (172.31.0.0/16) ▼

**Subnet A ID**  
A subnet in your VPC in which the public facing Console can be placed

subnet-2a24dc4c (172.31.0.0/20) ▼

**Subnet B ID**  
A subnet in your VPC in which the public facing Console can be placed

subnet-feec67b5 (172.31.16.0/20) ▼

**Console Security Group CIDR Block**  
The IP address range that can access the Console management website (e.g. X.X.X.X/24 for a single given IP, 0.0.0.0/0 for open access)

0.0.0.0/0

**Warning:** To create your CloudFormation successfully, you must ensure that Subnet A and Subnet B are different and that both subnets belong to the same VPC.

5. Under Console Configuration, configure the following settings:
  - a. From the Console VCPU drop-down menu, select the number of VPC units for the console.
  - b. From the Console Memory drop-down menu, select the amount of memory for the console.
  - c. In the Email address field, type the email address you want to associate with the console.

Here is an example of what your screen may look like after you have configured these options:

**Console Configuration**

**Console vCPU**  
The number of vCPU units for the Console (1024 vCPU units per 1 vCPU)

0.5vCPU

**Console Memory**  
The amount of memory for the Console. Must be 2-8x the vCPU units.

1GB

**UserName**  
Initial user name for the Console management website

admin

**Email**  
Email address for Console management website account

support@cloudstoragesec.com

**Warning:** You must select a memory amount in the range of 2x to 8x of the vCPU or the stack creation will fail.

**Note:** We recommend leaving the defaults for vCPU and Memory to get started. If you notice the Console lagging as you grow, you can increase these settings at that time.

6. Under Agent Sizing Configuration, configure the following settings:
  - a. From the Agent VCPU drop-down menu, select the number of VPC units for the agents.
  - b. From the Agent Memory drop-down menu, select the amount of memory for the agents.

Here is an example of what your screen may look like after you have configured these options:

**Agent Configuration**

**Agent vCPU**  
The number of vCPU units for the Agents (1024 vCPU units = 1 vCPU)

1vCPU


**Agent Memory**  
The amount of memory for the scanning Agent. Must be 2-8x the vCPU. i.e. 2 vCPU should have 4-16GB of memory.

3GB

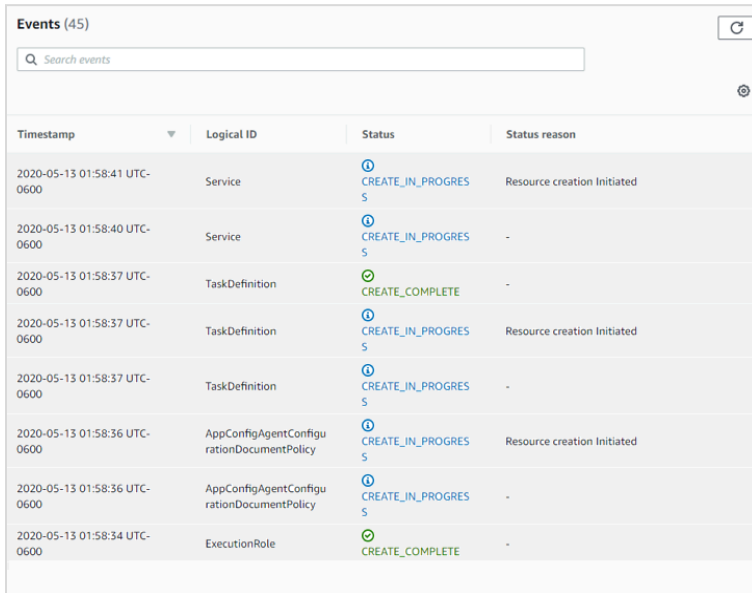
**Warning:** You must select a memory amount in the range of 2x to 8x of the vCPU or the stack creation will fail.

**Note:** Leave the defaults. For malware scanning there is NO benefit to increasing beyond 1 vCPU and 3GB of memory.

7. Other choices and considerations within the CloudFormation Template
  - a. We'll cover some of the choices here, but please review the [Advanced Considerations](#) within the Help Docs
  - b. Scanning Engine and Multi-Engine Scan
    - i. You have a choice amongst several engines for scanning your data
    - ii. You also have the choice to scan with multiple engines
  - c. Disabling Public Ips
    - i. AWS will auto-assign public IPs (even when running in a private subnet) unless you select Disable
    - ii. You can disable public IPs for the Console or the Scanning Agent or both. You could keep them public for the Console and set to private for the Scanning Agent  
**Note:** if you have private IPs you must have some form of outbound routing still (i.e. NAT Gateway)
  - d. Optional Load Balancer
    - i. This is often used for a hybrid public/private deployment or for a persistent access mechanism for a private deployment
    - ii. The other main advantage of this option is to be able to **Opt-Out** of any information sharing with Cloud Storage Security. Information we collect: email address for admin, IP address, product versions, number of files scanned, scan engine selected  
**Note:** Cloud Storage Security will **never** receive any information about the data you are scanning: file names or the data itself. Your data never leaves your set of accounts
  - e. Proxy Settings
    - i. This is required when using VPC Endpoints to allow for the 3 services (Amazon Marketplace, Amazon Cognito, Amazon AppConfig) which do not have Endpoints
8. Scroll down to the Capabilities section, select the **I Acknowledge that AWS CloudFormation might create IAM resources with custom names** check box, and click **Create stack**.



The next screen displays your stack status, which changes from **CREATE\_IN\_PROGRESS** to **CREATE\_COMPLETE** once AWS finishes creating the stack:

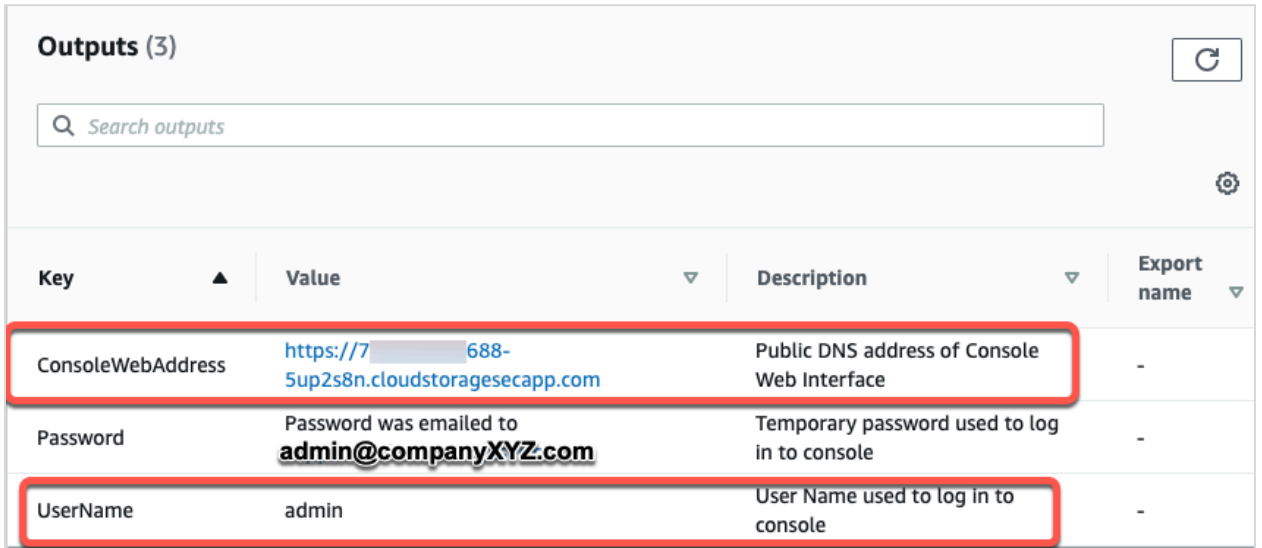


Timestamp	Logical ID	Status	Status reason
2020-05-13 01:58:41 UTC-0600	Service	CREATE_IN_PROGRESS	Resource creation Initiated
2020-05-13 01:58:40 UTC-0600	Service	CREATE_IN_PROGRESS	-
2020-05-13 01:58:37 UTC-0600	TaskDefinition	CREATE_COMPLETE	-
2020-05-13 01:58:37 UTC-0600	TaskDefinition	CREATE_IN_PROGRESS	Resource creation Initiated
2020-05-13 01:58:37 UTC-0600	TaskDefinition	CREATE_IN_PROGRESS	-
2020-05-13 01:58:36 UTC-0600	AppConfigAgentConfigurationDocumentPolicy	CREATE_IN_PROGRESS	Resource creation Initiated
2020-05-13 01:58:36 UTC-0600	AppConfigAgentConfigurationDocumentPolicy	CREATE_IN_PROGRESS	-
2020-05-13 01:58:34 UTC-0600	ExecutionRole	CREATE_COMPLETE	-

**Note:** If your stack creation fails, go to the [Troubleshooting](#) section to see if you can diagnose what went wrong during your deployment.

9. Click the **Output** tab.

The top row in the table will have the URL for the Antivirus for Amazon S3 Management Console in the format of **https://<accountID-appID>.cloudstoragesecapp.com**:



Key ▲	Value ▼	Description ▼	Export name ▼
ConsoleWebAddress	<a href="https://7688-5up2s8n.cloudstoragesecapp.com">https://7688-5up2s8n.cloudstoragesecapp.com</a>	Public DNS address of Console Web Interface	-
Password	Password was emailed to <b>admin@companyXYZ.com</b>	Temporary password used to log in to console	-
UserName	admin	User Name used to log in to console	-

Now that you have deployed the Cloud Storage Security Antivirus for Amazon S3 service, you can learn how to [configure the anti-malware detection settings](#).

## Security

In this section, we discuss Antivirus for Amazon S3 security considerations. Please note that pursuant to AWS's shared responsibility anyone who deploys Antivirus for Amazon S3 to their environment is responsible for the security and compliance of any resources they use for hosting our solution.

In this section, we discuss the AnyCompany application default configuration deployed pursuant to this guide, AWS general best practices, and options for securing your solution on AWS.

### IAM Roles

The CloudFormation template creates four roles: CloudStorageSecConsoleRole, CloudStorageSecAgentRole, CloudStorageSecExecutionRole, CloudStorageSecUserPoolRole. These roles allow the Console and Agent to perform their required actions.

Where possible we have followed the best practices of "least privilege", but you will notice some spots that grant with "\*" because of the actions we are taking. For example, listing all your Amazon S3 buckets requires us to see them all (\*), but delete permissions on a bucket are limited only to buckets the product creates (specific down to bucket names). We do not grant Antivirus for Amazon S3 permissions that are not needed.

Our [IAM Permissions Review](#) provides more details if needed. Look to [AWS documentation](#) for [IAM best practices](#).

### Console Access

The console is accessed through a unique URL generated during the deployment process that follows the convention <accountID-appID>.cloudstoragesecapp.com. This URL is tied to the IP assigned to the console task at boot and gets reassigned each time the console requires a reboot. The IP assignment is done via a Lambda execution in the Cloud Storage Security account.

**Note:** If you are using a private deployment, this URL will not be valid. You must access your console via the private IP or the load balancer URL if you added one

Whether you leverage the unique URL ([which you can rename](#)) or the IP, you will have to have access to the network from your locations. If the [AWS Security Groups](#) or the subnet doesn't allow public access, you must ensure you do have access from the network you are running.

**Note:** As indicated you can rename the URL which CSS suggests that you do. Then your AWS account number will not be in the URL.

### Outbound Access for Console and Agents

The console and the agent do not require public IPs, but they do require outbound internet access to get to the [Amazon Elastic Container Registry](#) and all other AWS services listed in the previous section as well as for AV signature updates. Outbound routing can be accomplished

through an [AWS Internet Gateway](#) (**not** private), an [AWS NAT Gateway](#) / [AWS Transit Gateway](#) or through [VPC Endpoints](#) + Proxy. A proxy is required because there are three services leveraged that do not have VPC Endpoints at this time: Amazon Marketplace, Amazon AppConfig and Amazon Cognito. Technically, you can use VPC Endpoints with a NAT, but most customers we've seen go down this road leverage a proxy.

**Note:** If you are leveraging an AWS Internet Gateway, then you **must** have public IPs assigned. AWS IGW will **not** work without public IPs.

### [API Access to the Console and Scanning Agents](#)

You can leverage REST APIs for management tasks as well as scanning your data. For management APIs add “/swagger” to the end of your Console URL to see what is available and how to use them. For data scanning APIs look to the [API Help Docs](#).

For either set of APIs, you are required to have a username and password within the AV Console. For data scanning you will also need to be [assigned the right](#) to scan data. Once the user(s) are setup, the first action will be to get an Auth Token. This token auto-expires after 3600 seconds (1 hour).

Cognito has no native ability to expire user passwords. The admin of the system can forcibly require the user to reset their password. This will allow for the rotation of user passwords along with the short-term nature of the authentication tokens.

### [Encrypting Amazon DynamoDB Data at Rest and In Transit](#)

Antivirus for Amazon S3 leverages DynamoDB for configuration and system information. Data at rest encryption is not turned on by default, but there is no technical reason it cannot be used. Please refer to the [DynamoDB Encryption at Rest documentation](#) to learn more. If you turn encryption on, make sure to provide both the CloudStorageSecConsoleRole and the CloudStorageSecAgentRole access to the KMS or the system will not function. By default, your data is also encrypted [in transit](#) over an HTTPS connection.

### [Encrypting Amazon SNS and SQS](#)

Antivirus for Amazon S3 leverages SNS and SQS for event-based scanning as well as [Proactive Notifications](#). Encryption for the Topics and Queues CSS creates is not turned on by default. Many customers do turn this on. You must ensure that access is granted to the keys leveraged for the encryption. If not properly configured, Amazon S3 will not be able to publish to SNS and the CloudStorageSecConsoleRole and CloudStorageSecAgentRole will not have permission to perform the necessary scanning tasks.

### [Amazon S3 Object Encryption](#)

Object encryption is simple to turn on and generally recommended, but it is not a requirement. If you already have it on or want to enable it, make sure the CloudStorageSecAgentRole has

access granted to it within the KMS. Please review the [Amazon S3 object encryption documentation](#) for more details.

## Amazon EBS Encryption

Large File Scanning is only used when scanning files larger than 195GB in size. When files of this size get scanned, they are processed as a Job by a dynamically created EC2 instance and EBS volume. Both are short lived and terminated after the scan is complete. We use AWS default encryption on the EBS volume for the duration of the scan.

## Logging/Auditing

All logging is done through [AWS CloudWatch Logs](#) from the Console and Agents. The CloudFormation template generates the Console focused logs, while the Console itself creates the logs in the various regions where agents are deployed. For more information on the log groups Antivirus for Amazon S3 create and for sample logs, please refer to the [documentation](#).

## Costs

The deployment of the Antivirus for Amazon S3 product will create and/or touch multiple AWS services as we have seen through this guide so far. The brunt of the costs associated (not including scanning software costs) will be around the running of the container tasks within AWS Fargate. [AWS Fargate pricing](#) comes down to the number of vCPU and number of memory GBs you run in each container each hour. Pricing for ECS Fargate containers is charged by the second (with a 1 minute minimum).

## Infrastructure Components

Deploying Antivirus for Amazon S3 will create the AWS resources listed below:

- Fargate components
  - Cluster
  - 1 Console Service and Task (required)
  - 1 Event Agent Service and Task (optional)
    - Optional by deployment needs
    - Agent tasks are created in each region you protect buckets. So you may end up with 1 or many agent tasks depending on your deployment needs
  - 1 API Agent Service and Task (optional)
    - Optional by deployment needs
    - Agent tasks are created regionally / geographically located based on how close you need the API endpoint to be
- DynamoDB tables
- CloudWatch Logs
  - Console Log groups

- Only deployed in the console region
  - Agent Log groups
    - These Log groups will be created in each region that deploys an agent
- SNS Topics
  - 1 Notifications topic only in the console region
  - 1 S3 Events topic will be created in each region that deploys a scanning agent
- SQS
  - 1 per scanning agent region
- AWS API calls

### Example Infrastructure Costs

The following assumes you run the containers full time during the month. AWS calculates this as 730 hours.

Default Management Console Configuration (.5vCPU, 1GB RAM):

$$(.5 * \$0.04048) + (1 * \$0.004445) = \$0.024685 / \text{hr} * 730\text{hr} = \mathbf{\$18.02005 / \text{month}}$$

Default Scanning Agent Configuration (1vCPU, 3GB RAM):

$$(1 * \$0.04048) + (3 * \$0.004445) = \$0.053815 / \text{hr} * 730\text{hr} = \mathbf{\$39.28495 / \text{month}}$$

### AWS API Calls – the other infrastructure costs

The AWS APIs we use for data scanning of any substantial cost are all around Amazon S3: LIST, PUT, GET, and S3 Tags (these count as a PUT as well as have their own cost). The following table shows the costs for the different calls CSS makes at various volumes.

S3 API Costs			
Number of Objects	List / Put Calls (\$0.005/1,000 calls)	Get Calls (\$0.004/1,000 calls)	S3 Tag Storage (\$0.01/10,000 tags)
1 million	\$5	\$0.40	\$2
10 million	\$50	\$4	\$20
100 million	\$500	\$40	\$200
1 billion	\$5,000	\$400	\$2,000

## Additional Cost Considerations

- Console generally needs to run all the time
- Agents do not have to run all the time, this depends on your workload and requirement for how quickly an object is scanned. Check out the [Smart Scan](#) efficiency mode available for infrastructure cost optimizations.
- Agents are required in each AWS region you protect buckets
  - So the above single-agent cost will be per region depending on your settings
  - Smart Scan can be mixed amongst regions. Regions with lower volume could be setup to run only when work is available and be shut down at all other times
- You may run more than 1 agent per region depending on volume and scan window requirements
  - Auto-scaling settings can be controlled on the [Agent Settings](#) page within the console. This allows you to control how much, if at all, you will scale up so you know the caps
- The other services leveraged will generate a negligible set of costs associated with them to run the solution (may add a few dollars)
  - DynamoDB, Systems Manager, SNS, SQS, CloudWatch
- **NOTE:** Storage and data transfer are not included as these vary depending on configuration. Please consult [AWS Pricing](#) for the latest information.

For a walk-through of the Total Cost of Ownership, please review the [TCO Overview Video](#).

## Sizing

Sizing and scalability for Antivirus for Amazon S3 are dependent on a few factors within your environment. You should ask yourself the following questions and then map into the scanning results data we've posted in our [Sizing discussion](#) within our Help Docs. This is at least a good starting point to walk yourself through the initial sizing, but final deployments should be tuned as needed because your mileage may vary.

### Questions to Ask

- Load
  - How many objects do I currently have?
  - How many objects flow into the system?
  - What size are the objects coming into the system?
  - Do they trickle in throughout the month or gush in at certain times?
- Time to Scan
  - Do you have a particular scan window?
  - Is there a requirement that files must be scanned within X time (seconds/minutes/hours) of receiving?

### Sizing Example

Let's walk through an example to better understand.

Assumptions:

- Receive 5 million files per month
- We'll look at 2 file sizes: 100 KB and 10 MB
  - Other file sizes can be sorted through in the documentation link listed above
- We'll assume all files come during workdays and 8am-5pm
- Sophos engine in use

For 100kb size files with a modest scanner size, we are seeing consistently ~22,000 files per hour. That is ~366 files per minute or a just over 6 files per second. For 10MB file sizes we are seeing ~10,000 per hour or ~167 per minute or ~2.7 every per second.

The scanners scale linearly. Each additional scanner added will add a similar number of files processed. Two scanners will process ~44,000 100KB files or 20,000 10MB files per hour. There is no limit to this beyond the number of objects in flight coming out of the SQS queues and your service quota limits for Fargate. One other limiter we have seen is the CIDR of the Subnets the solution is run in. A /24 will allow for 254 tasks to run. Two subnets will get you ~500 tasks running. If you need more than that, look at your service quotas and your CIDR ranges.

Using the assumptions from above, 5 million files a month broken down into 22 weekdays and 8 working hours in each gives us 176 hours where files will be deposited. That breaks down to 28,409 files per hour or 473 files per minute or 8 files per second.

If all of those were 100kb files, you could run 2 scanners (~44k file processing) and easily keep up with the load per hour. But, if you say what you really want is that every file is never any older than say 5 seconds before it is scanned, that is a different perspective. Keeping with current file size, 2 scanning agents would do ~12 files per second. So you could actually still keep up if there is a steady trickle of the files evenly dispersed across that hour. But, if the files were 10MB in size and we only do 3 files per second and you wanted to still stay on top a file never being older than 5 seconds, you would therefore run closer to 3-4 scanners. This way you are still keeping up with your numbers and the files are never more than 5 seconds old.

So it comes down to the math for processing and the age of the file before it is scanned. Adjust for how your environment works and the allowed age of files. If you get files all throughout the day and all 30 days of the month, then the load the scanners need to handle is even lighter (due to more processing hours available) and so fewer scanning agents could be run. And of course, with auto scaling in place, you don't always have to run a full load of scanning agents, you can have them scale up and down.

Please [contact us](#) if you have specific questions or concerns around sizing as we are happy to help you design and configure the solution.

## Testing and Deployment

The Antivirus for Amazon S3 solution has been created to scan your Amazon S3 objects for malware and viruses. Testing the product is straight forward. [Turn on protection](#) for a test

bucket and upload sample files to it. That's it. The files should be examples of what your environment sees regularly as well as known virus files so you can ensure the scanners are behaving as expected. [Eicar.org](http://Eicar.org) produces benign files that can be tested with. There are many other repositories out there where more harmful files can be downloaded. Please practice safe handling if you choose to use stronger files.

## Health Check

There are a number of ways to check the health of the system from logging into the application console to checking services and logs within the AWS Management Console.

The main questions to ask are:

- Am I up and running?
- Am I scanning objects?

Secondarily, you can also ask:

- Am I keeping up with the workload coming in?

## Verify Up and Running

1. From the application console
  - a. Login to the application console
    - i. If you can get in, your console is running and accessible
    - ii. If you can't, skip to 2. From the AWS Console
  - b. Review activity on the Dashboard page
    - i. Review the charts – do you see activity in the time slices (1hr, 24hr, 7 day, etc) you expect to. If so, the agents are up and running
  - c. Navigate to Configuration  Deployment Overview
    - i. This page indicates what aspects of the product you have deployed (Console, Event and API Agents) as well as where you have them deployed
    - ii. The region card details will indicate whether agents are running and indicate whether they should be running based on configuration
2. From the AWS Console
  - a. Login to the AWS Management Console
  - b. In the region you deployed the stack, go to Services  Containers  Elastic Container Service
  - c. Verify a CloudStorageSecCluster exists
    - i. At this dashboard, verify there is at least 1 service installed and at least 1 running task
      1. If there are 0 running tasks, then not even a console is running at this time
  - d. Click into the CloudStorageSecCluster
    - i. Verify the CloudStorageSecConsoleService exists and is running

1. Running Tasks column should contain '1'
- ii. You may or may not have a CloudStorageSecAgentService here based on whether you are protecting buckets within this region
  1. If you do have the CloudStorageSecAgentService, the running tasks may still be '0' based on your configuration settings around scaling
- e. Go to Services  Management & Governance  CloudWatch  Log Groups  CloudStorageSecurity.Console.Metering
  - i. In the latest Log Stream verify there are Log Events
    1. If the events are hourly in nature, the application console is running properly and submitting to AWS

## Verify Scanning

1. From the application console
  - a. Login to the application console
  - b. Review activity on the Dashboard page
    - i. Review the charts – do you see activity in the time slices (1hr, 24hr, 7 day, etc) you expect to. If so, the agents are up and running
    - ii. Instead of just looking for activity, do more specific sanity checks like verifying the number of new objects received in the last 1hr (or other time slice) and compare it against what is reflected in the charts
  - c. Review activity on the Scan Results page
    - i. Navigate to Monitoring  Scan Results
    - ii. Review the daily activity
      1. Is it the number of files you expect?
      2. Are they all accounted for?
    - iii. Break the activity review down by Account
      1. Expected volumes per account?
    - iv. Break the activity review down by Bucket
      1. Expected volumes per bucket?
  - d. Verify “problem files” are reflected on the Problem Files page for the time slice desired
  - e. Navigate to Configuration  Agent Settings to verify auto-scaling configuration
    - i. Verify scaling configurations to ensure agents are running as expected
    - ii. You can also determine here, which agents (Event or API) are deployed and in which regions
2. From the application notifications
  - a. Antivirus for Amazon S3 creates an SNS Topic (CloudStorageSecNotificationsTopic) you can subscribe to in order to receive [Proactive Notifications](#)
  - b. If you are receiving notifications for the scan results, the agent scanners are operating as expected
3. From the AWS Console
  - a. Login to the AWS Management Console

- b. In the regions you have protected buckets (this may be more than just the console region), go to Services  Containers  Elastic Container Service
- c. Verify a CloudStorageSecCluster exists
  - i. At this dashboard, verify there is at least 1 service installed
    - 1. You may or may not see running tasks and that can be ok
      - a. Agents can run all the time or be controlled, by auto-scaling groups which are setup and modified within the application console, to [run only when work is available](#)
- d. Click into the CloudStorageSecCluster
  - i. Verify the CloudStorageSecAgentService exists
    - 1. Based on the Agent Settings for the given region you are in, verify the agents are running as configured
      - a. For example, if [Agent Settings](#) are configured to have Min Agents set to '3', ensure there are at least 3 running agents
- e. Go to Services  Management & Governance  CloudWatch  Log Groups  CloudStorageSecurity.Agent.ScanResults
  - i. In the latest Log Stream(s) verify there are Log Events
    - 1. There are 4 Streams created: Clean, Infected, Error and Unscannable
    - 2. Review log streams to verify activity
    - 3. The CloudWatch Logs are created in each region with the agent, so this step can/should be done in each region

### Keeping Up with the Work

You must understand your data in order to know if the system is keeping up. Can you determine how many objects are coming in to compare against how many objects are being processed by the system. Whether you have a clear grasp on this side of your data, there are two places you can look to get a feel of the throughput.

1. From the application console
  - a. Login to the application console
  - b. Review activity on the Dashboard page
    - i. Review the 'Total Objects Scanned' information widget by clicking on the "zoom" icon
      1. This gives you a total count by 1 Hour, 24 Hours and longer. Seeing the 1 Hour and 24 Hour will give you a feel for the throughput in a given day
      2. Does this match up to expectation
  - c. Review activity on the Scan Results page
    - i. Navigate to Monitoring  Scan Results
    - ii. Review the daily activity
      1. Is it the number of files you expect?
      2. Are they all accounted for?
    - iii. Break the activity review down by Account

1. Expected volumes per account?
  - iv. Break the activity review down by Bucket
    1. Expected volumes per bucket?
2. From the AWS Console
  - a. Login to the AWS Management Console
  - b. In the regions you have protected buckets (this may be more than just the console region), go to Services  Application Integration  Simple Queue Service
  - c. Review the CloudStorageSecQueue items
    - i. Depending on your object volume and flow, the Messages Available and Messages In Flight entries should approach 0 or at least be decreasing over time
    - ii. If you find the number always increasing, it could be an indication that Antivirus for Amazon S3 is not keeping up
    - iii. Modify the [Agent Settings](#) scaling options, if so desired, to scale the application further and it will process a greater number of objects and keep up with your data volumes

## Backup and Recovery

### Backup

The “brains” of the deployment are the DynamoDB tables. All other aspects of the solution are expendable and better to replace than restore. Much of the configuration and all of the historical data is kept within DynamoDB. DynamoDB can be backed up leveraging [on-demand backup](#) or [point-in-time recovery](#). Follow AWS best practices and internal requirements to back these tables up.

Should you need to restore from these table backups, make sure the Console task and all Agent tasks are shutdown beforehand. Follow the AWS Documentation for [how to restore a DynamoDB backup](#). If you have been using Point in Time backups, follow this AWS Documentation for [how to restore a DynamoDB backup](#).

Currently, [Amazon Cognito](#) does not have a built-in way to back up your users. There are scripts that can be run to capture most of the data via export from Cognito which can then be imported back in if needed. There is no way to capture the password hashes, so all users will have to “reset password” in the event the User Pool has to be recreated.

### Instance Failure

The AWS Fargate container tasks are managed by policy and will recover from failures. The Console task is set to always have one running, so if it shuts down for any reason it will be started back up. The Agent task is set based on scaling policies that leverage CloudWatch Alarms. The agents will scale up and contract back down based on these rules.

## Availability Zone Failure

It is recommended that you provide at least two subnets to both the Console and the Agents and make it so the subnets are in different Availability Zones. In the event an AZ goes down, the console or agent will spin back up in a different AZ typically within two minutes.

## Region Failure

In the event of region failure, only the components running in that region would face issues. If it were the console region, the console would not be up and running to manage the solution. With that said, the scanning agents in all other regions would continue to operate and scan data. Antivirus for Amazon S3 is designed to sustain durations of disconnected function between the scanning agents and the console. You can simply wait for the region to come back online and start the console backup (auto-scaling should do this automatically) and the system will normalize itself and continue to operate.

If the region appears it will be down longer than what is acceptable, you could leverage your DynamoDB backups and the CloudFormation template to recreate the architecture in another region. The time to restore varies on the volume of your system processing to date as that impacts the DynamoDB table sizes. A <4-hour recovery time objective (RTO) and a <24 hour recovery point objective (RPO) are certainly achievable. Unless critically needed, your easiest solution is to wait for the region to recover as the system itself is self-healing and will come back up.

If absolutely necessary, to restore operations in another region:

- Run AWS CloudFormation template to recreate console deployment
- Ensure Console task is not running
- Restore DynamoDB tables, with the exception of the Console table, over the top of the new installation tables
  - Email us the row from your backed up Console table. We need to create a new row (decrypt and re-encrypt) the data for you to place back into your new table
- Start console back up
- Remove “All create events” event from each of your protected buckets
- Shut down current running Event and API Agents across all regions
- Re-enable bucket protection for all required buckets
- Run a [Scan Existing](#) scan for the window of time between the original console region going offline and the re-enabled bucket protection
- You can delete the original clusters for the agents

## General Failure Considerations

For a stateless application, most of your business critical data will be housed in Amazon DynamoDB. We highly recommend you do daily back-ups to ensure minimal loss and downtime for your application. Work performed is driven by events and queue entries. If the agents go

offline, whether intentionally or unintentionally, the work queues will continue to build. Once the agents are back online they will start chugging through the entries.

## Routine Maintenance

Antivirus for Amazon S3 is a “black box” solution from the container perspective. There won’t be times you have to login directly to those tasks to maintain and update them. As the underlying components require updating, Cloud Storage Security will update the ECR image with the latest components. We build with “no cache” whenever we update our product images.

Product updates are rolled out regularly. Make sure to perform the [updates through the application console](#) for ease of use. AV signature updates are done every one hours for the ClamAV engine and every 15 minutes for the Sophos engine. Both sets of updates are done automatically. Look to subscribe to [Proactive Notifications](#) to be made aware when new releases are pushed out.

Deploying the software requires the use of multiple AWS Services. In some cases (VPC for example) you can reuse previously created entities, but you may also choose to create new entities specifically for the solution. Please evaluate your [AWS Service Quotas](#) in each region you plan to deploy a component of the product. You can also go to the [Services Quotas](#) centralized view within the AWS Console to see what your current limits are at and make changes. It is best practice to do this evaluation regularly.

## Emergency Maintenance

Please [contact us](#) if you are experiencing an emergency with the product. We are happy to help you. In the event you need to replace the existing deployment with a new deployment and want to keep your historical data, proceed with the following steps:

- Run AWS CloudFormation template to recreate console deployment
- Ensure Console task is not running
- Restore DynamoDB tables, with the exception of the Console table, over the top of the new installation tables
  - Email us the row from your backed up Console table. We need to create a new row (decrypt and re-encrypt) the data for you to place back into your new table
- Start console back up
- Remove “All create events” event from each of your protected buckets
- Shut down current running Event and API Agents across all regions
- Re-enable bucket protection for all required buckets

- Run a [Scan Existing](#) scan for the window of time between the original console region going offline and the re-enabled bucket protection
- You can delete the original clusters for the agents

## Support

	Standard Plan	Premium Plan
<b>Live Support Coverage</b>		
Email & Phone	9-5 ET, M-F	8-8 ET, M-F
Dedicated Slack Channel		8-8 ET, Daily
<b>Support Response Times</b>		
General Guidance	< 24 hours, M-F	< 24 hours, M-F
System Impaired	< 12 hours, M-F	< 4 hours, M-F; < 12 hours, Weekends
<b>Training &amp; Updates</b>		
Initial Training (5 Named Users)		3, 90 Minute Sessions
Ongoing Training (5 Named Users)		45 Minutes, Monthly
Product & Roadmap Reviews		Quarterly
Recorded Sessions	On Demand	On Demand
	<b>Included</b> with subscription	<b>Contact us for Pricing</b>

## Troubleshooting

Please refer to our [Troubleshooting](#) section for details.

## FAQ

Please refer to our [FAQ](#) section for details.

## Clean Up

Antivirus for Amazon S3 deploys some infrastructure through the CloudFormation template and some infrastructure through code within the console. As a result, doing a [CloudFormation delete stack](#) is not enough to completely clean up the solution. As a result, Antivirus for Amazon S3 provides a cleanup mechanism within the console that will allow you to “clean” particular regions or the complete solution. The Console will take the action to delete everything it has created and then indicate to you when to go initiate a stack deletion on the CloudFormation stack in the AWS Management Console to complete the process.

The application cleanup process will not remove any backups you have performed as the application will be completely unaware of them.

## Contact Us

We would love to hear from you. If you would like to provide feedback or need any help with the product, please email [support@cloudstoragesec.com](mailto:support@cloudstoragesec.com) or call +1 801-410-0408. Support is included with all AWS Marketplace offerings at no additional cost. We will respond within 24 hours (usually much faster).

If you are a software vendor and would like to integrate and partner with us, please reach out to [partners@cloudstoragesec.com](mailto:partners@cloudstoragesec.com).