AT&T Cybersecurity

AT&T Business

USM Anywhere™

Deployment Guide

Copyright © 2024 AT&T Intellectual Property. All rights reserved.

AT&T, the AT&T logo and all other AT&T marks contained herein are trademarks of AT&T Intellectual Property and/or affiliated companies. All other marks are the property of their respective owners.

Updated April 03, 2024

Contents

About USM Anywhere Deployment 6

USM Anywhere Deployment Guide 7

USM Anywhere Architecture 7

USM Anywhere Data Security 10

USM Anywhere Log Data Enhancement 12

USM Anywhere Deployment Types and Scalability 17

USM Anywhere Deployment Requirements 18

USM Anywhere Deployment Process 19

USM Anywhere Updates 21

VMware Sensor Deployment 24

About VMware Sensor Deployment 25

Requirements for VMware Sensor Deployment 26

Create the VMware Virtual Machine 36

The OVF Package Is Invalid and Cannot Be Deployed - SHA256 Error 42

Set Up USM Anywhere on the VMware Virtual Machine 43

Connect the VMware Sensor to USM Anywhere 46

Complete the VMware Sensor Setup 51

USM Anywhere Sensor Deployment on Microsoft Hyper-V 69

About Hyper-V Sensor Deployment 70

Requirements for Hyper-V Sensor Deployment 70

Create the Hyper-V Virtual Machine 80

Set Up USM Anywhere on the Hyper-V Virtual Machine 89

Connect the Hyper-V Sensor to USM Anywhere 92

Complete the Hyper-V Sensor Setup 97

USM Anywhere Sensor Deployment on AWS 114

About AWS Sensor Deployment 115

Requirements for AWS Sensor Deployment 116
Deploy the AWS Sensor 129
Connect the AWS Sensor to USM Anywhere 133
Complete the AWS Sensor Setup 137
Enable Connections in an AWS VPC 148
VPC Traffic Mirroring with an AWS Sensor 150
Collect Logs from Amazon S3 Buckets with KMS Encryption 158
AWS Log Discovery and Collection in USM Anywhere 159

USM Anywhere Sensor Deployment on Microsoft Azure 182
About Azure Sensor Deployment 183
Requirements for Azure Sensor Deployment 184
Deploy the USM Anywhere Sensor from the Azure Marketplace 194
Create an Application and Obtain Azure Credentials 197
Connect the Azure Sensor to USM Anywhere 203
Complete the Azure Sensor Setup 207
Azure Log Discovery and Collection in USM Anywhere 225

USM Anywhere Sensor Deployment on GCP 249
About GCP Sensor Deployment 250
Requirements for GCP Sensor Deployment 251
Preparing Your GCP Environment for Sensor Deployment 261
Deploy the GCP Sensor 267
Connect the GCP Sensor to USM Anywhere 269
Complete the GCP Sensor Setup 274
GCP Log Discovery and Collection in USM Anywhere 287

The AWS Cloud Connector Deployment in USM Anywhere 300
Differences Between an AWS Cloud Connector and a Sensor 300
Activating an AWS Cloud Connector 302
Uploading AWS CloudFormation Templates 302
AWS Cloud Connector Resources 306

Network Setup and Configuration 308
Configure Network Interfaces for On-Premises Sensors 309
Configure USM Anywhere to Receive ERSPAN Traffic 320
Port Mirroring Configuration on Network Devices 321
Granting Access to Active Directory for USM Anywhere 332
Proxy Configuration on the USM Anywhere Sensor 334
Data Sources and Log Processing 341
Data Sources and Log Collection 342
File Integrity Monitoring 394
Scheduling Active Directory Scans from the Job Scheduler Page 401
Alarm and Event Notifications 407
Sending USM Anywhere Notifications to Slack 407
Sending USM Anywhere Notifications to Datadog 420
Sending USM Anywhere Notifications to PagerDuty 429
Sending Notifications Through Amazon SNS 438
Troubleshooting and Remote Sensor Support 450
Checking Connectivity to the Remote Server 451
Creating a Remote Support Session 453
Sensor System Menu 456
Collecting Debug Information 466
View Network Testing Information 468
Retrieve Unique Identifier Information 476

About USM Anywhere Deployment

USM Anywhere is a software as a service (SaaS) security monitoring solution that centralizes threat detection, incident response, and compliance management across your on-premises, cloud, or hybrid environments. Data collection, security analysis, and threat detection are centralized in the AT&T Cybersecurity Secure Cloud and provide you with a single view into all of your critical infrastructure.

This section includes the following topics:

USM Anywhere Deployment Guide 7

USM Anywhere Architecture 7

USM Anywhere Data Security 10

USM Anywhere Log Data Enhancement 12

USM Anywhere Deployment Types and Scalability 17

USM Anywhere Deployment Requirements 18

USM Anywhere Deployment Process 19

USM Anywhere Updates 21

USM Anywhere Deployment Guide

USM Anywhere consists of a modular, scalable, two-tier architecture to manage and monitor every aspect of cloud security. Software sensors collect and normalize data from all of your on-premises and cloud environments, while USM Anywhere provides centralized cloud security management, analysis, correlation, detection, alerting, log management, and reporting.

Purpose-built USM Anywhere Sensors deploy natively into each environment and help you gain visibility into all of your on-premises and cloud environments. These sensors collect and normalize logs, monitor networks, and collect information about the environments and assets deployed in your hybrid environments.

USM Anywhere Architecture

USM Anywhere has a modular and scalable two-tier architecture.

USM Anywhere Architecture Diagram

Tier 1 — USM Anywhere Sensors and Agents

USM Anywhere Sensors deploy natively into each environment and help you gain visibility into all of your on-premises and cloud environments. Sensors collect and normalize logs, monitor networks, and collect information about the environments and assets deployed in your hybrid environments.

Sensors are a key component of the USM Anywhere solution. They operate either on-premises or in the cloud, performing the following tasks:

- Discovering your assets
- Scanning assets for vulnerabilities
- Monitoring packets on your networks and collecting data
- Collecting log data and normalizing it before securely sending it to USM Anywhere

USM Anywhere Agents deploy on your network host and provide the following:

- Endpoint detection and response
- Network asset monitoring
- File integrity monitoring (FIM)
- Log collection

Tier 2 — USM Anywhere Cloud

The USM Anywhere cloud instance is deployed in one of the Amazon Web Services (AWS) endpoint regions based on your location. The following table lists the available AWS regions.

AWS Regions where USM Anywhere Instance Is Available

Code	Name
ap-northeast-1	Asia Pacific (Tokyo)
ap-south-1	Asia Pacific (Mumbai)
ap-southeast-1	Asia Pacific (Singapore)
ap-southeast-2	Asia Pacific (Sydney)
ca-central-1	Canada (Central)
eu-central-1	Europe (Frankfurt)
eu-west-1	Europe (Ireland)
eu-west-2	Europe (London)
me-central-1	Middle East (UAE)
sa-east-1	South America (São Paulo)
us-east-1	US East (N. Virginia)
us-west-2	US West (Oregon)
us-gov-west-1	AWS GovCloud (US-West)

USM Anywhere receives data from USM Anywhere Sensors and uses it to provide essential security capabilities in a single SaaS platform:

- Centralized system security management
- Log data analysis and correlation
- Detection
- Alerting
- Log management
- Reporting

USM Anywhere also retains raw logs long-term for forensic investigations and compliance mandates.

USM Anywhere Data Security

As a security-first organization, AT&T Cybersecurity makes your data protection and privacy a top priority. USM Anywhere architecture and processes are designed to protect your data in transit and at rest.

Data Collection

All data sent from the USM Anywhere Sensor deployed in your on-premises or cloud environment to the USM Anywhere service in the AT&T Cybersecurity Secure Cloud is encrypted and transferred over a secure TLS 1.2 connection. Each sensor generates a certificate to communicate with the USM Anywhere service. This means that all communication is uniquely encrypted between each sensor and USM Anywhere.

All forensic data (raw logs) is backed up on an hourly basis. The data collected in USM Anywhere is secured using AES-256 encryption for both hot (online) storage and cold (offline) storage.

Data Access

Your data in USM Anywhere is treated as highly confidential, and only a select few AT&T Cybersecurity staff members have access. This group of employees uses multi-factor authentication (MFA) to access the AT&T Cybersecurity Secure Cloud. Strict internal controls and automation enable support for the service while minimizing administrative access.

AT&T Cybersecurity also has a formal information security program that implements various security controls to the National Institute of Standards Technology (NIST) Cyber Security

Framework. Key controls include: Inventory of Devices, Inventory of Software, Secure Configurations, Vulnerability Assessment, and Controlled Use of Administrative Privileges. Additionally, AT&T Cybersecurity conducts security self-assessments on a regular basis.

Cold Storage Data Integrity

USM Anywhere offers secure long-term log retention, known as *cold storage*. By default, USM Anywhere stores all data associated with a customer's subdomain in cold storage for the life of the active USM Anywhere subscription at no additional charge, while AT&T TDR for Gov customer data are kept for three years or longer (if requested).

Important: The retention period set on the license (30-days standard or 90-days standard) only applies to regular events. The retention policy for system events is 30 days and for user activities is 180 days, while the user activities related to investigations never expire.

USM Anywhere uses a *write once, read many* (WORM) approach in log storage to prevent log data from being modified or otherwise tampered with. You can download your raw logs at any time. If you do not renew your subscription, AT&T Cybersecurity will keep the raw logs for 14 days after your subscription expires, giving you a grace period to restart your service. Within the 14 days, no data is collected until your license is reactivated. Therefore, data is lost between license expiration and reactivation. After 14 days, your data will be destroyed.

End-of-Contract Shut Down

If your subscription expires and you decide not to renew, your USM Anywhere instance will be decommissioned 14 days after the expiration. All data, including asset information, orchestration rules, user credentials, events and vulnerabilities (hot storage), and raw logs (cold storage), will be destroyed.

Business Continuity Plan

To ensure business continuity, USM Anywhere executes a backup procedure 2 times a day, encrypts the data, and stores it for 15 days. The Recovery Point Objective (RPO) is up to 12 hours and the Recovery Time Objective (RTO) is approximately an hour, depending on the size of the data being restored.

Password Policy

USM Anywhere stores and encrypts user credentials using the latest industry standards for securing passwords.

Keep in mind these points when you are logging in:

- The login credentials that you set will apply to any USM Anywhere™ and USM Central™ you have access to.
- USM Anywhere requires all passwords to have a minimum length of 8 characters and a maximum length of 128 characters.
- The password must contain numerical digits (0-9).
- The password must contain uppercase letters (A-Z).
- The password must contain lowercase letters (a-z).
- The password must contain special characters, such as hyphen (-) and underscore (_).

Note: USM Anywhere passwords expire after 90 days. When your password expires, USM Anywhere enforces a password change when you next log in. A new password must be different from the previous four passwords.

After 45 days of inactivity, your user account will be locked. Manager users can unlock inactive accounts.

A user account is locked for 30 minutes after 5 consecutive failed login attempts (GovCloud users are locked out after 3 consecutive failed login attempts).

USM Anywhere Log Data Enhancement

When evaluating threats to your systems, the more complete and clear the context of an incident is, the more accurate and efficient USM Anywhere can be in identifying and responding to those threats. Log data is one of the key sources of this threat data context, providing a tremendous amount of information about network events. Every network connection, authentication request, file transfer, and privilege escalation generates a log message.

However, many of these log messages were not originally designed to be used for security purposes. There are no official standards for log contents (although there are best practices); therefore, log message content is often inconsistent and incomplete.

For example, look at a typical log message generated by an authentication event:

```
{
       "outcome" : "Allow",
       "type" : "Authentication",
       "source" : "13.107.4.50",
       "destination" : "10.60.5.94",
       “time” : “2018-10-17T19:03:26+00:00”
}
```

This message is brief and doesn't provide enough context for incident analysis. USM Anywhere can improve that context by normalizing and enriching the data provided in the log message.

Data Normalization

The first step USM Anywhere takes when it analyzes your system logs is to normalize them so that all incoming data uses the same terminology. In this context, normalization means mapping it to a standard terminology. For example, a vendor may use the terms "outcome" or “result” to describe the success or failure of the authentication attempt. USM Anywhere normalizes these two different attributes, replacing them with a single, standard term. Likewise, things like source, source_ip, client, and client_ip all need to be mapped to the same set of terminology so events from different vendors can be used for correlation and alarm generation.

The following is an example of how normalization works. Note that USM Anywhere preserves the original log message as a best practice in case you need to share it with a vendor or need to refer to the original alert. This means that the normalization phase of message processing likely increases the size of the log message by around 100%.

```
{
       "log" : "{ \"outcome\" : \"Allow\",
                  \"type\" : \"Authentication\",
                   \"source\" : \"13.107.4.50\",
                  \"destination\" : \"10.60.5.94\" }",
       "source_address" : "13.107.4.50",
       "destination_address" : "10.60.5.94",
       “event_outcome” : “ALLOW”,
       "event_name" : "Authentication",
       "timestamp_occured" : “2018-10-17T19:03:26+00:00”
}
```

Data Enrichment

Normalization enables you to analyze all the log messages USM Anywhere receives. Given the incomplete nature of so many log messages, it also makes sense to use this same process to add valuable information to the log messages, which helps USM Anywhere perform better incident detection.

Data enrichment is the process by which valuable information is added to log messages. The USM Anywhere infrastructure has a large amount of contextual data about the network and systems that it can attach to the log messages to fill in the gaps and enhance threat detection. It also has access to many databases of things like the location of specific IP addresses, device types, and threats it can also leverage.

These are examples of information that can be added through data enrichment:

- Device identity
- Geolocation
- Collection details and flags

Device Identity

Most servers rely on Dynamic Host Configuration Protocol (DHCP) for dynamic IP address allocation. From a security point of view, this means that identifying and containing threats is much more difficult. By the time a system is identified as compromised, it may be on the network in a completely different place with a completely different IP address. To address that problem, USM Anywhere uses the network context it has to collect and includes the media access control (MAC) address, fully qualified domain name (FQDN), and a unique identifier for the system, depending on which are known:

```
"source_asset_id" : "f8ebb373-b551-43d0-a628-a00771b5d0c1",
"source_mac" : "98:01:A7:B4:D8:47",
"destination_fqdn": "ip-10-6-255-129.ec2.internal",
"source_fqdn": "ip-10-6-2-102.ec2.internal",
```

Geolocation

Knowing where your network connections are terminating is important when deciding if traffic should be permitted, blocked, or more carefully monitored. Geolocation can play a role in deciding if a given incident is worthy of more attention. USM Anywhere augments logs with geolocation information of source and destination. In the following example, this data enables an operator to quickly determine that this particular destination is probably not an issue:

```
"destination_address" : "10.60.5.94",
"destination_name" : "AD Server",
"destination_asset_id" : "8cdf98a1-533d-9ec2-b5bc-3424caecef15",
"destination_organisation : "Microsoft Azure",
"destination_city" : "Redmond",
"destination_fqdn" : "ad.alienvault.com",
"destination_hostname" : "ad",
"destination_organisation : "Microsoft Azure",
"destination_latitude" : "47.6801",
"destination_longitude" : "-122.1206",
"destination_region" : "WA",
"destination_country" : "US",
"destination_country_registered" : "US",
```

Collection Details and Flags

USM Anywhere also includes some additional information about how the log message was acquired and processed. This information is included to give the security analyst and correlation algorithms insight into the source of the log, when a sensor received it, and how it was processed. For example, `was_fuzzied = true` means that the log message was received from a source that USM Anywhere doesn't have a specialized plugin for and, therefore, it may not have normalized all the fields. If the log is key to an investigation, the operator should look at the original log message and ensure nothing was overlooked.

Impact on Log Storage

Because USM Anywhere adds data to log messages, the size of the original log message inevitably grows. Very sparse messages can grow as much as 1860%. However, the messages themselves are still relatively small, typically growing from less than 250 B to as much as 2.6 KB, adding up over time. The good news is that the amount of metadata added is stable, which means it doesn't grow much larger or shrink in size for different event classes. So with careful planning, storage use can still be quite predictable. For larger events (for example, events coming from network-based intrusion detection systems [NIDS] and Amazon Web Services [AWS]), the percentage goes down significantly since the messages start out quite large. However, for small events such as the one in the previous example, it can have a noticeable impact on the total amount of data stored.

These are some syslog- and AWS-heavy data points for planning purposes.

Syslog-heavy deployment

From a sample size of 599,979 events

- **Total size including enriched data in bytes:** 1,612,790,164
- **Total size of just log data in bytes:** 145,781,057
- **Average log size in bytes:** 243
- **Average log size with enriched data:** 2,688
- **Increase in size:** 1106%

AWS-heavy deployment

From a sample size of 500,000 events

- **Total size including enriched data in bytes:** 1,934,740,282
- **Total size of just log data in bytes:** 711,502,141
- **Average log size in bytes:** 1,423
- **Average log size with enriched data:** 3,868
- **Increase in size:** 272%

What Happens When You Reach the Tier Limit?

If you find yourself running into problems with inadequate storage space, your first step should be to review your logging strategy with AT&T Cybersecurity Technical Support or your service provider. It may be that you don't need to send as many logs as you are. However, it's better to err on the side of logging too much rather than logging too little, since lost logs can't be recovered and security investigations can lead in unexpected directions.

Important: Tier options do not have unlimited processing power, memory allotment, or disk input/output (I/O) speeds. In addition to storage per month, your deployment size's impact on any of these factors will influence which tier option is right for your environment. AT&T Cybersecurity recommends pre-deployment sizing discussions with your sales representative to help select the right tier for you.

AT&T Cybersecurity strives to guarantee that no data is lost, even when you're facing inadequate storage space or processing power. Because of this, USM Anywhere always makes data storage a top priority. When you exceed your data tier, or are projected to far exceed your tier, your system tries to store as much data as possible, even if functionality must be reduced to preserve the data. For instance, if you find that you are over your data tier, you may find that your USM Anywhere has transitioned into one of four possible data consumption tiers. In these tiers, your USM Anywhere may experience some small limitations to its functionality, such as paused correlation, asset counters, and more. All functionality is restored once your USM Anywhere is no longer experiencing resource limitations.

See Understanding Your Data Consumption Status in the *USM Anywhere User Guide* for more information.

Event Filtering

If you want to be proactive with your data consumption, consider reducing the amount of data stored by using filters. Event filtering enables packets to be dropped before they enter correlation and persistence and consume any of the monthly storage allotment. Filtering enables you to define a set of rules for fields, which, when matched, are dropped. This enables you to easily pick certain types of packets that you don't want to enter the system. When filtering, it's important to realize the impact:

- Filtered events are not stored within cold storage.
- Filtered events are not correlated. Alarms are not generated off filtered events.
- Filtered events are dropped from going into hot storage. You will not see them within your events view.

When using filters, it's important to make sure that you're precisely defining the criteria for events to be dropped. If the filter rule is too broad, there is a chance you may drop packets that you are interested in keeping.

USM Anywhere Deployment Types and Scalability

USM Anywhere scales with your business needs. Using the following deployment types you can add or remove sensors, bring on additional cloud services, and scale central log management as your business needs change.

On-Premises Deployment

USM Anywhere provides VMware ESXi and Microsoft Hyper-V Sensors to support an on-premises (private cloud) deployment. The following table summarizes the capabilities each type of deployment has:

VMware ESXi	Microsoft Hyper-V
• ESXi API asset discovery • ESXi log monitoring and alerting • Network-based intrusion detection system (NIDS) packet inspection • Network asset discovery	• NIDS packet inspection • Network asset discovery

Cloud Deployment

USM Anywhere provides Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) Sensors to support deployment on a public cloud.

AT&T Threat Detection and Response for Government (AT&T TDR for Gov), the Federal Risk and Authorization Management Program (FedRAMP)-authorized version of USM Anywhere, provides AWS, Azure, and GCP Sensors for the corresponding government cloud environment.

If your organization possesses resources in more than one cloud environments, you can deploy multiple sensors to monitor your assets. The following table summarizes the capabilities each type of deployment has:

AWS	Azure	GCP
• AWS API asset discovery • AWS CloudTrail monitoring and alerting • Amazon Simple Storage Service (S3) access log monitoring and alerting • Elastic Load Balancing (ELB) access log monitoring and alerting • AWS infrastructure assessment • NIDS packet inspection	• Azure API asset discovery • Azure Monitor Representational State Transfer (REST) API monitoring and alerting • Azure infrastructure assessment • Azure security alerts • Azure Microsoft Windows log locations	• GCP API asset discovery • Cloud Pub/Sub monitoring and alerting • Audit logs • Stackdriver audit logs

Hybrid Cloud Deployment

A hybrid cloud deployment uses a combination of on-premises sensors (VMware, Hyper-V, or both) and cloud sensors (AWS, Azure, or GCP).

USM Anywhere Deployment Requirements

Before you deploy a USM Anywhere Sensor, you must configure your firewall permissions to enable the required connectivity for the new sensor. Initial deployment of a sensor requires that you open egress or outbound ports and protocols in the firewall for communication with USM Anywhere and AT&T Cybersecurity Secure Cloud resources. The sensor receives no inbound connections from outside the firewall.

Note: To launch the USM Anywhere Sensor web user interface (UI) during the initial setup, you need to allow inbound traffic to the sensor IP address through TCP port 80. You can remove access to this port after the sensor successfully connects to USM Anywhere. You do not need to allow inbound traffic to this port from the Internet.

Each USM Anywhere Sensor has unique requirements. See the following topics for detailed information about these sensor-specific requirements:

- Requirements for AWS Sensor Deployment
- Requirements for Azure Sensor Deployment
- Requirements for GCP Sensor Deployment
- Requirements for Hyper-V Sensor Deployment
- Requirements for VMware Sensor Deployment

Supported Web Browsers

USM Anywhere works best in the latest desktop version of the following web browsers:

- Google Chrome
- Mozilla Firefox

Change the Domain Name

If you want to change the domain name of your environment, you need to contact the AT&T Cybersecurity Technical Support department to open a ticket and indicate the current name and the new one.

Warning: Keep in mind that after this change, all logs and configurations of your environment will be lost.

USM Anywhere Deployment Process

The USM Anywhere registration kicks off the deployment process. There are four basic tasks to complete your initial USM Anywhere deployment.

Task 1: Receive the USM Anywhere Sensor Link and Activation Code

After registering for USM Anywhere online, the system displays a page with the following information you will use to deploy your initial USM Anywhere Sensor:

- A link used to access the sensor template
- An authentication code

You also receive an email with the same information in case you want to do the deployment another time.

Task 2: Provision the Initial USM Anywhere Sensor

Use the provided link to access the USM Anywhere Sensor template for your chosen deployment type, create the new sensor virtual machine (VM) within your cloud account or network, connect to the sensor URL, and then provide your authentication code to provision your USM Anywhere instance within the AlienVault Secure Cloud.

After several minutes, the USM Anywhere provisioning process is complete and you will receive a system message with a URL and password. Access this URL from your browser window and enter your login credentials, including the password you received in the message box.

USM Anywhere prompts you to create a new password for this initial user account. After password verification, the Setup Wizard prompts you to complete the next task.

Task 3: Configure the USM Anywhere Sensor with the Setup Wizard

A Setup Wizard that is specific to your sensor deployment type guides you through the initial configuration of your sensor to initiate the following:

- Initial log collection
- Log management
- Authenticated scans of single assets, an asset group, or a network range

After you create and set up the sensor, it communicates with USM Anywhere in the cloud about the assets in your network. The sensor then transfers any available raw data to USM Anywhere in the cloud for normalization, correlation, and event generation.

Task 4: Configure Your Network for Data Collection

You should configure your network to ensure that the sensor performs optimally and collects the data that you want.

Use the following links to learn about the individual network configuration tasks that may apply to your deployment:

- Collecting Linux System Logs
- Collecting Windows System Logs
- File Integrity Monitoring
- Configure Network Interfaces for On-Premises Sensors
- System Settings for Authenticated Scans
- Granting Access to Active Directory for USM Anywhere
- Direct Traffic from Your Physical Network to the VMware Sensor

Note: Some tasks are specific to the sensor deployment type or the data sources that you have.

USM Anywhere Updates

The USM Anywhere updates occur automatically as new versions become available. The update is transparent, requires no action on your part, and runs in the background. If you want to be notified about the USM Anywhere updates and incidents, subscribe to the USM Anywhere Status page, which provides information like this:

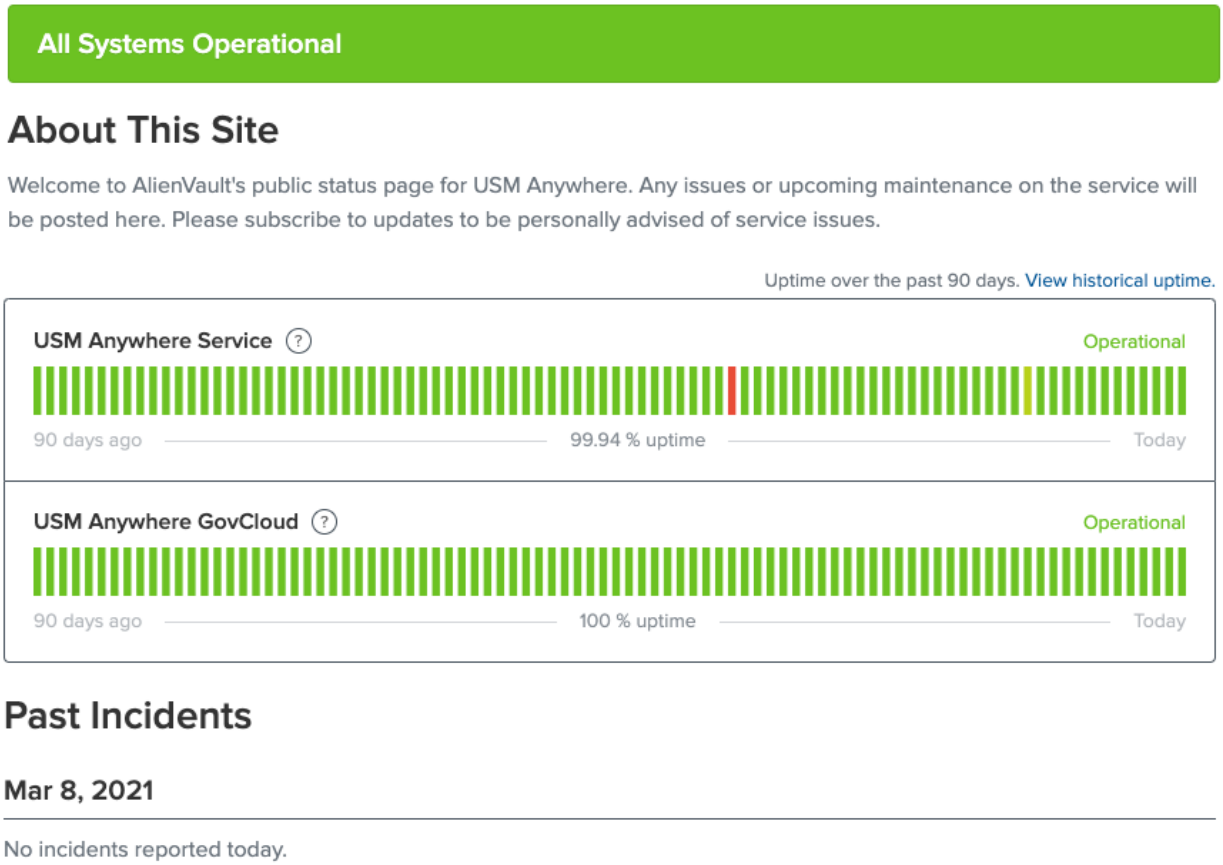

AT&T Cybersecurity recommends that you follow the USM Anywhere Product Announcements to learn what's included in each update and the latest sensor version. When you subscribe to receive updates from USM Anywhere, you will receive two kinds of notifications at the email address used in your subscription:

- **USM Anywhere Provisioning Update**: This is to update the provisioning system (for new customers). Therefore, new deployments are not created until maintenance is complete.
- **USM Anywhere Service Update**: This is to update existing USM Anywhere services. Your service may be offline during this period.

For each update, you receive an email when it is scheduled (with start time and estimated duration), when it is in progress, and when it is completed. The estimated duration is the time it takes for all USM Anywhere services to be updated, also referred to as the maintenance window. Your service restarts after its update has completed.

When your USM Anywhere Sensor tries to connect to the USM Anywhere service after it has completed an update, the sensor downloads and installs the update immediately. Your USM Anywhere Sensor restarts after its update has completed. You do not need to reconfigure the sensor or restart it after an update.

Note: Logs are cached locally during the sensor update process, which will be forwarded after the update has completed and the connection resumes.

In general, your USM Anywhere Sensor can complete its update during the USM Anywhere Service Update maintenance window. You can confirm the sensor version when you log in to your USM Anywhere instance and go to Data Sources > Sensors.

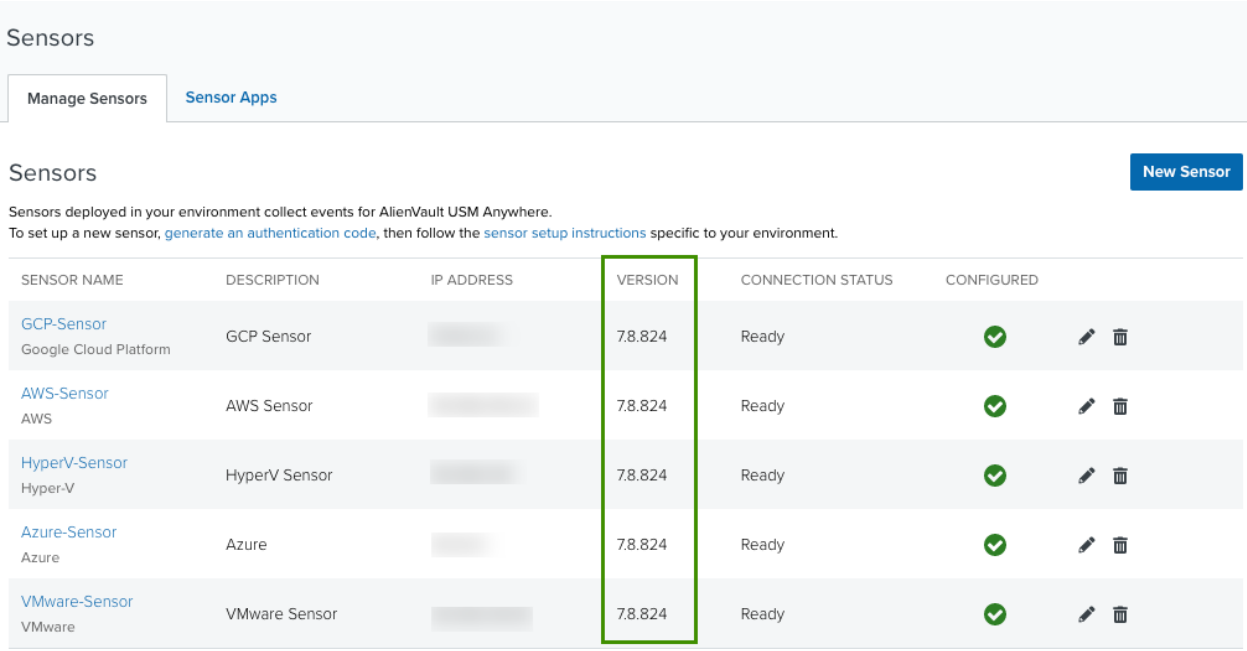

Warning: The VMware Sensor and Hyper-V Sensor require *all five network interface cards (NICs)* to be enabled; otherwise, the USM Anywhere update will fail. The NICs can remain disconnected.

VMware Sensor Deployment

AT&T Cybersecurity provides a VMware Sensor to monitor your virtual and physical on-premises infrastructure. When this USM Anywhere Sensor is deployed and configured for your USM Anywhere instance, security-related data is collected and sent to the AT&T Cybersecurity Secure Cloud for security analysis, threat correlation, and secure, compliance-ready data storage. You can also create jobs to collect log data through VMware, including operating system (OS) and database-level logs.

The VMware Sensor deployment includes a network-based intrusion detection system (NIDS) that monitors the networks connected to the listening interfaces. A deployed VMware Sensor supports a NIDS throughput of 600 Mbps, but this performance may vary depending on your environment, configurations, and other variables.

This section includes the following topics.

About VMware Sensor Deployment 25

Requirements for VMware Sensor Deployment 26

Create the VMware Virtual Machine 36

The OVF Package Is Invalid and Cannot Be Deployed - SHA256 Error 42

Set Up USM Anywhere on the VMware Virtual Machine 43

Connect the VMware Sensor to USM Anywhere 46

Complete the VMware Sensor Setup 51

About VMware Sensor Deployment

Through VMware, you can deploy a USM Anywhere Sensor in any of the virtual networks that you want to instrument for a network-based intrusion detection system (NIDS), including standard sensor features:

- Log data collection
- Authenticated asset scans
- Unauthenticated asset discovery scans

AT&T Cybersecurity distributes the VMware Sensor as an open virtual format (OVF) file that can be deployed through VMware vCenter or directly to a VMware ESX Hypervisor version 6.5 and later.

Important: Use VMware ESXi 6.5, you must have build 7388607 or later. Earlier builds have an issue with the OVF tools that will cause the sensor OVF deployment to fail.

If the OVF package is invalid and can't be deployed, and you get a SHA256 Error message, see The OVF Package Is Invalid and Cannot Be Deployed - SHA256 Error for more information.

The USM Anywhere Sensor deployed on VMware provides the ability to monitor the packets on networks that you select by attaching one of the Sensor network interfaces to a port configured in promiscuous mode on a virtual switch. This also requires that port mirroring is enabled on the upstream physical switch to which the ESXi host is connected.

Note: If your organization uses multiple subnets to enable communication between headquarters and remote offices, you do not need a sensor for each subnet. However, you will need a deployed VMware Sensor for each physical location that you want to monitor.

There is an option for you to enter credentials for either your vCenter or ESXi servers, which will allow the sensor to discover the virtual machines (VMs) registered on the ESXi servers through the VMware vSphere API. This enables the discovery of assets and also monitors user logins within your vSphere environment and feeds the information back to USM Anywhere.

Deployment Process Overview

The deployment process for an initial USM Anywhere Sensor on VMware consists of these primary tasks:

1. Review requirements for a VMware Sensor deployment.
2. Deploy a VMware Sensor by executing the `USM_sensor-node.ovf` file.
3. Configure the sensor on the VM.
4. Register the new sensor with your sensor authentication code to provision the USM Anywhere instance and connect the deployed sensor.
5. Complete your VMware Sensor configuration, including initial asset discovery.

Requirements for VMware Sensor Deployment

To ensure that you can successfully deploy a USM Anywhere Sensor in VMware and monitor all of your VMware resources, make sure the following requirements are met.

Minimum Requirements

These are the minimum requirements to set up and configure the USM Anywhere Sensor on VMware:

- Access to VMware ESXi 6.5 or later.
- Dedicated 4 CPUs and 12 GB of reserved memory.
- Dedicated 178 GB of disk space (128 GB data device and 50 GB root device).
- Internet connectivity to the network where you plan to install the VMware Sensor.

Important: Because the needs of a sensor differ based on the varying demands of different deployment environments and the complexity of events being processed, the number of events per second (EPS) a sensor can process varies.

Depending on your environment, you may need to deploy additional sensors to ensure that all events are processed.

Recommended Requirements

These are the recommended requirements to set up and configure the USM Anywhere Sensor on VMware:

- A VMware vSphere or VMware vCenter user account to use for USM Anywhere Sensor configuration with an assigned role that has permissions equivalent to the read-only default role.

  Note: The read-only role enables a user limited read access to the system without any other privileges. Credentials with this assigned role enable the deployed USM Anywhere Sensor to collect vCenter and vSphere events and run asset discovery.

- Installed VMware Tools for hosts in your vSphere or vCenter environment.

  With configured vSphere or vCenter credentials, the VMware Sensor uses the VMware APIs to run asset discovery. For hosts that do not have VMware Tools installed, the asset does not have an assigned IP address. This can result in the asset being missed from asset discovery or in duplicate assets created during subsequent discoveries. These tools also enable the sensor to collect more detailed information about the asset.

- If Dynamic Host Configuration Protocol (DHCP) is not available, a configured static IP for the management interface and local Domain Name System (DNS) information.

  Important: AT&T Cybersecurity strongly recommends assigning a static IP address to deploy the USM Anywhere Sensor.

  If DHCP changes the IP address of the sensor, you must update all the IP addresses on all the devices that are forwarding logs to the sensor through syslog.

- Port mirroring set up for network monitoring (see Direct Traffic from Your Physical Network to the VMware Sensor for more information).
- Administrative credentials for devices that require configuration to forward logs to the VMware Sensor.
- Administrative credentials for remote hosts to support authenticated asset scans.
- Configuration on firewall or other security device to send UDP or TCP syslog (if it is capable of exporting security logs through UDP or TCP syslog).
- Network topology information to run asset discovery.

- To access network-based intrusion detection system (NIDS) functionality on the sensor, an ethernet port on the host must be available to receive data from a Switched Port Analyzer (SPAN) or Test Access Point (TAP) port.

Sensor Ports and Connectivity

Before deploying a USM Anywhere Sensor, you must configure your firewall permissions to enable the required connectivity for the new sensor. Initial deployment of a sensor requires that you open egress and outbound ports and protocols in the firewall for communication with USM Anywhere and AT&T Cybersecurity Secure Cloud resources. The sensor receives no inbound connections from outside the firewall.

Note: To launch the USM Anywhere Sensor web UI during the initial setup, you need to allow inbound traffic to the sensor IP address through TCP port 80. You can remove access to this port after the sensor successfully connects to USM Anywhere. You do not need to allow inbound traffic to this port from the Internet.

The following tables list the inbound and outbound ports.

Sensor Ports and Connectivity (Outbound Ports)

Type	Ports	Endpoints	Purpose
TCP	443	update.alienvault.cloud	Communication with AT&T Cybersecurity for initial setup and future updates of the sensor.
TCP	443	reputation.alienvault.com	Ongoing communication with AT&T Alien Labs™ Open Threat Exchange® (OTX™).

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	otx.alienvault.com	Ongoing communication with OTX to retrieve vulnerability scores. Connecting to otx.alienvault.com is not required but highly recommended. OTX uses the AWS CloudFront services. Refer to the AWS IP address ranges page when you deploy a new sensor. This page contains the current IP address ranges for the service and instructions on how to filter the addresses.
TCP	443	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
TCP	9443	vCenter Server	Authenticate sensor to ESXi. Connect to the vCenter for VMware configuration to gather data directly from vCenter.
SSL	443	storage-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send and retrieve backups.
SSL	443	metrics-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send metrics and messages.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
SSL/TCP	443	api-parameters-<REGION>-prod.alienvault.cloud[1] api-message-proxy-<REGION>-prod.alienvault.cloud api.message-proxy.<REGION>.prod.alienvault.cloud	Ongoing communication with USM Anywhere. It is only necessary to allowlist the address that corresponds to the region where your USM Anywhere instance is hosted.
SSL/TCP	7100	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
UDP	53	DNS Servers (Google Default)	Ongoing communication with USM Anywhere.
UDP	123	0.pool.ntp.org 1.pool.ntp.org 2.pool.ntp.org 3.pool.ntp.org	Sync with network time protocol (NTP) services.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	22 and 443	prod-usm-saas-tractorbeam.alienvault.cloud prod-gov-usm-saas-tractorbeam.gov.alienvault.us (for AT&T TDR for Gov)	SSH communications with the USM Anywhere remote support server. See Troubleshooting and Remote Sensor Support for more information about remote technical support through the USM Anywhere Sensor console.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	geoip-ap-northeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-south-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-2-prod.alienvault.cloud/geo-ip/sensor geoip-ca-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-me-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-sa-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-us-gov-west-1-prod-gov.gov.alienvault.us/geo-ip/sensor (for AT&T TDR for Gov)	Allows resolution of IP addresses for geolocation services. It is only necessary to allowlist the GeoIP address that corresponds to the region where your USMA instance is hosted.

1

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

2

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

3

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

4

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

5

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

Sensor Ports and Connectivity (Inbound Ports)

Type	Ports	Purpose
SSH	22	Inbound method for secure remote login from a computer to USM Anywhere.
HTTP	80	Inbound communication for HTTP traffic.
UDP (RFC 3164)	514	USM Anywhere collects data through syslog over UDP on port 514 by default.
TCP (RFC 3164)	601	Inbound communication for reliable syslog service. USM Anywhere collects data through syslog over TCP on port 601 by default.
TCP (RFC 5424)	602	USM Anywhere collects data through syslog over TCP on port 602 by default.
Traffic Mirroring	4789	Inbound communication for virtual extensible local area network (VXLAN).

Sensor Ports and Connectivity (Inbound Ports) (Continued)

Type	Ports	Purpose
WSMANS	5987	Inbound WBEM WS-Management HTTP over Secure Sockets Layer/Transport Layer Security (SSL/TLS) (NXLog).
TLS/TCP (RFC 3164)	6514	USM Anywhere collects TLS-encrypted data through syslog over TCP on port 6514 by default.
TLS (RFC 5424)	6515	USM Anywhere collects data through syslog over TLS on port 6515 by default.
TCP	9000	Inbound communication used internally for HTTP sensor traffic.
Graylog	12201	Inbound communication for Graylog Extended Log Format (GELF).

USM Anywhere IP Addresses for Allowlisting

Your sensor is connected to a USM Anywhere instance deployed in one of the Amazon Web Services (AWS) endpoint regions based on your location. If you need to configure your firewall to allow communication between the sensor and the USM Anywhere instance, refer to the following table with the reserved IP address ranges for each region.

Important: The Update Server and the AlienVault Agent always use the 3.235.189.112/28 range no matter where your USM Anywhere is deployed. The AT&T TDR for Gov Update Server uses the 3.32.190.224/28 range.

Note: The regional IP ranges listed in this table are limited to the control nodes (subdomain). You must also meet all requirements provided in the Sensor Ports and Connectivity (Outbound Ports) table.

AWS Regions Where USM Anywhere Instance Is Available

Code	Name	Reserved Static IP Address Ranges
ap-northeast-1	Asia Pacific (Tokyo)	18.177.156.144/28 3.235.189.112/28 44.210.246.48/28
ap-south-1	Asia Pacific (Mumbai)	3.7.161.32/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-1	Asia Pacific (Singapore)	18.143.203.80/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-2	Asia Pacific (Sydney)	3.25.47.48/28 3.235.189.112/28 44.210.246.48/28
ca-central-1	Canada (Central)	3.96.2.80/28 3.235.189.112/28 44.210.246.48/28
eu-central-1	Europe (Frankfurt)	18.156.18.32/28 3.235.189.112/28 44.210.246.48/28
eu-west-1	Europe (Ireland)	3.250.207.0/28 3.235.189.112/28 44.210.246.48/28

AWS Regions Where USM Anywhere Instance Is Available (Continued)

Code	Name	Reserved Static IP Address Ranges
eu-west-2	Europe (London)	18.130.91.160/28 3.235.189.112/28 44.210.246.48/28
me-central-1	Middle East (UAE)	3.29.147.0/28 3.235.189.112/28 44.210.246.48/28
sa-east-1	South America (São Paulo)	18.230.160.128/28 3.235.189.112/28 44.210.246.48/28
us-east-1	US East (N. Virginia)	3.235.189.112/28 44.210.246.48/28
us-west-2	US West (Oregon)	44.234.73.192/28 3.235.189.112/28 44.210.246.48/28
us-gov-west-1	AWS GovCloud (US-West)	3.32.190.224/28

Create the VMware Virtual Machine

Role Availability ✗ Read-Only ✗ Investigator ✗ Analyst ✓ **Manager**

AT&T Cybersecurity provides a download package, which contains the VMware Open Virtualization Format (OVF) template that you can use to import and deploy the USM Anywhere Sensor on a VMware ESXi host.

Important: Use VMware ESXi 6.5, you must have build 7388607 or later. Earlier builds have an issue with the OVF tools that will cause the sensor OVF deployment to fail.

If the OVF package is invalid and can't be deployed, and you get a SHA256 Error message, see The OVF Package Is Invalid and Cannot Be Deployed - SHA256 Error for more information.

The following procedure describes the standard VMware ESXi Embedded Host Client, which is a native HTML and JavaScript application served directly from your ESXi host. Before you begin this procedure, make sure that your ESXi 6.5 host is updated to build 7388607 or later and that the web client is updated to build 7119706 or later. Refer to these VMware online resources for the latest download files and information:

- VMware ESXi Patch Tracker: https://esxi-patches.v-front.de/ESXi-6.5.0.html
- VMware ESXi Embedded Host Client: https://labs.vmware.com/flings/esxi-embedded-host-client

If you are using VMware vCenter to manage your VMware ESXi hosts and using the VMware vSphere web client, refer to the documentation provided by VMware and extrapolate from this procedure.

To load the OVF and deploy the USM Anywhere Sensor Virtual Machine (VM)

1. Go to the USM Anywhere Sensor Downloads page and click the icon of your specific sensor. After clicking, your browser starts to download the USM Anywhere Sensor package. Depending on your Internet connection, the download can take 30 minutes or more.
2. Extract the USM Anywhere Sensor package to any folder on the machine where you are using the vSphere client.
3. In your ESXi Web Client, click **Create/Register VM**.

   This opens the *New virtual machine* wizard.

4. In the *Select creation type* page, choose **Deploy a virtual machine from an OVF or OVA file** and click **Next**.

5. Enter a name for the new VM and select the template files.

6. Browse to the location where you extracted the files from the sensor download package, select the OVF and VMDK files, and click **Next**.

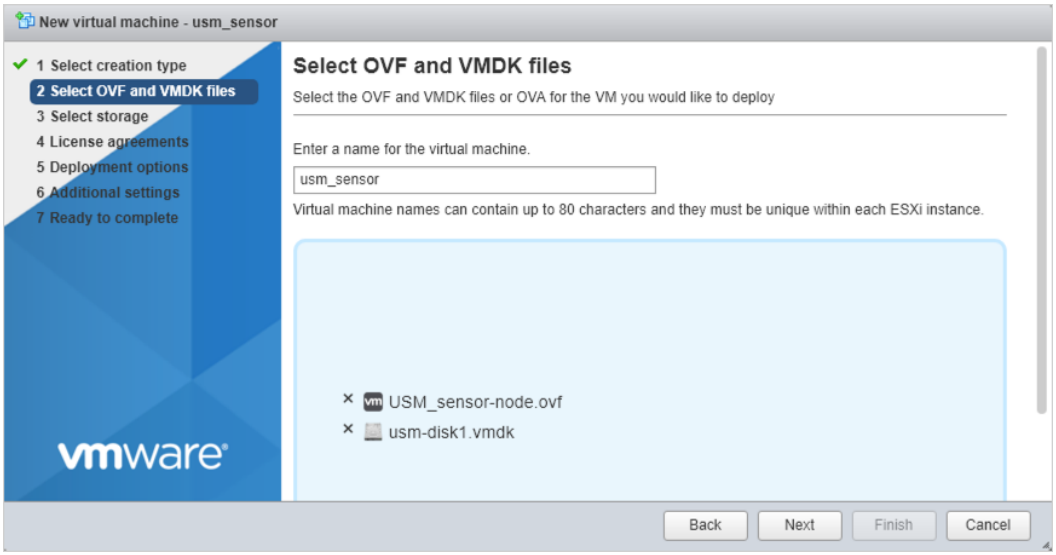

7. For each of the wizard pages, set the parameters as needed for your network and click **Next**:

   - **Select storage**: Select the datastore you want to use for the VM.
   - **Deployment options**: Set the networking and deployment for the VM.

     The primary network requires internet connectivity and an IP address that is routed to provide the access to USM Anywhere. The other interfaces passively monitor network traffic in promiscuous mode.

     Warning: The VMware Sensor requires *all five network interface cards (NICs)* to be enabled; otherwise, the USM Anywhere update will fail. The NICs can remain disconnected.

     See Configure Network Interfaces for On-Premises Sensors for more information about these interfaces.

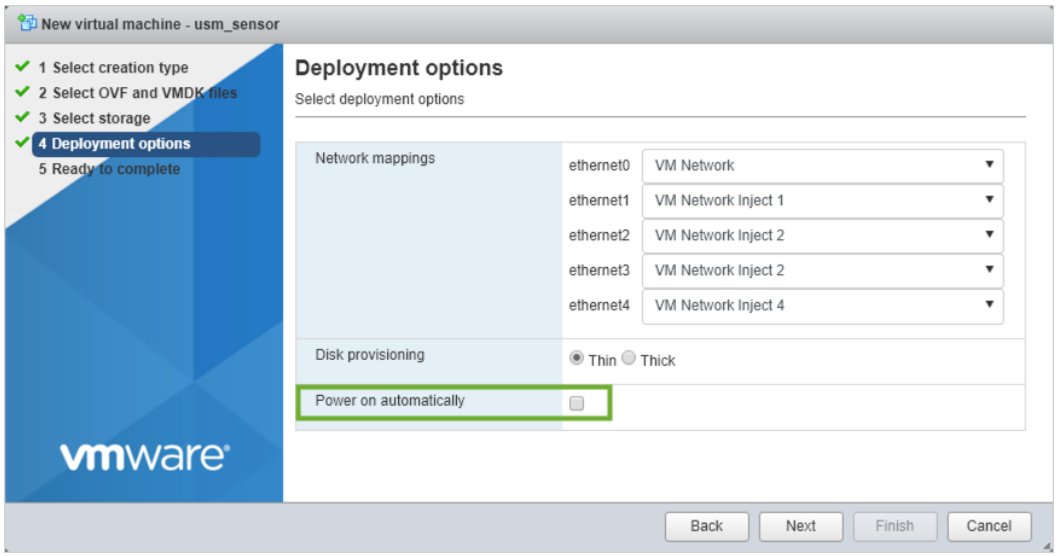

- Clear the **Power on automatically** option. It is important to create the VM without powering it on so that you can configure the ISO file before the initial boot.

8. In the Ready to complete panel, review the configuration and click **Finish**.

   An alert appears that says "A required disk image was missing". Ignore this message, because you will address the disk image in the next step.

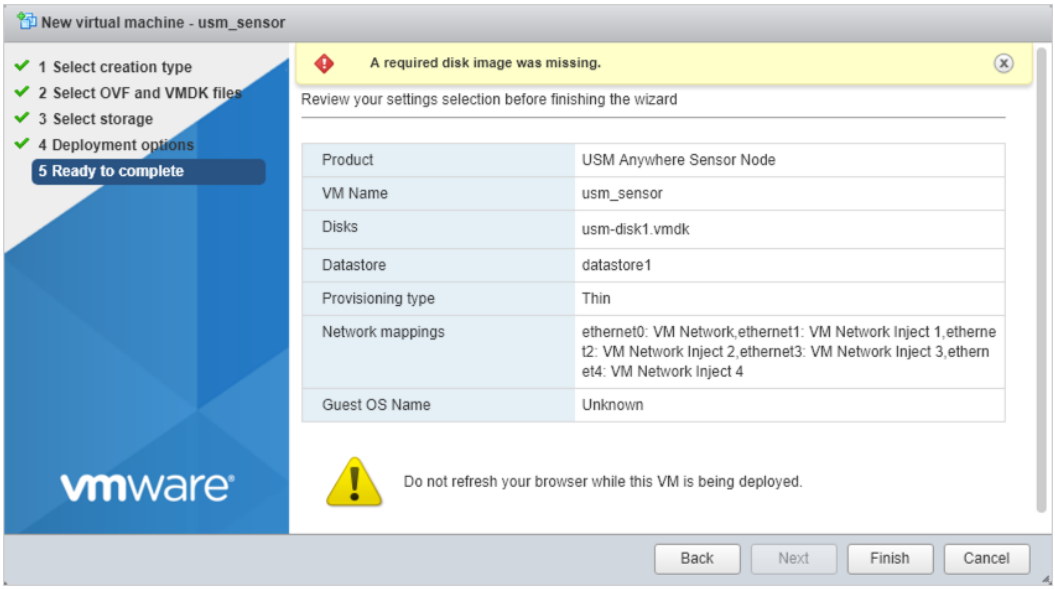

Import of the OVF and VMDK files and the creation of the virtual image can take some time. You can check the status in the Recent Tasks window.

9. After the VM is created but not yet powered on, configure the correct ISO file, `deploy_config.iso`, for the datastore:

> **Note:** Sometimes a different ISO file is selected by default causing the deployment to fail.

> **Warning:** You must complete this step and ensure that the ISO is mounted before you start the sensor VM for the first time.
>
> If you see `REPLACEME` as the initial login password in the sensor welcome screen when you connect to the VM, it is most likely that the ISO was not mounted before the sensor was started. If this happens, you must shut down the VM, complete this step so that the ISO is configured for the datastore, and then begin the deployment process anew.

- Upload the `deploy_config.iso` file to your datastore. You can use the datastore browser in the web client to select the ISO file and upload it.
- Select the new sensor VM in the left pane and scroll to the Hardware Configuration section.
- Locate CD/DVD drive 1 in the hardware list and click **Select disc image**.

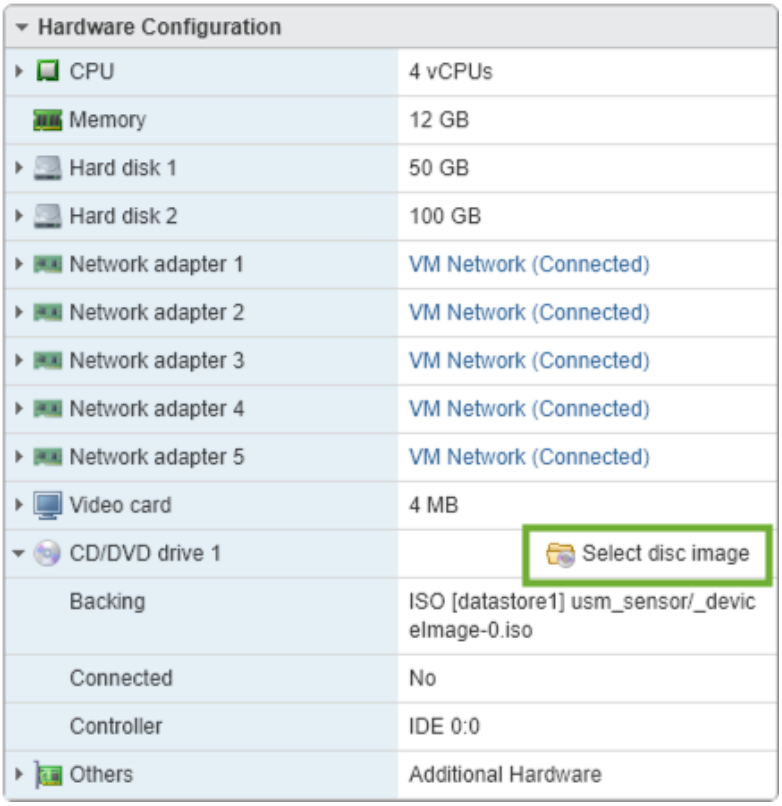

- Navigate the datastore and select the `deploy_config.iso` file.

- Click **Select**.

10. In the toolbar, click **Power on** to start the USM Anywhere Sensor VM.

> After starting the sensor initialization process, the USM Anywhere Sensor VM thumbnail displays a green startup screen during this process, which can take a few minutes to complete.

11. Connect to the console for the USM Anywhere Sensor using one of the following methods:

    - In the toolbar, click **Console**.
    - Click the thumbnail for the sensor VM.

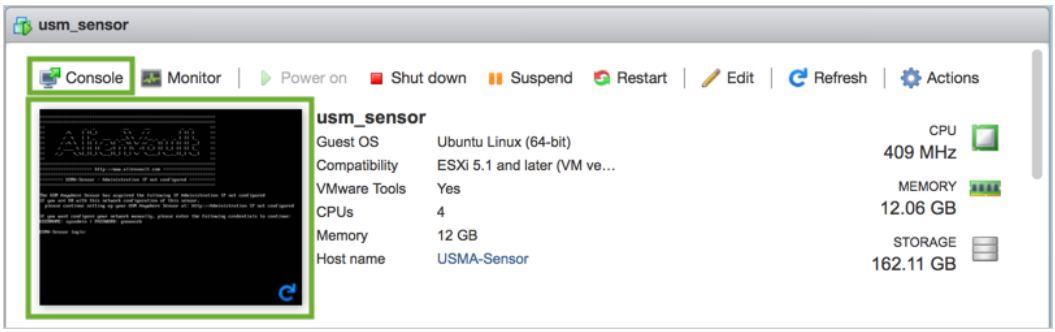

The USM Anywhere Sensor screen provides the initial login password to use when you complete the sensor setup. It also displays the URL that you use to access USM Anywhere and complete the sensor registration and connection.

The OVF Package Is Invalid and Cannot Be Deployed - SHA256 Error

You may receive an error during the VMware Open Virtualization Format (OVF) deployment with certain older versions of VMWare (sub 6). This is due to the legacy SHA changes that were made by VMware.

You can also receive this error when deploying the VMware Open Virtual Appliance (OVA) via the VMware vSphere Client fails:

```
The OVF package is invalid and cannot be deployed.
The following manifest file entry (line 1) is invalid: SHA256 (xxxxxxxx.ovf).
```

This issue occurs because the vSphere Client does not support the SHA256 hashing algorithm, which the vSphere Integrated Containers (VIC) OVA was made of. This also affects any OVA deployments via VMware PowerCLI when using the Get-Ovf-Configuration cmdlet.

To resolve this issue, deploy vIC via the vSphere Web Client or VMware ESXI Embedded Host Client because they both support SHA256. However, if you still want to automate your deployments, you must convert the OVA from the Cryptographic Hash Algorithm SHA256 to SHA1. To do this, you can use OVF Tool, which is available for all operating systems (OSes) at https://developer.vmware.com/web/tool/4.4.0/ovf.

To do the conversion, run the following command:

```
ovftool.exe --shaAlgorithm=SHA1 /path/to/the/original/ova_file.ova
/path/to/the/new/ova/file-SHA1.ova
```

Note: The OVF Tool doesn't install on the OS. You must run an elevated command prompt from the folder that contains the OVF Tool.

Important: If you need more information, contact AT&T Cybersecurity Technical Support for assistance.

Set Up USM Anywhere on the VMware Virtual Machine

Role Availability	Read-Only	Investigator	Analyst	Manager

There is some configuration required within the sensor console on the virtual machine (VM). The sensor console also provides tools for troubleshooting the USM Anywhere Sensor. After this initial configuration, you complete the sensor configuration in the USM Anywhere web user interface (UI).

Perform these initial configuration tasks on the VMware VM, using the USM Anywhere Sensor console.

Change the Administrative Password and Keyboard Layout

Follow these instructions to change the administrative password and keyboard layout.

To change the administrative password and keyboard layout

1. Log in using the credentials displayed in the console screen.

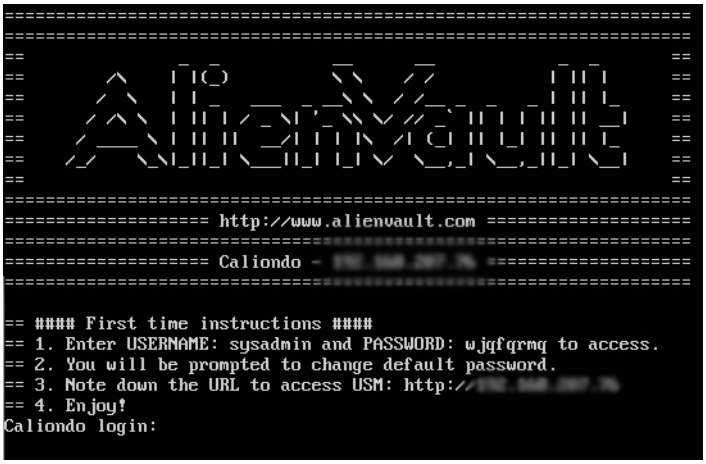

2. (Optional.) Configure the keyboard if you use a keyboard layout other than the U.S. default.

3. Set a new password for the *sysadmin* user.

> **Important:** During the installation, your system acquires the initial IP address through Dynamic Host Configuration Protocol (DHCP). If DHCP is not enabled, you must configure it manually.
>
> AT&T Cybersecurity *strongly* recommends assigning a static IP address to the USM Anywhere Sensor as a best practice. This allows for proper log forwarding and network architecture.

- If your system sets an IP address automatically, note the web URL (IP address). You will need the URL when you exit from the console and follow the instructions in Connect the VMware Sensor to USM Anywhere.

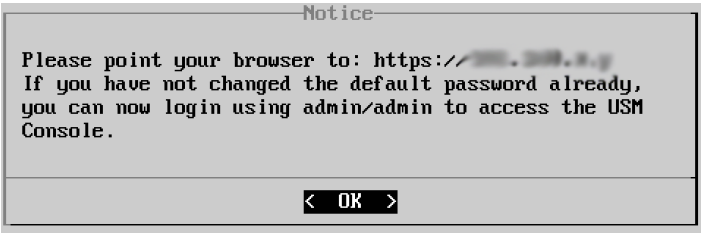

- If your system does not set an IP address automatically, a message box confirms that the system was unable to acquire an IP address from a DHCP server after you change the sysadmin password.
- In this case, you must manually set a static IP address so that it remains unchanged in the future.

Configure a Static IP Address

Follow these instructions to configure a static IP address.

To configure a static IP address

1. Go to **Network Configuration > Configure Management Interface > Set a Static Management IP Address**.
2. Enter the IP address, subnet, and gateway information in each screen.
3. Press **Enter**.

Important: DNS settings are not maintained when a static IP address is configured. If you configure a static IP address, you must configure the DNS network settings for successful sensor activation.

Configure Domain Name System

Follow these instructions to configure the Domain Name System (DNS).

Important: When the USM Anywhere Sensor performs an asset scan, it must access the local Domain Name System (DNS) server to resolve local host names. The sensor uses reverse DNS to look up the hostname through the discovered IP address.

Note: When deploying your VMware Sensor in a DHCP environment, the DNS server is automatically set to retrieve via DHCP. This can be configured later in your sensor's settings. See Deploying Your Sensor in a DHCP Environment for more information about USM Anywhere Sensors in a DHCP environment.

To configure DNS

1. Go to **Network Configuration > Configure DNS**.
2. Enter the primary DNS and press **Enter**.

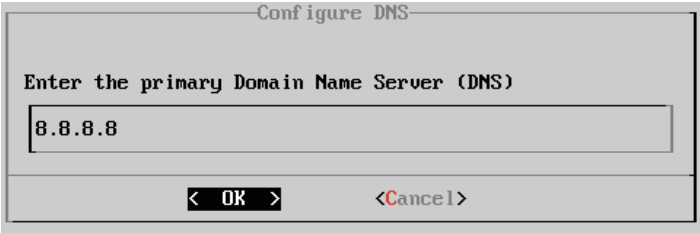

3. (Optional.) Enter the secondary DNS and press **Enter**.

   A text box opens to confirm that you want to apply changes.

4. Press **Enter**.

> **Note:** Check your settings through **Network Configuration > View Network Configuration**.

Connect the VMware Sensor to USM Anywhere

Role Availability	Read-Only	Investigator	Analyst	Manager
	✗	✗	✗	✓

After deploying the VMware Sensor, you must connect it to USM Anywhere through registration.

Obtain the Authentication Code

You must enter an authentication code when registering the USM Anywhere Sensor. How to obtain the authentication code depends on your USM Anywhere instance and whether this is the first sensor you're deploying.

Instructions for USM Anywhere customers:

If this is your first USM Anywhere Sensor, you must register the sensor using the initial authentication code (starts with a "C") received from AT&T Cybersecurity. With this code, the registration process provisions a new USM Anywhere instance and defines its attributes, such as how many sensors to allow for connection, how much storage to provide, and what email address to use for the initial user account. After registration, you will gain access to the

sensor through the USM Anywhere web user interface (UI), where you can complete the sensor setup.

If you are deploying additional sensors, you must generate the authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Instructions for AT&T TDR for Gov customers:

AT&T Cybersecurity has already provisioned the AT&T Threat Detection and Response for Government (AT&T TDR for Gov) instance for you, therefore you won't receive an authentication code for your sensor. This is true regardless if it's the first sensor or additional sensors you're deploying. However, for the first sensor, you'll receive a link to access your instance.

For every sensor you deploy, you must generate an authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Register Your Sensor

You perform this procedure after deploying the USM Anywhere Sensor within your VMware environment. You can acquire the IP address when you set your management interface settings in the sensor console.

To register your sensor

1. Open a web browser and enter the IP address.

   This opens the *Welcome to USM Anywhere Sensor Setup* page, which prompts you to provide the information for registering the sensor with your new USM Anywhere instance.

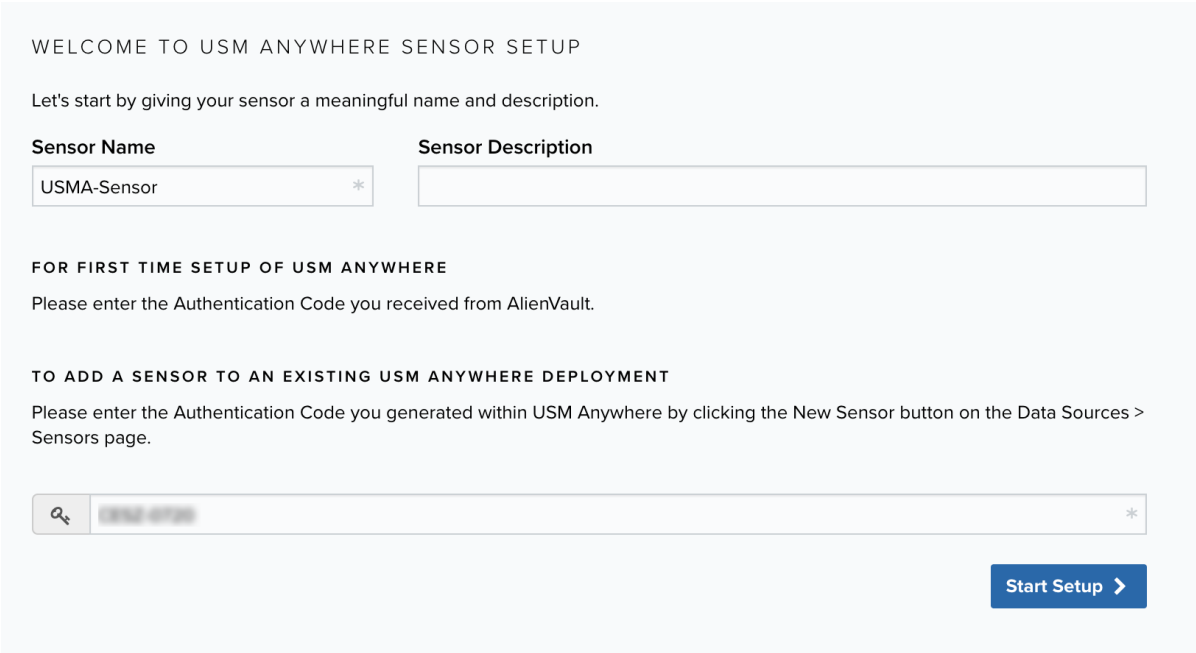

2. Enter a sensor name and sensor description.
3. Paste the authentication code into the field with the key icon ( ).
4. Click **Start Setup** to start the process of connecting the USM Anywhere Sensor.

   It takes about 20 minutes to provision your USM Anywhere instance upon registration of your initial sensor. When this instance is provisioned and running, you'll see a welcome message that provides an access link.

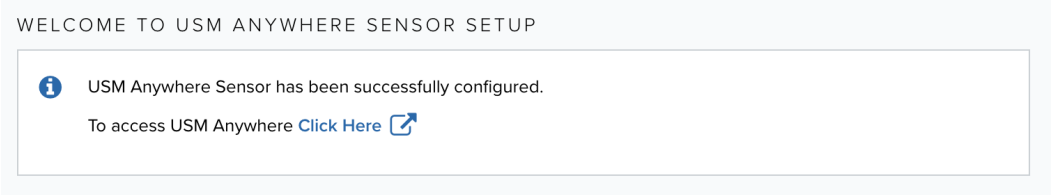

   Use this link to open the secured web console for your USM Anywhere instance. You and the other USM Anywhere users in your organization can access this console from a web

browser on any system with internet connectivity.

> **Note:** If this is your first deployment, you'll also receive an email from AT&T Cybersecurity that provides the access link to USM Anywhere.

Configure the Initial Login Credentials

When you link to a newly provisioned USM Anywhere instance, you must configure the password for the initial user account. This is the default administrator as defined in your subscription.

To configure login credentials

1. In the welcome message, click the link.

   This displays a prompt to set the password to use for the default administrator of USM Anywhere.

2. Enter the password, and then enter it again to confirm.

   Keep in mind these points when you are logging in:

   - The login credentials that you set will apply to any USM Anywhere™ and USM Central™ you have access to.
   - USM Anywhere requires all passwords to have a minimum length of 8 characters and a maximum length of 128 characters.
   - The password must contain numerical digits (0-9).
   - The password must contain uppercase letters (A-Z).
   - The password must contain lowercase letters (a-z).
   - The password must contain special characters, such as hyphen (-) and underscore (_).

> **Note:** USM Anywhere passwords expire after 90 days. When your password expires, USM Anywhere enforces a password change when you next log in. A new password must be different from the previous four passwords.After 45 days of inactivity, your user account will be locked. Manager users can unlock inactive accounts.

3. Click **Save & Continue**.
4. When the login page opens, enter the password you just set and click **Login**.

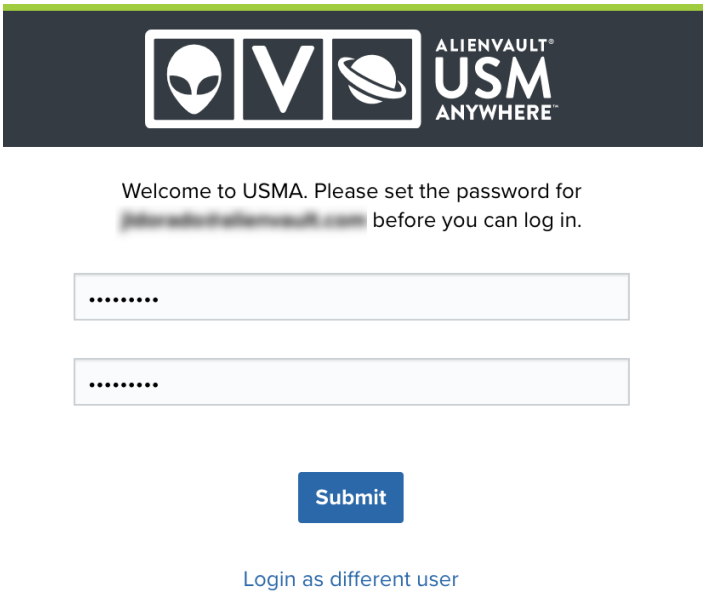

Verify That Your Sensor Is Running

It's a good idea to verify that the USM Anywhere Sensor is running. It also gives you the chance to watch the sensor actively working to find all of your assets and to record events from the start.

> **Note:** Verify that the sensor is running before performing the configuration. You can keep one web browser tab with the Welcome to USM Anywhere page in the background while you perform the verification on a different tab.

To verify that your new sensor is running

1. In USM Anywhere, go to **Data Sources > Sensors**.

   You should now see your sensor in the page. See in the *USM Anywhere User Guide* for more information.

   After a few minutes, USM Anywhere locates your assets and starts generating events.

2. You can review the activity in two locations:
   - From the primary task bar, select **Environment > Assets**.
   - From the primary task bar, select **Activity > Events**.

> **Note:** It could take up to six minutes before events appear. Make sure to refresh your browser from time to time to display the current data.

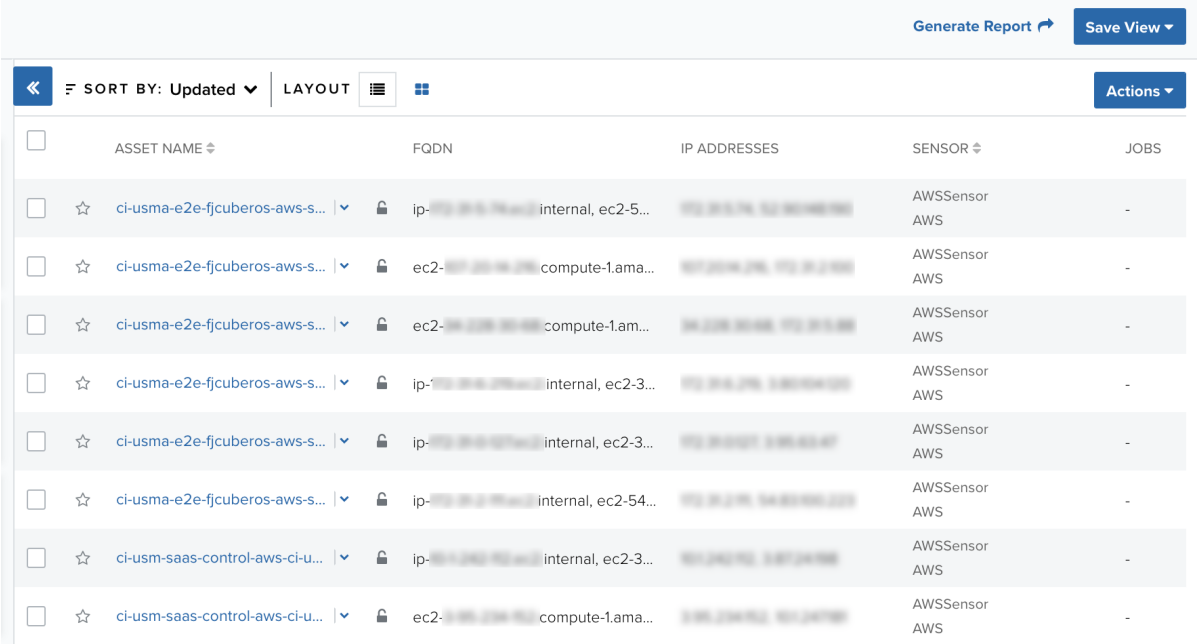

See the *USM Anywhere User Guide* for more information about using the Assets and Events pages in USM Anywhere.

Complete the VMware Sensor Setup

Role Availability	Read-Only	Investigator	Analyst	Manager

After you initialize a new USM Anywhere Sensor, you must configure it in the Setup Wizard. As you complete the VMware Sensor configuration, USM Anywhere performs specific actions, like running an asset discovery scan and collecting logs.

Accessing the Setup Wizard

The Setup Wizard is accessible under the following circumstances:

- After you first log in to the USM Anywhere web user interface (UI) and see the Welcome to USM Anywhere page, click **Get Started** to launch the Setup Wizard.
- If you have already registered one USM Anywhere Sensor but did not complete the setup before logging out, the USM Anywhere Sensor Configuration page launches automatically at your next login to remind you to finalize configuration of the sensor. From that page, you click **Configure** to launch the Setup Wizard and complete the sensor configuration.
- If you registered an additional USM Anywhere Sensor, but did not complete the setup, the Sensors page displays an error (⊗) in the Configured column. See in the *USM Anywhere User Guide* for more information.

  Go to **Data Sources > Sensors**, and then click the sensor name to complete the sensor configuration. See in the *USM Anywhere User Guide* for more information.

Configuring the VMware Sensor in the Setup Wizard

The first time you log in from the Welcome to USM Anywhere web page, the Setup Wizard prompts you to complete the configuration of the first deployed sensor. Thereafter, you can use the Sensors page to configure an additional sensor or to change the configuration options for a deployed sensor. See in the *USM Anywhere User Guide* for more information.

> **Note:** You must have already configured your network interfaces for VMware. See Set Up USM Anywhere on the VMware Virtual Machine for more information.

Within the Setup Wizard, complete the configuration on each page.

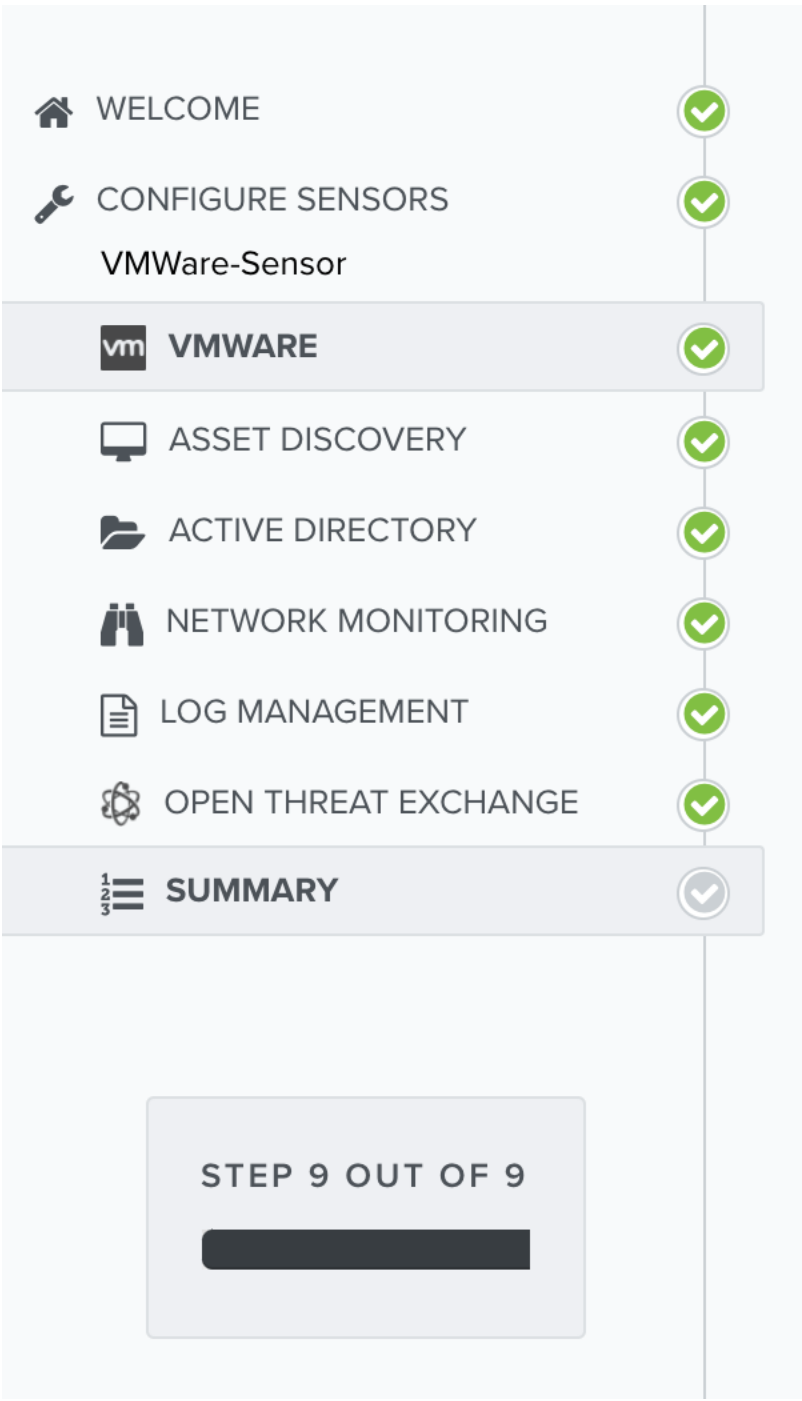

VMware Configuration

The first page in the Setup Wizard for a VMware Sensor is VMware Configuration. The information that you provide on this page enables USM Anywhere to discover assets in your VMware environment and also collect events from that environment.

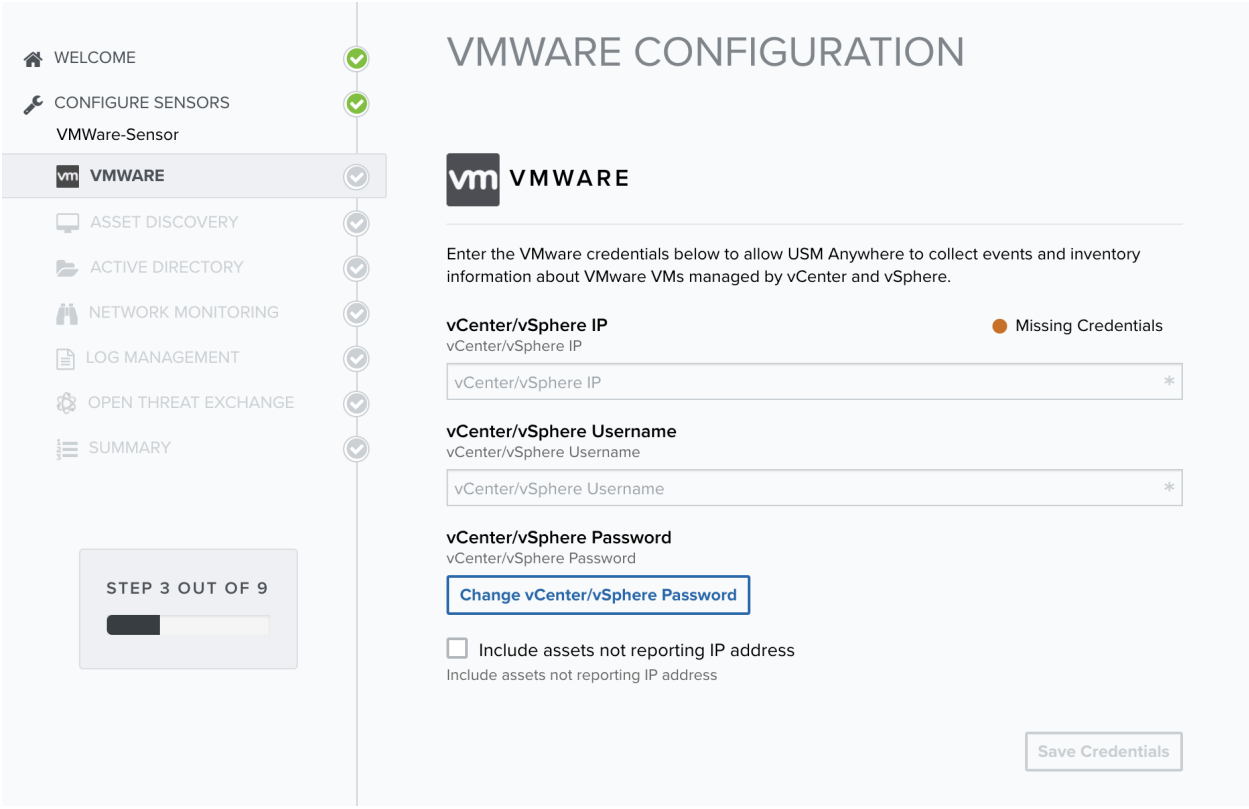

To complete the VMware configuration step

1. Enter your VMware vCenter or VMware vSphere IP address and user credentials.

   This should be a user account for the VMware environment with an assigned role that has permissions equivalent to the Read Only default role, which enables limited read access to the system without any other privileges.

2. (Optional.) If you do not have VMware Tools installed on all virtual machines (VMs) in your vCenter or VMware vSphere environment but want the sensor to be able to discover them, select the **Include assets not reporting IP address** checkbox.

Important: VMware Tools identifies and associates all network assets reporting to the sensor, including the assets that don't report their IP addresses. However, if you have not used VMware Tools to map those assets on your network and you click the **Include assets not reporting IP address** checkbox during the configuration process, using asset discovery on the VMware Sensor can create duplicates of those unmapped assets that don't report their IP addresses because they were not uniquely identified and mapped previously by VMware Tools.

AT&T Cybersecurity recommends that you install VMware Tools on all hosts in your vSphere or vCenter environment. This option is provided to handle situations where you are unable to do this.

Note: If you change this option while modifying the settings of a sensor that is already configured, you must provide the IP address and user credentials again.

3. Click **Save Credentials**.
4. Click **Next**.

   The wizard opens the next page in the setup process, Asset Discovery.

Asset Discovery

When you move forward to Asset Discovery, a dialog box automatically opens and prompts you to allow asset scanning. USM Anywhere must discover your assets to enable security monitoring on them.

To complete the asset discovery task

1. Click **Yes** to start the automatic asset discovery.

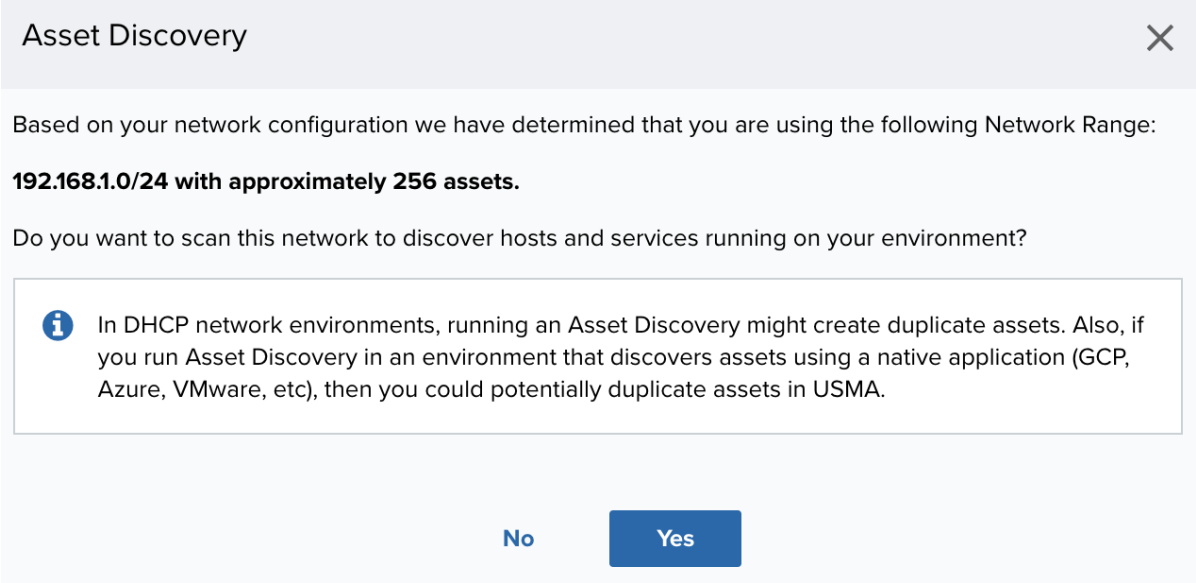

Or if you prefer to add the assets manually or scan another network, click **No** and skip to the next step.

During the automated scan, the Scan Networks status bar opens and displays the number of assets detected in your network range.

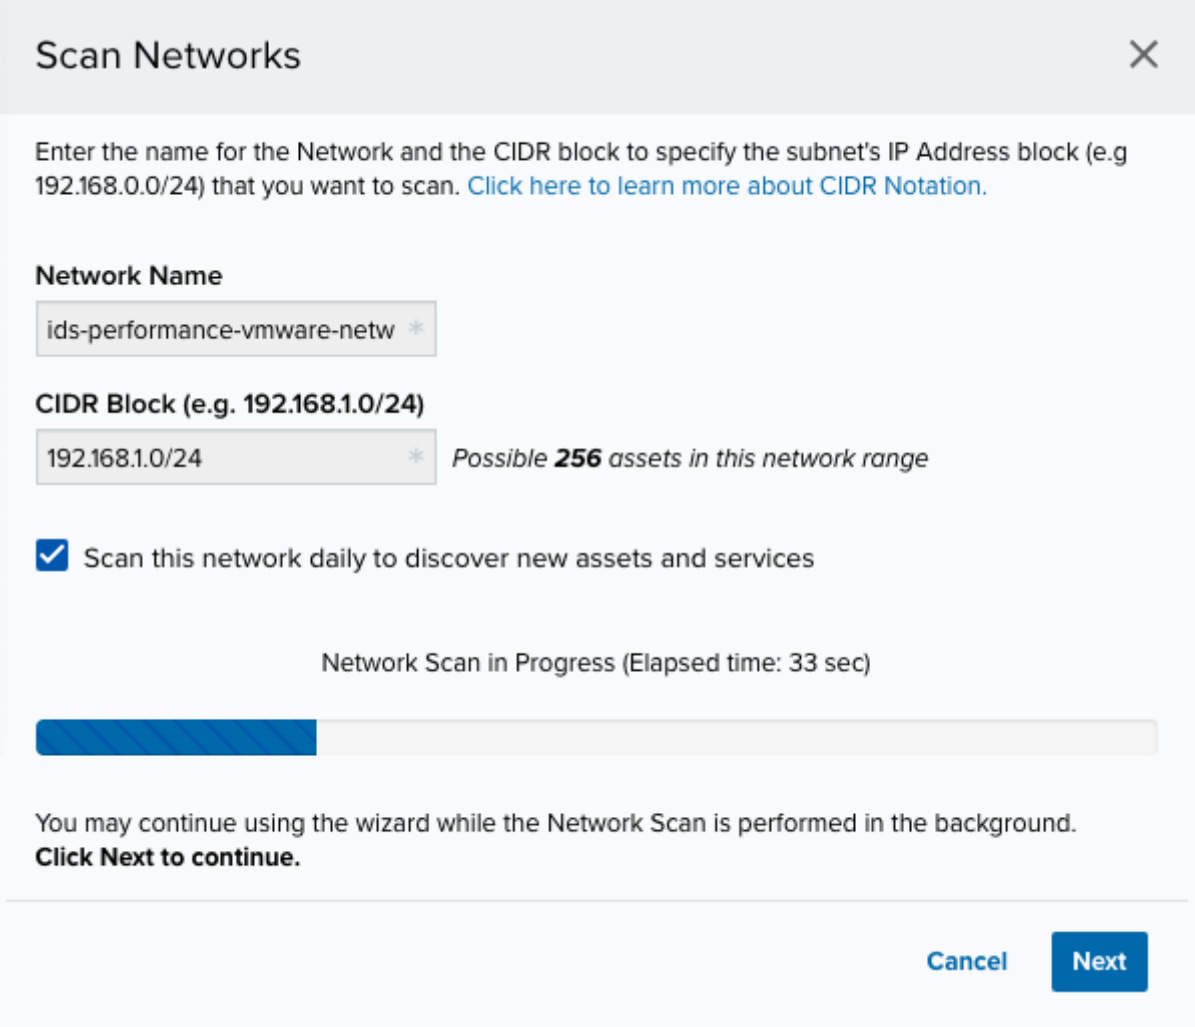

When the scan stops, you have these options:

- Click **Scan Another** to scan a different set of assets
- Click **Next** to continue with asset discovery setup options

When the initial asset scan dialog box closes, the Asset Discovery page displays status information for an ongoing scan or any discovered assets for completed scans.

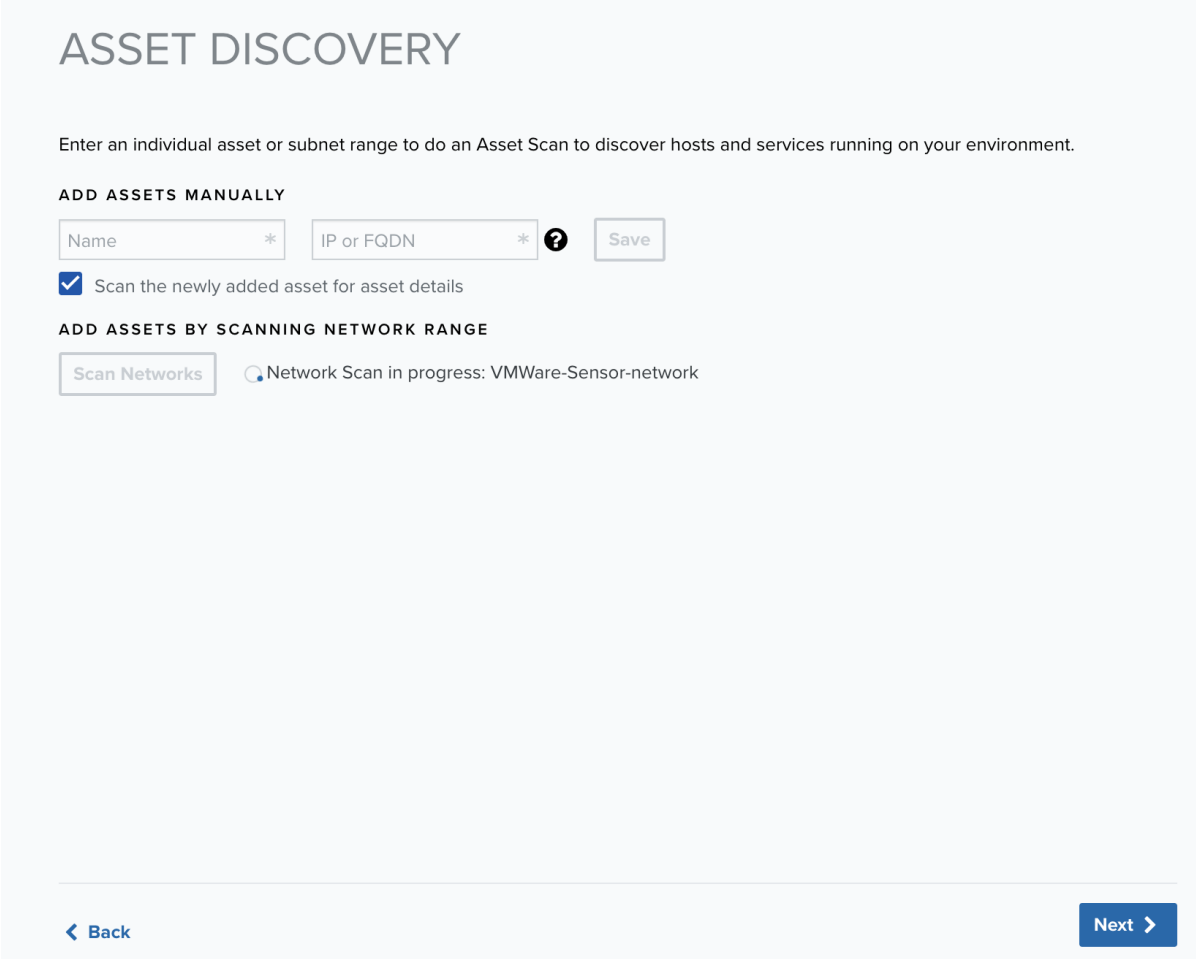

2. (Optional.) **Add assets manually**

   Enter the name and IP address or fully qualified domain name (FQDN) to specify an asset for discovery. The scan option is selected by default. Click **Save** to add the asset.

   You can repeat this for each individual asset you want to add.

3. (Optional.) **Add assets by scanning network range**

   Click **Scan Networks** to scan a network range that you specify. This runs asset discovery to scan hosts and services running on the specified network range.

4. When all the needed assets are discovered, click **Next** at the bottom of the page.

   The wizard opens the next page in the setup process, Active Directory.

Active Directory

The optional Active Directory (AD) setup page configures USM Anywhere to collect information from your AD account. To monitor Microsoft Windows systems effectively, USM Anywhere needs access to the AD server to collect inventory information.

Note: This configuration is only for one AD server. If you want to scan different AD servers, you must create an AD scan job for each of them. See Scheduling Active Directory Scans from the Job Scheduler Page for more information.

AT&T Cybersecurity recommends that you create a dedicated AD account with membership in the Domain Admins group to be used by USM Anywhere to log in to the Windows systems. You also need to activate Microsoft Windows Remote Management (WinRM) in the domain controller and in all the hosts that you want to scan. You can do this by using a group policy for all the systems in your AD.

Important: Before this feature is fully functional, you must configure access to the USM Anywhere Sensor on the AD server. See Granting Access to Active Directory for USM Anywhere for more information.

To complete the AD access configuration

1. Provide the AD credentials for USM Anywhere:
   - **Active Directory IP Address**: Enter the IP address for the AD server.
   - **Username**: Enter your username as admin of the account.
   - **Password**: Enter your admin's password.
   - **Domain**: Enter the domain for the AD instance.

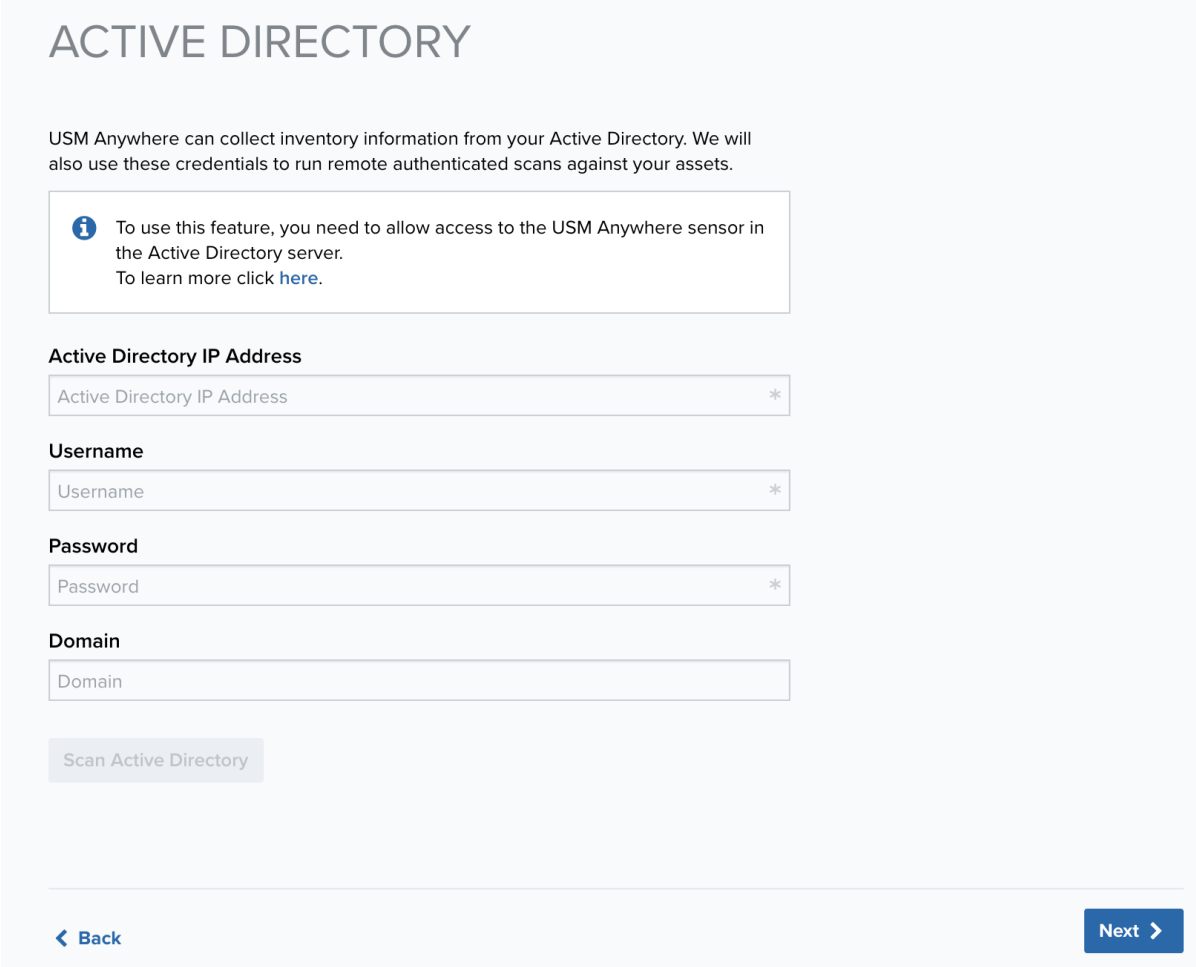

2. Click **Scan Active Directory**.

   After a successful launch of the scan, a confirmation dialog box opens.

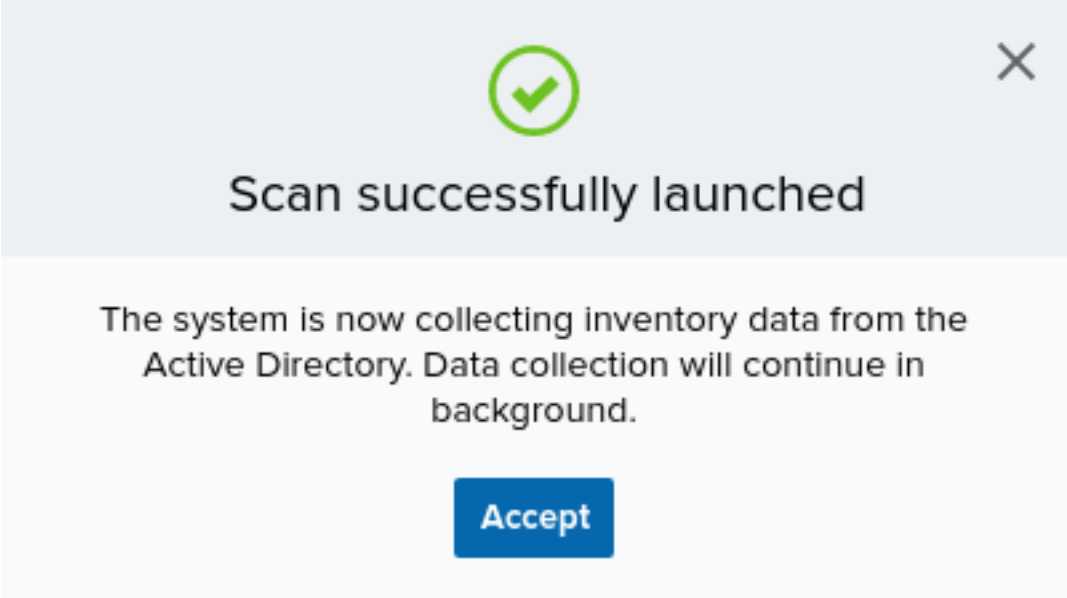

3. Click **Accept**.

   The scan continues in the background.

   Upon completion, another dialog box opens and provides information about the number of assets USM Anywhere discovered. It also prompts you to decide if you want to scan for hosts and services running in your environment.

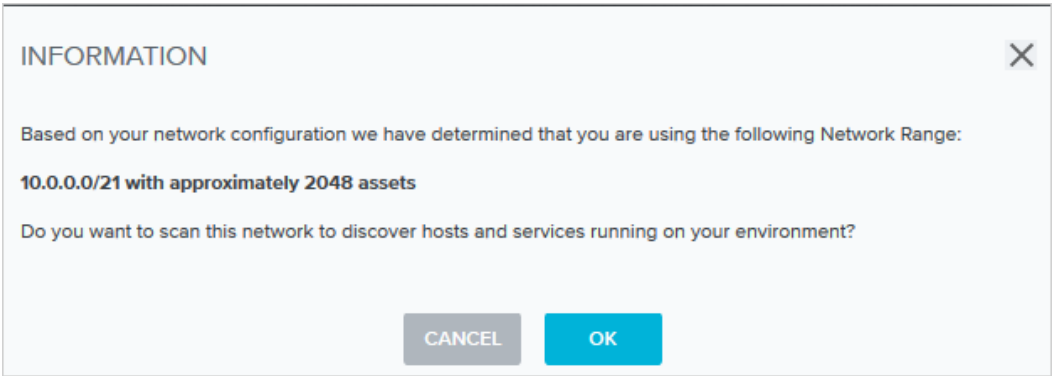

   Click **Cancel** to opt out of this scan.

4. (Optional.) If you want to scan for other hosts and services, click **OK**.

5. Click **Next** after the scan ends.

   The wizard opens the next page in the setup process, Network Security Monitoring.

Network Security Monitoring

The Network Security Monitoring page shows the status of the network interfaces monitored by the sensor (it could take a few moments to load the interfaces). All network adapters are configured for network monitoring by default.

You must manually enable port mirroring or port spanning, promiscuous mode, or both in a virtual switch to send a copy of the network traffic you want to analyze to these interfaces. This page provides links to documentation about how to configure your networking to allow for the interfaces to see the network traffic and perform network intrusion detection.

> **Note:** You must have already configured your network interfaces for VMware. See Set Up USM Anywhere on the VMware Virtual Machine for more information.

USM Anywhere connectivity and communications are handled by the first network interface connection on the Network Security Monitoring page. This is the primary network that provides asset scanning and log collection for the particular network.

You can connect additional interfaces to other networks for monitoring, or connect them to individual vSwitch port groups for virtual networks. Each interface should be connected to a vSwitch that mirrors a different subnet within your network.

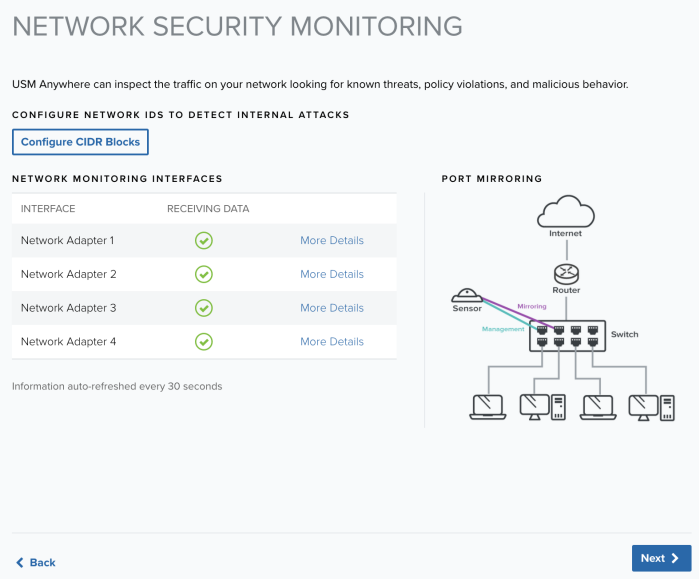

Use this page to verify that USM Anywhere can monitor your network traffic for security events.

Note: You can see red X icons next to the interfaces if the port mirroring or promiscuous mode is not configured. You might also see these icons if the network interfaces have not seen any traffic in the past 30 seconds.

Log Management

On the Log Management page are syslog port numbers. (These ports are the same for all USM Anywhere Sensors.)

USM Anywhere collects third-party device, system, and application data through syslog over UDP on port 514 and over TCP on ports 601 or 602 by default. It collects Transport Layer Security (TLS)-encrypted data through TCP on ports 6514 or 6515 by default. These ports support the RFC 3164 and RFC 5424 formats. To configure any third-party devices to send data to USM Anywhere, you must provide the IP address and the port number of your USM Anywhere Sensor.

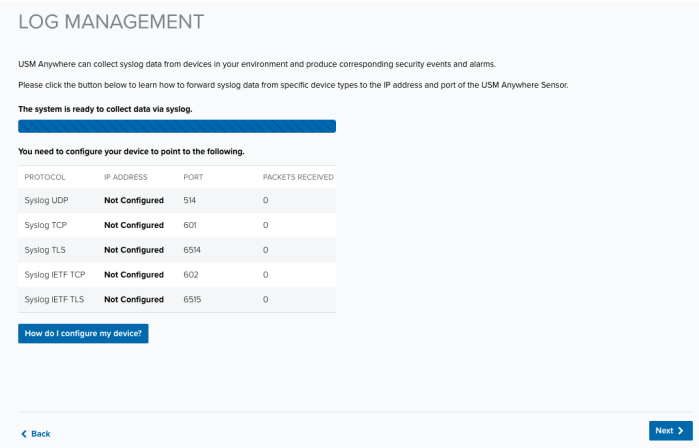

To enable log collection and configure your log management

1. Make sure that you have granted the necessary permissions for your OS to allow USM Anywhere to access its logs. You can also integrate a wide variety of data sources to send log data over syslog to the USM Anywhere Sensor.

   To learn how to configure your operating systems and supported third-party devices to forward syslog log data, see the following related topics:

   - **The Syslog Server Sensor App**: Log collection (UDP, TCP, and TLS-encrypted TCP) from rsyslog
   - **Collecting Linux System Logs**: Log collection from a Linux system
   - Collecting Windows System Logs: Log collection from a Windows system

- Go to the specific AlienApp in USM Anywhere for instructions about syslog forwarding

> **Note:** Because the log scan can take some time, you might not see all of the automatically discovered log sources immediately after deploying the first USM Anywhere Sensor.

2. When you have finished the log collection setup and integrated any needed plugins, verify that the data transfer is occurring.
3. Click **Next** when this step is complete.

OTX

AT&T Alien Labs™ Open Threat Exchange® (OTX™) is an open information-sharing and analysis network providing users with the ability to collaborate, research, and receive alerts on emerging threats and indicators of compromise (IoCs) such as IP addresses, file hashes, and domains.

You must have an OTX account to receive alerts based on threats identified in OTX. This account is separate from your USM Anywhere account. Go to The World's First Truly Open Threat Intelligence Community to create an OTX account.

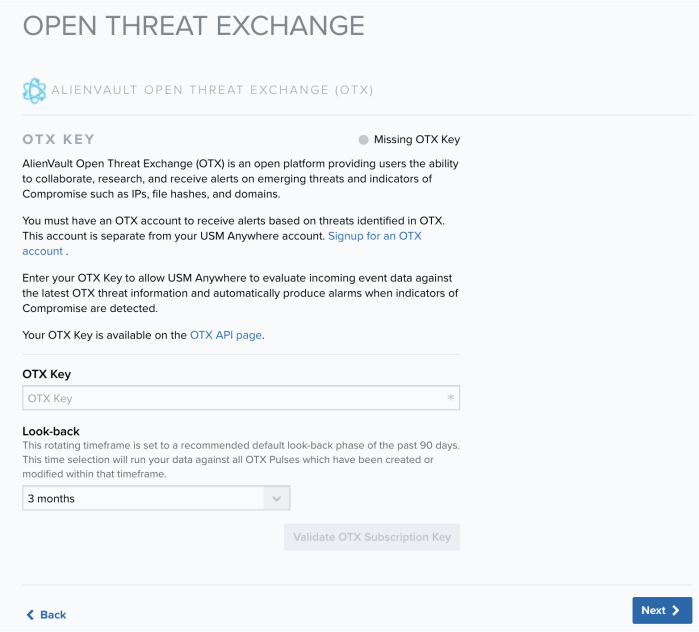

Note: If you do not already have an OTX account, click the **Sign up** link. This opens another browser tab or window that displays the OTX signup page. After you confirm your email address, you can log in to OTX and retrieve the unique API key for your account.

See Open Threat Exchange® and USM Anywhere in the *USM Anywhere User Guide* for more information about OTX integration in USM Anywhere.

To enable USM Anywhere to evaluate event data against the latest OTX intelligence

1. Log in to OTX and open the API page (https://otx.alienvault.com/api).
2. In the DirectConnect API Usage pane, click the icon to copy your unique OTX connection key.

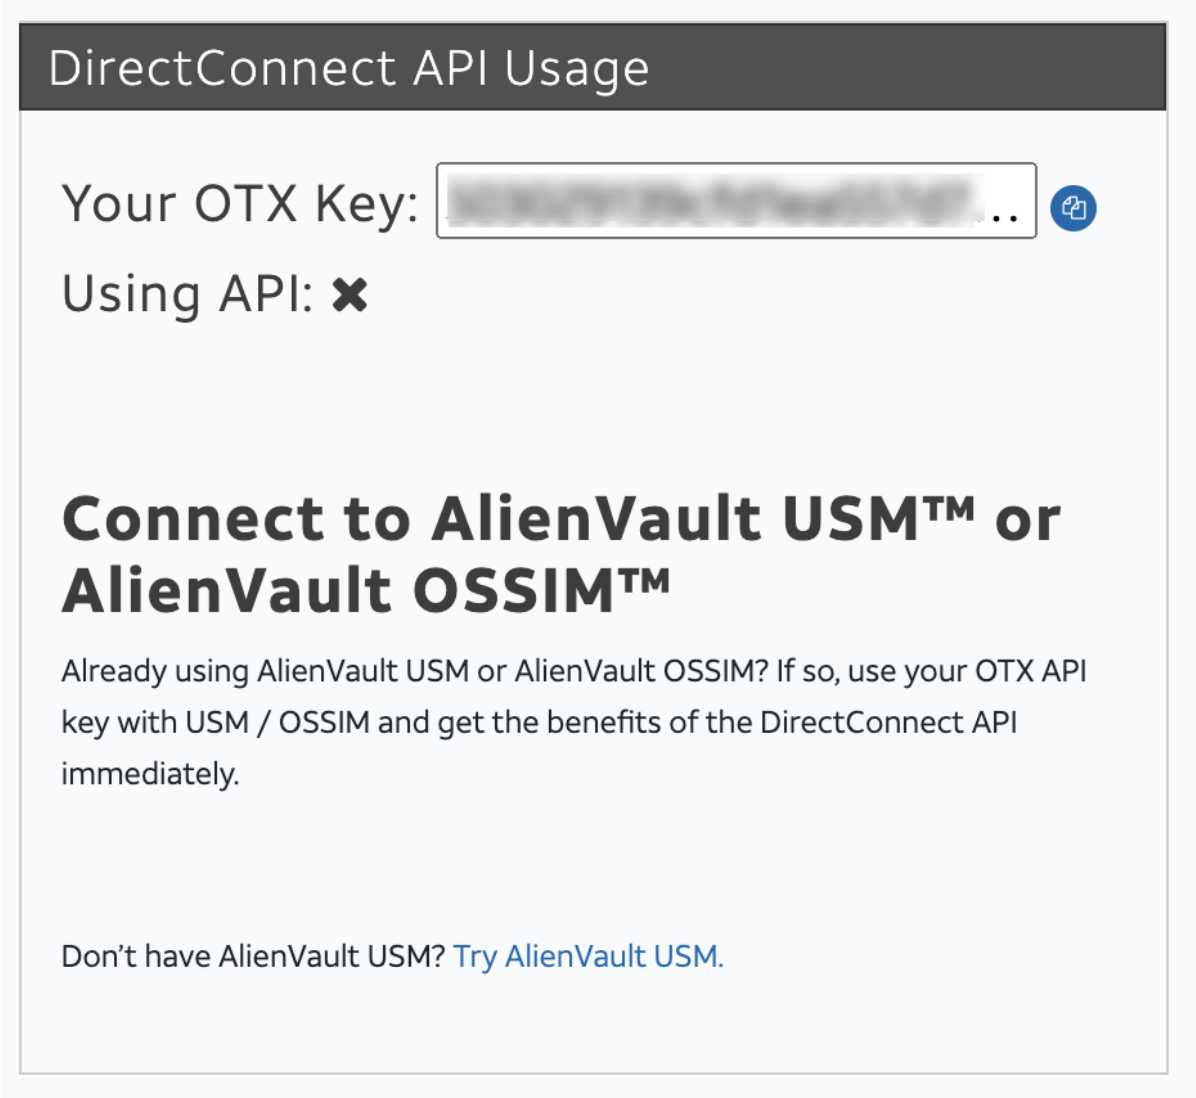

3. Return to the Open Threat Exchange (OTX) page of the USM Anywhere Sensor Setup Wizard and paste the value in the OTX Key text box.

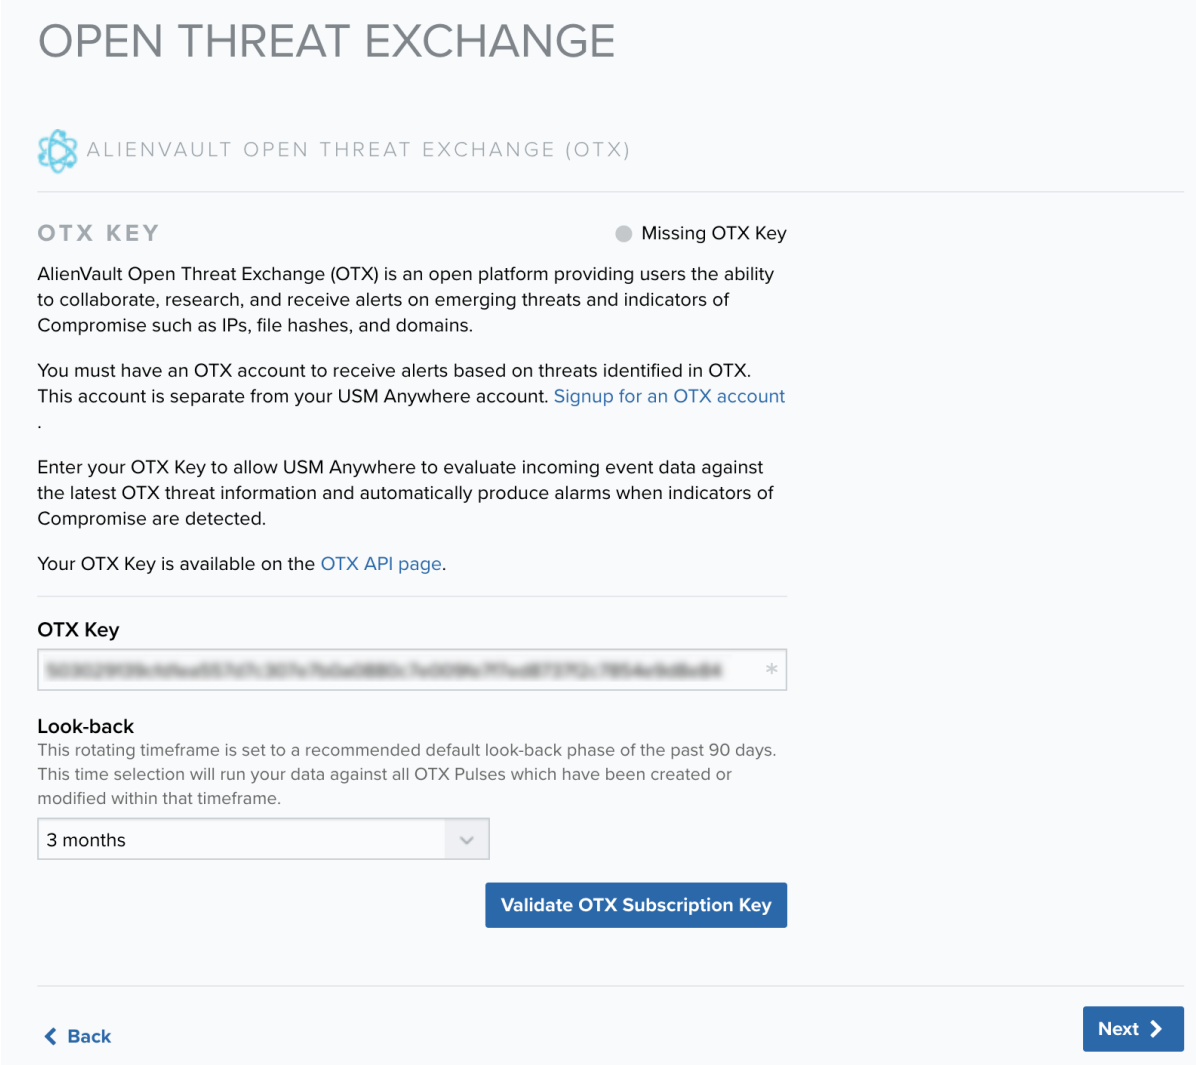

4. Click **Validate OTX Subscription Key**.

   With a successful validation of the key, the status at the top of the page changes to "Valid OTX key".

5. Click **Next** when this task is complete.

Setup Complete

The Congratulations page summarizes the status of your configuration.

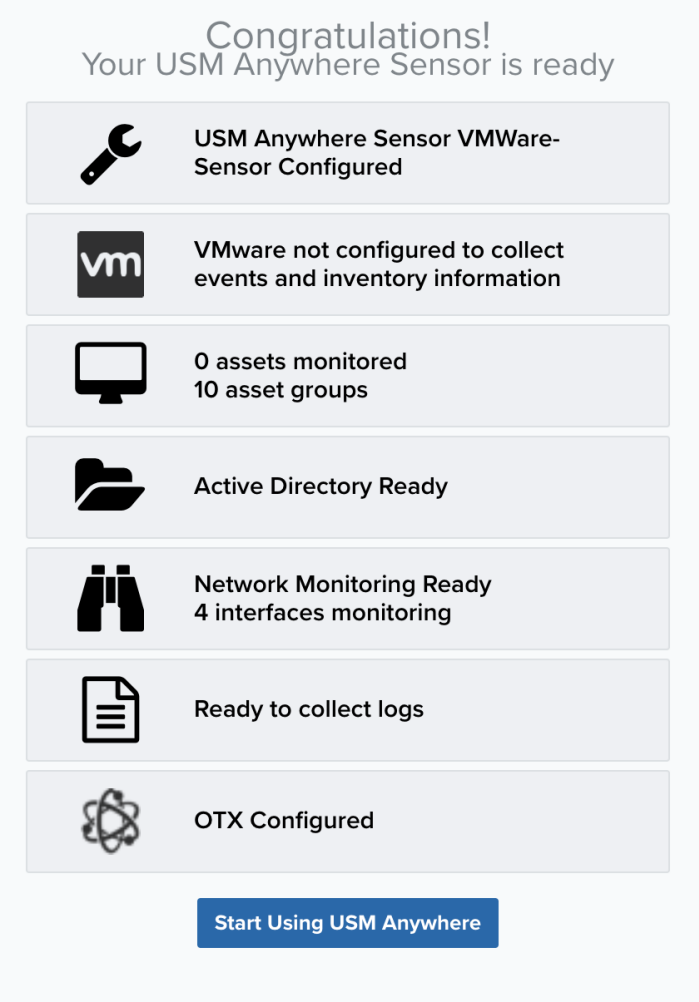

Click **Start Using USM Anywhere**, which takes you to the Overview dashboard.

> **Note:** You can disable the VMware Discover Virtual Machine Scheduled Task and can still get events without any issues because this task is only to discover the virtual machines.

Next...

Now is a great time to run a vulnerability scan. See Vulnerability Assessment in the *USM Anywhere User Guide* for detailed information about running a vulnerability scan.

USM Anywhere Sensor Deployment on Microsoft Hyper-V

Microsoft Hyper-V is a hypervisor that lets you create and manage a virtualized computing environment by using virtualization technology that is built into Microsoft Windows Server. Through Hyper-V, you can deploy a USM Anywhere Sensor in any of the virtual networks that you want to instrument for threat monitoring.

The Hyper-V Sensor deployment includes a network-based intrusion detection system (NIDS) that monitors the networks connected to the listening interfaces. A deployed Hyper-V Sensor supports a NIDS throughput of 600 Mbps, but this performance may vary depending on your environment, configurations, and other variables.

This section includes the following topics.

About Hyper-V Sensor Deployment 70

Requirements for Hyper-V Sensor Deployment 70

Create the Hyper-V Virtual Machine 80

Set Up USM Anywhere on the Hyper-V Virtual Machine 89

Connect the Hyper-V Sensor to USM Anywhere 92

Complete the Hyper-V Sensor Setup 97

About Hyper-V Sensor Deployment

If your organization uses multiple subnets to allow communication between headquarters and remote offices, you do not need a sensor for each subnet. However, you will need a Hyper-V Sensor for each physical location that you want to monitor.

Note: See the Microsoft TechNet website (https://technet.microsoft.com/en-us/library/mt169373(v=ws.11.aspx) for more information about Hyper-V.

Deployment Process Overview

The deployment process for an initial USM Anywhere Sensor on Hyper-V consists of these primary tasks:

1. Review requirements for a Hyper-V Sensor deployment
2. Create, configure, and start the Hyper-V virtual machine
3. Configure the sensor on the virtual machine
4. Register the new sensor with your sensor authentication code to provision the USM Anywhere instance and connect the deployed sensor
5. Complete your Hyper-V Sensor configuration, including initial asset discovery

Requirements for Hyper-V Sensor Deployment

Review the following prerequisites to ensure an efficient setup and configuration of a USM Anywhere Sensor on Microsoft Hyper-V.

Minimum Requirements

These are the minimum requirements needed to set up and configure a USM Anywhere Sensor on Hyper-V:

- Operating system (OS) must be Windows Server 2012 R2, 2016, 2019, or 2022
- Using Hyper-V Manager or Microsoft System Center Virtual Machine Manager (SCVMM)
- Dedicated 178 GB of disk space (128 GB data device and 50 GB root device)
- Dedicated 4 CPUs and 12 GB of statically assigned memory
- Internet connectivity from the virtual machine

Important: Because the needs of a sensor differ based on the varying demands of different deployment environments and the complexity of events being processed, the number of events per second (EPS) a sensor can process varies.

Depending on your environment, you may need to deploy additional sensors to ensure that all events are processed.

Recommended Requirements

These are the recommended requirements needed to set up and configure a USM Anywhere Sensor on Hyper-V:

- If Dynamic Host Configuration Protocol (DHCP) is unavailable, a static IP for the management interface and local Domain Name System (DNS) information

  Important: AT&T Cybersecurity strongly recommends assigning a static IP address to deploy the USM Anywhere Sensor.

  If DHCP changes the IP address of the sensor, you must update all the IP addresses on all the devices that are forwarding logs to the sensor through syslog.

- Network topology information to run asset discovery
- Port mirroring setup for network monitoring (see Direct Traffic from Your Physical Network to the Hyper-V Sensor for more information)
- Administrative credentials for remote hosts to support authenticated asset scans
- Administrative credentials for devices that require configuration to forward logs to the Hyper-V Sensor
- To access network-based intrusion detection system (NIDS) functionality on the sensor, an ethernet port on the host must be available to receive data from a Switched Port Analyzer (SPAN) or Test Access Point (TAP) port

Sensor Ports and Connectivity

Before deploying a USM Anywhere Sensor, you must configure your firewall permissions to enable the required connectivity for the new sensor. Initial deployment of a sensor requires that you open egress and outbound ports and protocols in the firewall for communication with USM Anywhere and AT&T Cybersecurity Secure Cloud resources. The sensor receives no inbound connections from outside the firewall.

Note: To launch the USM Anywhere Sensor web UI during the initial setup, you need to allow inbound traffic to the sensor IP address through TCP port 80. You can remove access to this port after the sensor successfully connects to USM Anywhere. You do not need to allow inbound traffic to this port from the Internet.

The following tables list the inbound and outbound ports.

Sensor Ports and Connectivity (Outbound Ports)

Type	Ports	Endpoints	Purpose
TCP	443	update.alienvault.cloud	Communication with AT&T Cybersecurity for initial setup and future updates of the sensor.
TCP	443	reputation.alienvault.com	Ongoing communication with AT&T Alien Labs™ Open Threat Exchange® (OTX™).
TCP	443	otx.alienvault.com	Ongoing communication with OTX to retrieve vulnerability scores. Connecting to otx.alienvault.com is not required but highly recommended. OTX uses the AWS CloudFront services. Refer to the AWS IP address ranges page when you deploy a new sensor. This page contains the current IP address ranges for the service and instructions on how to filter the addresses.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
SSL	443	storage-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send and retrieve backups.
SSL	443	metrics-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send metrics and messages.
SSL/TCP	443	api-parameters-<REGION>-prod.alienvault.cloud[2] api-message-proxy-<REGION>-prod.alienvault.cloud api.message-proxy.<REGION>.prod.alienvault.cloud	Ongoing communication with USM Anywhere. It is only necessary to allowlist the address that corresponds to the region where your USM Anywhere instance is hosted.
SSL/TCP	7100	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
UDP	53	DNS Servers (Google Default)	Ongoing communication with USM Anywhere.
UDP	123	0.pool.ntp.org 1.pool.ntp.org 2.pool.ntp.org 3.pool.ntp.org	Sync with network time protocol (NTP) services.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	22 and 443	prod-usm-saas-tractorbeam.alienvault.cloud prod-gov-usm-saas-tractorbeam.gov.alienvault.us (for AT&T TDR for Gov)	SSH communications with the USM Anywhere remote support server. See Troubleshooting and Remote Sensor Support for more information about remote technical support through the USM Anywhere Sensor console.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	geoip-ap-northeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-south-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-2-prod.alienvault.cloud/geo-ip/sensor geoip-ca-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-me-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-sa-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-us-gov-west-1-prod-gov.gov.alienvault.us/geo-ip/sensor (for AT&T TDR for Gov)	Allows resolution of IP addresses for geolocation services. It is only necessary to allowlist the GeoIP address that corresponds to the region where your USMA instance is hosted.

1

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

2

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

3

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

4

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

5

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

Sensor Ports and Connectivity (Inbound Ports)

Type	Ports	Purpose
SSH	22	Inbound method for secure remote login from a computer to USM Anywhere.
HTTP	80	Inbound communication for HTTP traffic.
UDP (RFC 3164)	514	USM Anywhere collects data through syslog over UDP on port 514 by default.
TCP (RFC 3164)	601	Inbound communication for reliable syslog service. USM Anywhere collects data through syslog over TCP on port 601 by default.
TCP (RFC 5424)	602	USM Anywhere collects data through syslog over TCP on port 602 by default.
Traffic Mirroring	4789	Inbound communication for virtual extensible local area network (VXLAN).

Sensor Ports and Connectivity (Inbound Ports) (Continued)

Type	Ports	Purpose
WSMANS	5987	Inbound WBEM WS-Management HTTP over Secure Sockets Layer/Transport Layer Security (SSL/TLS) (NXLog).
TLS/TCP (RFC 3164)	6514	USM Anywhere collects TLS-encrypted data through syslog over TCP on port 6514 by default.
TLS (RFC 5424)	6515	USM Anywhere collects data through syslog over TLS on port 6515 by default.
TCP	9000	Inbound communication used internally for HTTP sensor traffic.
Graylog	12201	Inbound communication for Graylog Extended Log Format (GELF).

USM Anywhere IP Addresses for Allowlisting

Your sensor is connected to a USM Anywhere instance deployed in one of the Amazon Web Services (AWS) endpoint regions based on your location. If you need to configure your firewall to allow communication between the sensor and the USM Anywhere instance, refer to the following table with the reserved IP address ranges for each region.

Important: The Update Server and the AlienVault Agent always use the 3.235.189.112/28 range no matter where your USM Anywhere is deployed. The AT&T TDR for Gov Update Server uses the 3.32.190.224/28 range.

Note: The regional IP ranges listed in this table are limited to the control nodes (subdomain). You must also meet all requirements provided in the Sensor Ports and Connectivity (Outbound Ports) table.

AWS Regions Where USM Anywhere Instance Is Available

Code	Name	Reserved Static IP Address Ranges
ap-northeast-1	Asia Pacific (Tokyo)	18.177.156.144/28 3.235.189.112/28 44.210.246.48/28
ap-south-1	Asia Pacific (Mumbai)	3.7.161.32/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-1	Asia Pacific (Singapore)	18.143.203.80/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-2	Asia Pacific (Sydney)	3.25.47.48/28 3.235.189.112/28 44.210.246.48/28
ca-central-1	Canada (Central)	3.96.2.80/28 3.235.189.112/28 44.210.246.48/28
eu-central-1	Europe (Frankfurt)	18.156.18.32/28 3.235.189.112/28 44.210.246.48/28
eu-west-1	Europe (Ireland)	3.250.207.0/28 3.235.189.112/28 44.210.246.48/28

AWS Regions Where USM Anywhere Instance Is Available (Continued)

Code	Name	Reserved Static IP Address Ranges
eu-west-2	Europe (London)	18.130.91.160/28 3.235.189.112/28 44.210.246.48/28
me-central-1	Middle East (UAE)	3.29.147.0/28 3.235.189.112/28 44.210.246.48/28
sa-east-1	South America (São Paulo)	18.230.160.128/28 3.235.189.112/28 44.210.246.48/28
us-east-1	US East (N. Virginia)	3.235.189.112/28 44.210.246.48/28
us-west-2	US West (Oregon)	44.234.73.192/28 3.235.189.112/28 44.210.246.48/28
us-gov-west-1	AWS GovCloud (US-West)	3.32.190.224/28

Hyper-V Machine Deployment

You can deploy a Hyper-V virtual machine (VM) using either of the following management tools:

- **Microsoft Hyper-V Manager**, which is an administrative tool for managing local and remote Hyper-V servers. See Create the VM with Hyper-V Manager for more information.
- **System Center Virtual Machine Manager 2012**, which is designed for managing large numbers of virtual servers, based on Microsoft Virtual Server and Hyper-V. See Create the VM with SCVMM 2012 for more information.

Create the Hyper-V Virtual Machine

The first tasks in deploying the Microsoft Hyper-V Sensor are to create, configure, and start the virtual machine (VM) using either the Microsoft System Center Virtual Machine Manager (SCVMM) 2012 or Microsoft Hyper-V Manager. Before you begin these tasks, download the Hyper-V Sensor package from AT&T Cybersecurity.

To download the Hyper-V package

1. Go to the USM Anywhere Sensor Downloads page and click the icon of your specific sensor. After clicking, your browser starts to download the USM Anywhere Sensor package. Depending on your Internet connection, the download can take 30 minutes or more.
2. Extract the USM Anywhere Sensor package and place the files where they are accessible from your Hyper-V management tool.

Create the VM with SCVMM 2012

Follow these procedures to create and start the VM using SCVMM 2012.

Creating the VM

To create a VM with SCVMM 2012

1. Connect to the SCVMM 2012.
2. Select **Library > VM Templates**.
3. Select the AlienVault VM template under VM Templates, and then click **Create Virtual Machine**.

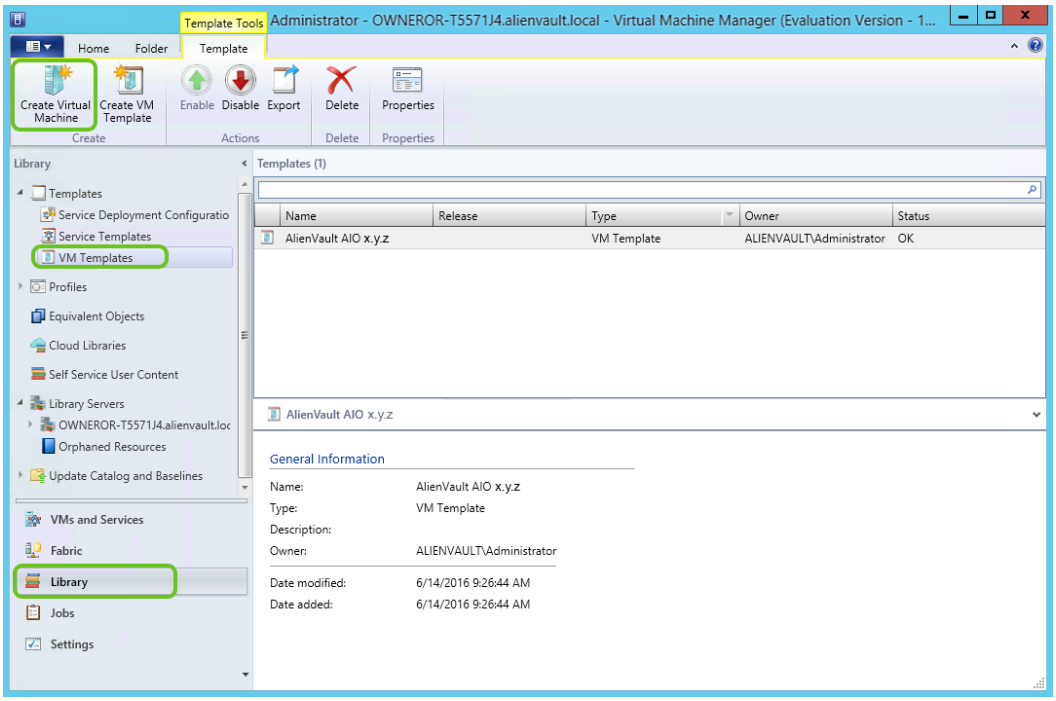

4. Enter a name for the VM, and then click **Next**.
5. In Configure Hardware, click **Next**.
6. In Select Destination, choose whether to deploy or store the VM, and then click **Next**.
7. In Select Host, select a destination for the VM, and then click **Next**.
8. In Configure Settings, review the VM settings, and then click **Next**.
9. In Add Properties, change the automatic actions if necessary, and then click **Next**.
10. In Summary, confirm the settings and then click **Next**.

    The Jobs window opens and shows the VM being created.

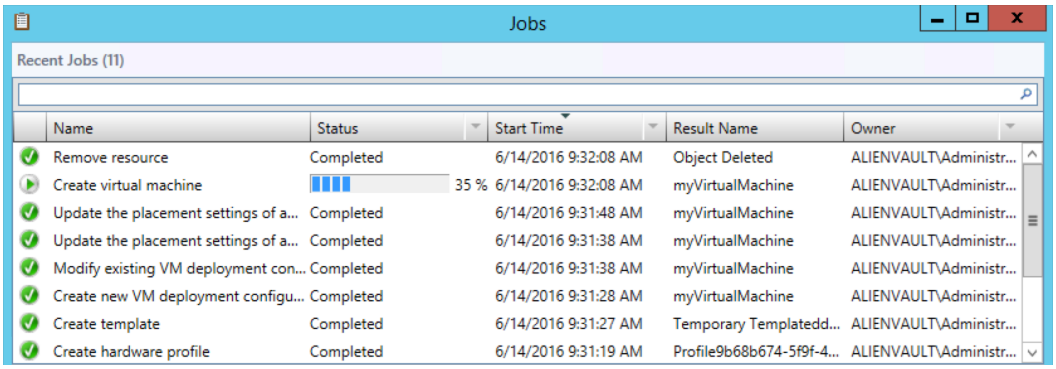

After the VM is created, you can close the window.

Starting the VM

To start the VM and connect

1. Right-click the VM and select **Power On**.

   This process could take a few minutes. The Virtual Machine State will display as Running when this is complete.

2. Right-click the VM and select **Connect or View > Connect via Console**.

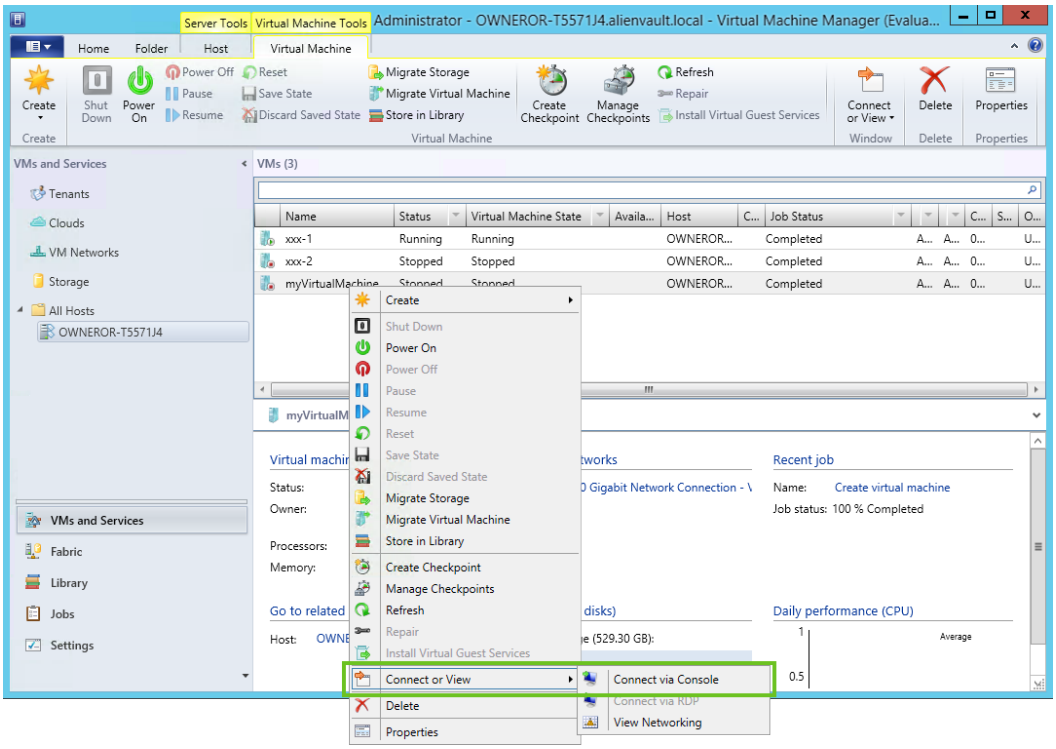

   Use the information in the Welcome screen to log in to the sensor. This screen also displays the URL you can use to access USM Anywhere and register the sensor.

Create the VM with Hyper-V Manager

Follow these procedures to create and start the VM using Hyper-V Manager.

Creating the VM

To create a new VM with Hyper-V Manager

1. Open the Hyper-V Manager and connect to the server.
2. From the Actions panel, go to **New > Virtual Machine**.
   The New Virtual Machine Wizard launches.

3. Click **Next**.
4. In Specify Name and Location, enter a name for the new VM, and then click **Next**.

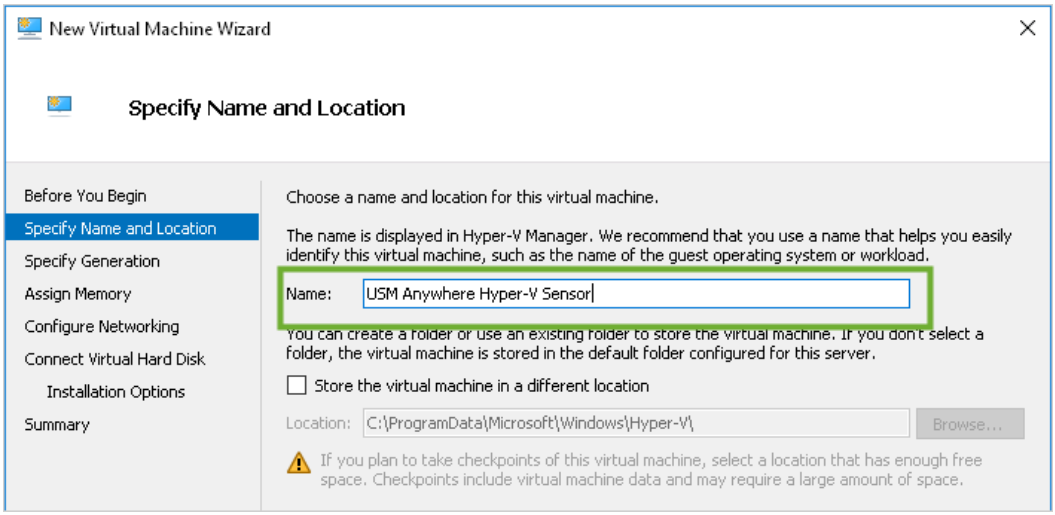

5. In Specify Generation, choose **Generation 1** for the virtual machine, and then click **Next**.
6. In Assign Memory, change the value of the *Startup memory* to **12288 MB**.

> **Important:** Make sure that the Use Dynamic Memory option is *not* selected.

Click **Next** when complete.

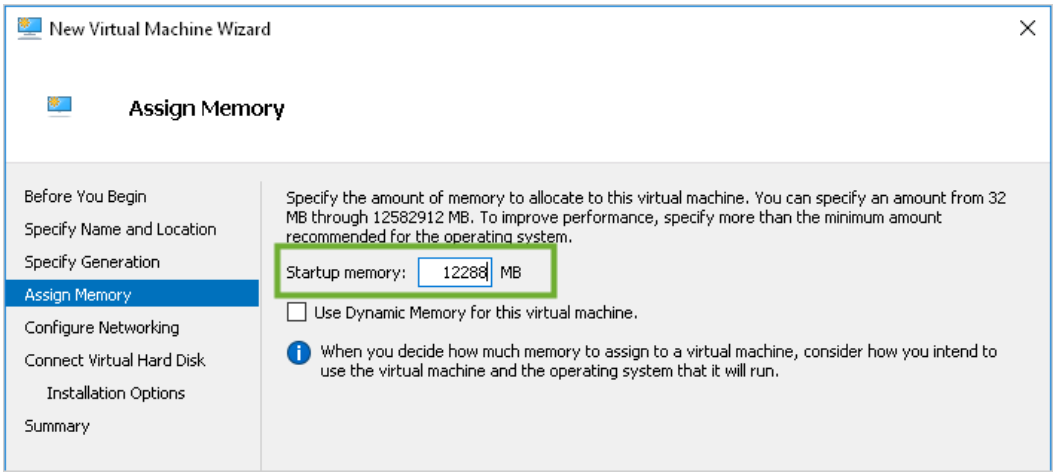

7. In Configure Networking, connect the new VM to the desired network, and then click **Next**.

8. In Connect Virtual Hard Disk, select **Use an existing virtual hard disk**, and then click **Browse** to locate the `usm-os-disk.vhd` file that was extracted from the USM Anywhere Sensor package download.

> **Note:** You will add the data disk later in another step because the wizard doesn't support this.

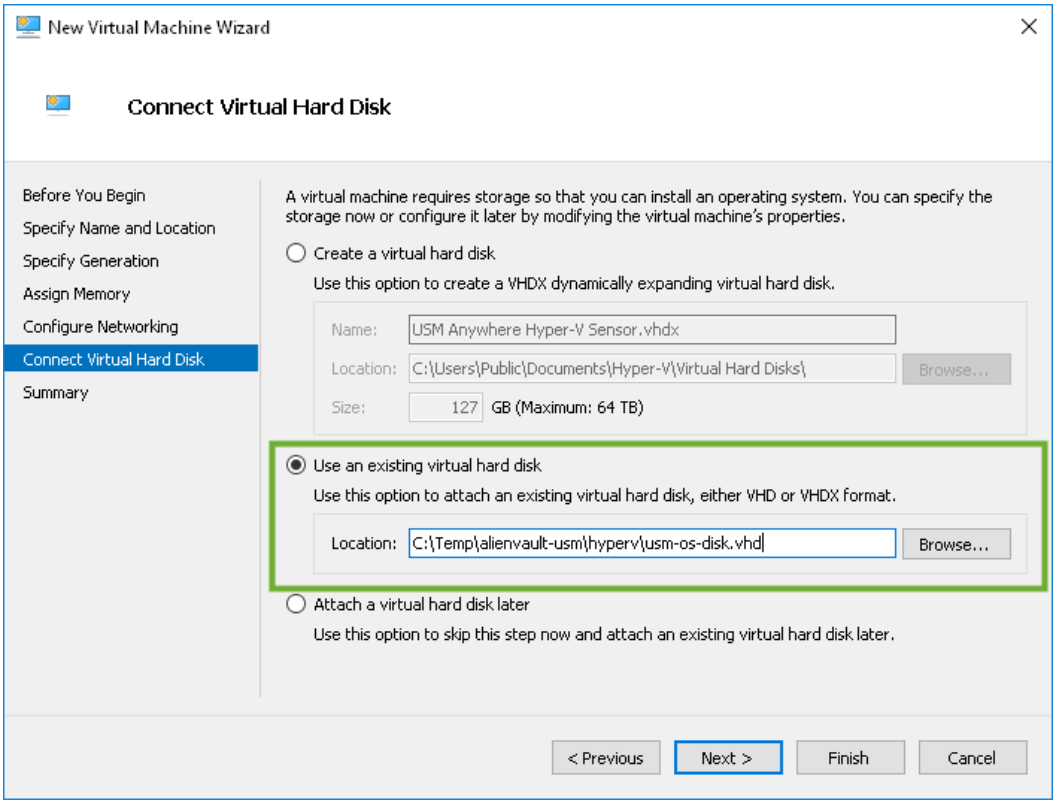

9. In Complete the New Virtual Machine Wizard, click **Next** and then **Finish**.

Configuring the VM

To configure a VM using the Hyper-V Manager

1. Select the VM you previously created, and then select **Action > Settings**.
2. In the left navigation menu of the dialog box, select **Processor**, and then set the number of virtual processors to **4**.

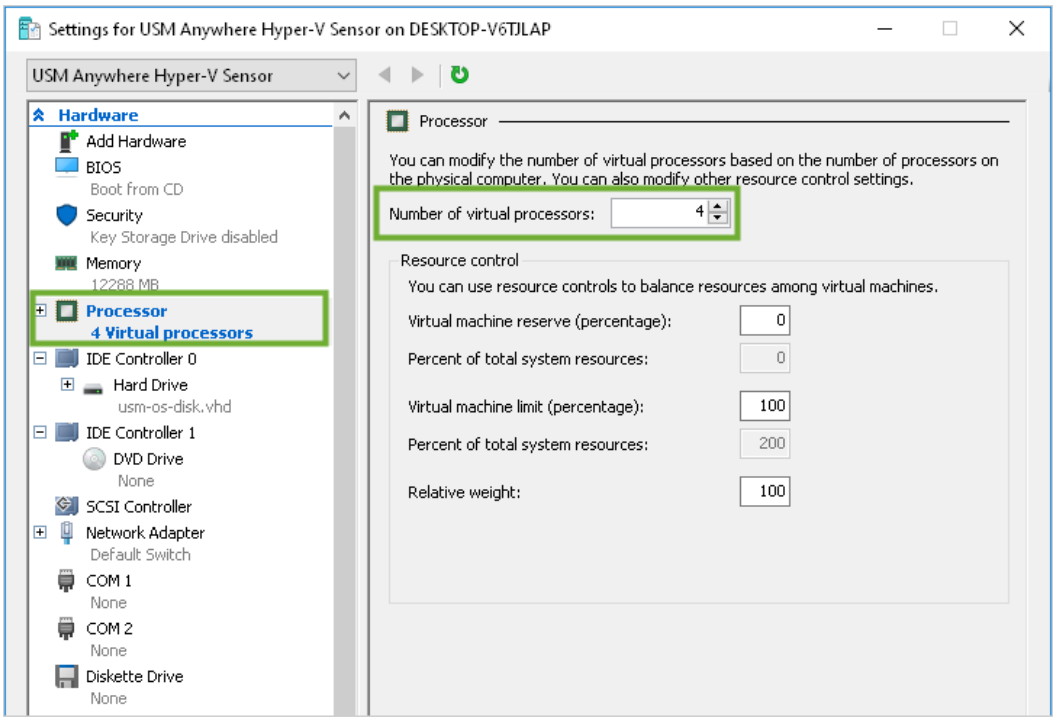

3. Click **Apply**.
4. In the left navigation menu, select **IDE Controller 0**.
5. Select **Hard Drive**, click **Add**, and then click **Browse** to locate the `usm-data-disk.vhdx` file that was part of the USM Anywhere Sensor download.

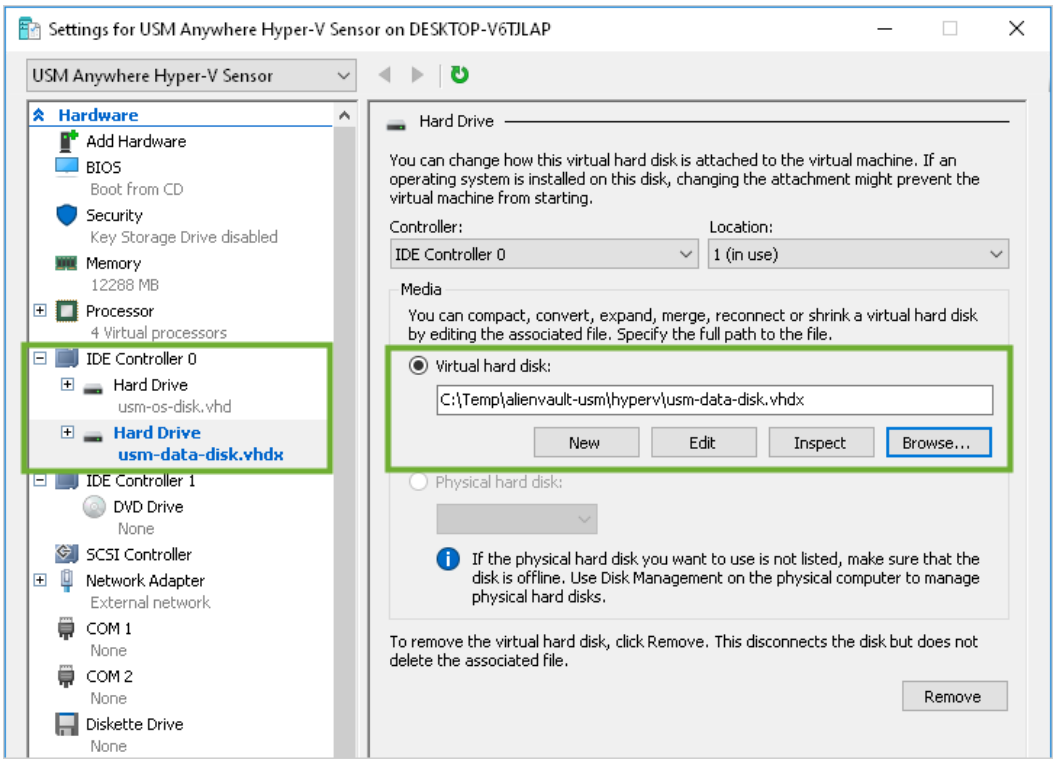

6. In the top-left corner of the dialog box, click **Add Hardware**, and then select **Network Adapter > Add** to add additional network adapters.

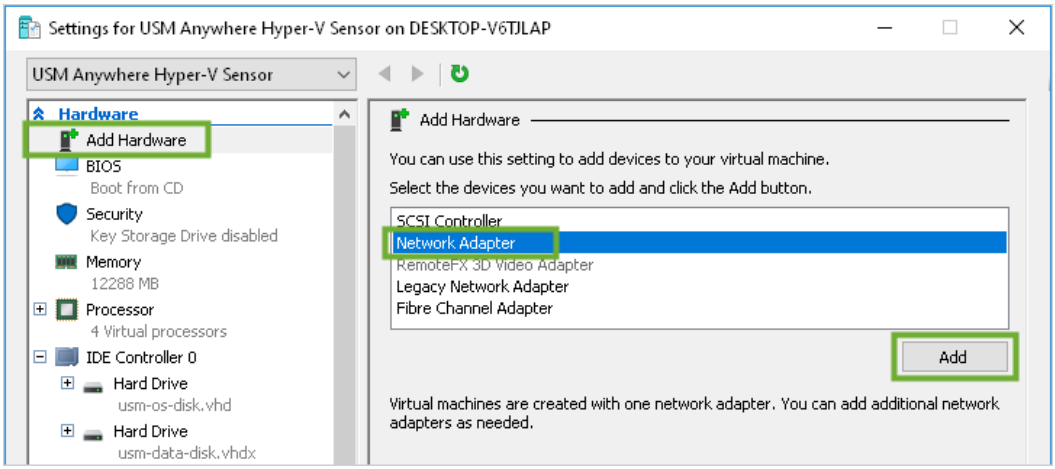

Repeat this function so that you have a total of five network adapters.

Warning: The Hyper-V Sensor requires *all five network interface cards (NICs)* to be enabled; otherwise, the USM Anywhere update will fail. The NICs can remain disconnected.

See Configure Network Interfaces for On-Premises Sensors for more information about these interfaces.

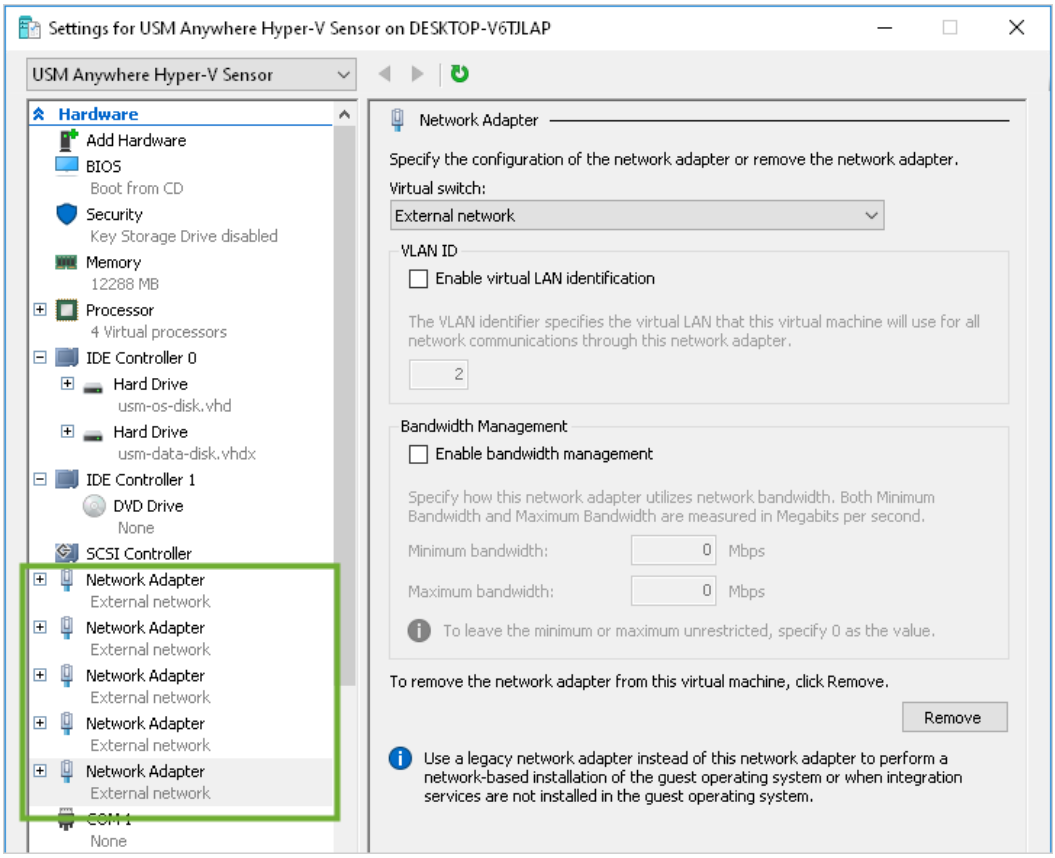

7. Click **OK**.

8. Configure the ISO file for the datastore.

> **Warning:** You must complete this step and ensure that the ISO is mounted before you start the sensor VM for the first time.
>
> If you see `REPLACEME` as the initial login password in the sensor welcome screen when you connect to the VM, it is most likely that the ISO was not mounted before the sensor was started. If this happens, you must shut down the VM, complete this step so that the ISO is configured for the datastore, and then begin the deployment process anew.

- Select **IDE Controller 1 > DVD Drive**.
- Select **Image file**, and then click **Browse** to locate the `deploy_config.iso` image file that was extracted from the USM Anywhere Sensor package download.
- Click **Apply**.

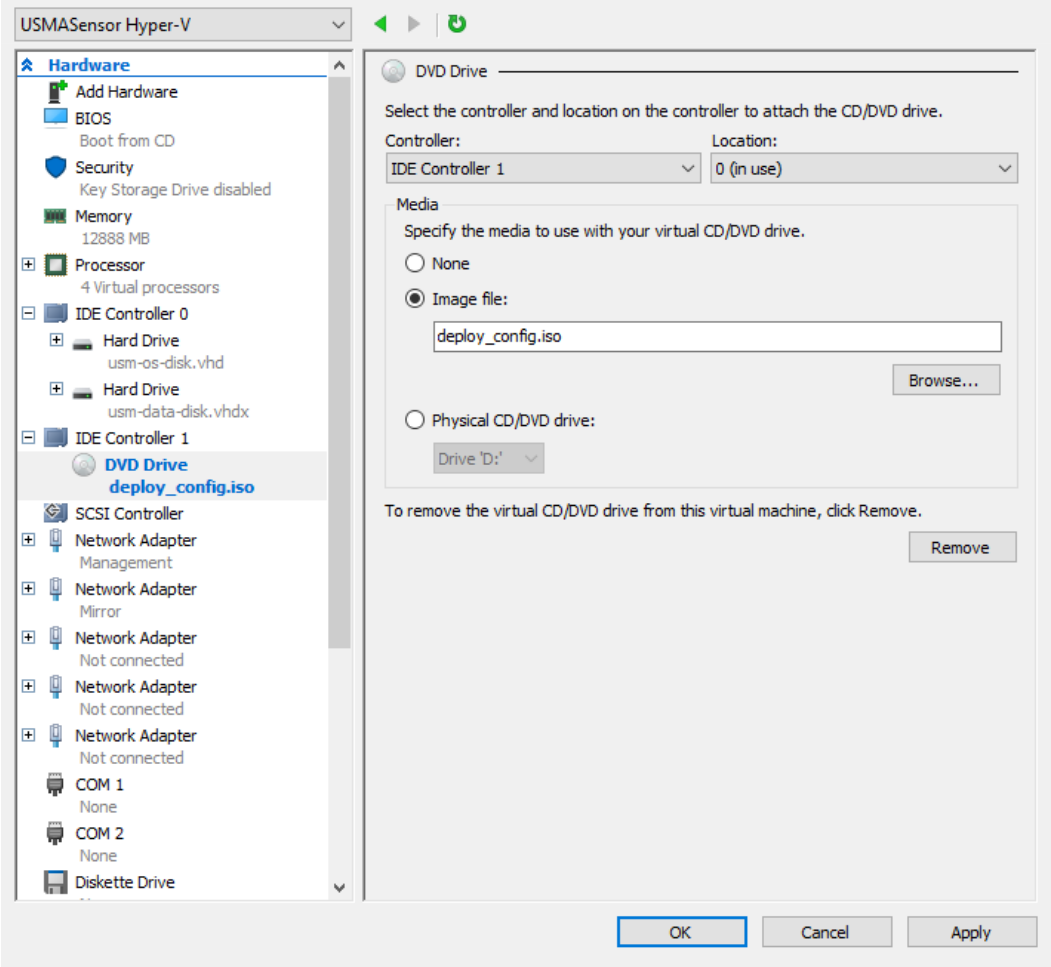

9. Click **OK** and close the dialog box.

Starting the VM

To start the VM and connect

1. Select the VM and select **Action > Start**.

   This process can take a few minutes. The status displays as Running when this is complete.

2. Select your VM and select **Action > Connect**.

   Use the information in the Welcome screen to log in to the sensor. This screen also displays the URL you can use to access USM Anywhere and register the sensor.

Set Up USM Anywhere on the Hyper-V Virtual Machine

Role Availability	Read-Only	Investigator	Analyst	Manager

There is some configuration required within the console of the sensor. After this initial configuration, you use the USM Anywhere web UI to further configure the sensor, and all other sensors that you connect to USM Anywhere.

Perform these initial configuration tasks on the Hyper-V virtual machine, using the USM Anywhere Sensor console.

Change the Administrative Password and Keyboard Layout

Follow these instructions to change the administrative password and keyboard layout.

To change the administrative password and keyboard layout

1. Log in using the credentials displayed in the console screen.

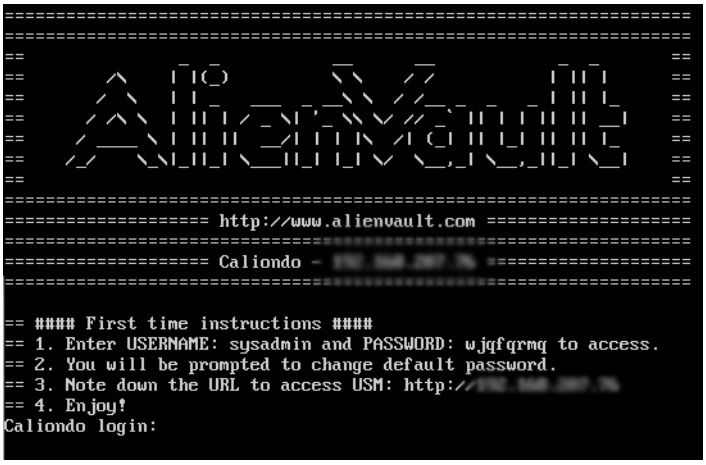

2. (Optional.) Configure the keyboard if you use a keyboard layout other than the U.S. default.

3. Set a new password for the *sysadmin* user.

> **Important:** During the installation, your system acquires the initial IP address through Dynamic Host Configuration Protocol (DHCP). If DHCP is not enabled, you must configure it manually.
>
> AT&T Cybersecurity *strongly* recommends assigning a static IP address to the USM Anywhere Sensor as a best practice. This allows for proper log forwarding and network architecture.

- If your system sets an IP address automatically, note the web URL (IP address). You will need the URL when you exit from the console and follow the instructions in Connect the Hyper-V Sensor to USM Anywhere .

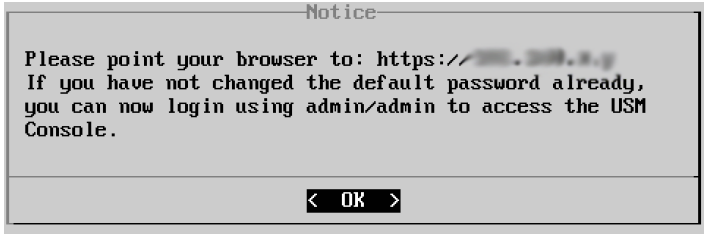

- If your system does not set an IP address automatically, a message box confirms that the system was unable to acquire an IP address from a DHCP server after you change the sysadmin password.
- In this case, you must manually set a static IP address so that it remains unchanged in the future.

Configure a Static IP Address

Follow these instructions to configure a static IP address.

To configure a static IP address

1. Go to **Network Configuration > Configure Management Interface > Set a Static Management IP Address**.
2. Enter the IP address, subnet, and gateway information in each screen.
3. Press **Enter**.

Important: DNS settings are not maintained when a static IP address is configured. If you configure a static IP address, you must configure the DNS network settings for successful sensor activation.

Configure Domain Name System

Follow these instructions to configure the Domain Name System (DNS).

Important: When the USM Anywhere Sensor performs an asset scan, it must access the local Domain Name System (DNS) server to resolve local host names. The sensor uses reverse DNS to look up the hostname through the discovered IP address.

Note: When deploying your VMware Sensor in a DHCP environment, the DNS server is automatically set to retrieve via DHCP. This can be configured later in your sensor's settings. See Deploying Your Sensor in a DHCP Environment for more information about USM Anywhere Sensors in a DHCP environment.

To configure DNS

1. Go to **Network Configuration > Configure DNS**.
2. Enter the primary DNS and press **Enter**.

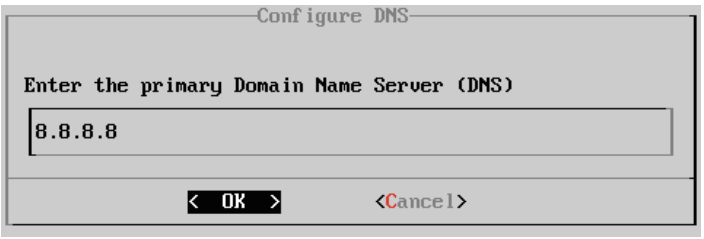

3. (Optional.) Enter the secondary DNS and press **Enter**.

   A text box opens to confirm that you want to apply changes.

4. Press **Enter**.

Note: Check your settings through **Network Configuration > View Network Configuration**.

Connect the Hyper-V Sensor to USM Anywhere

Role Availability ✗ Read-Only ✗ Investigator ✗ Analyst ✓ **Manager**

After deploying the Hyper-V Sensor, you must connect it to USM Anywhere through registration.

Obtain the Authentication Code

You must enter an authentication code when registering the USM Anywhere Sensor. How to obtain the authentication code depends on your USM Anywhere instance and whether this is the first sensor you're deploying.

Instructions for USM Anywhere customers:

If this is your first USM Anywhere Sensor, you must register the sensor using the initial authentication code (starts with a "C") received from AT&T Cybersecurity. With this code, the registration process provisions a new USM Anywhere instance and defines its attributes, such as how many sensors to allow for connection, how much storage to provide, and what email address to use for the initial user account. After registration, you will gain access to the sensor through the USM Anywhere web user interface (UI), where you can complete the sensor setup.

If you are deploying additional sensors, you must generate the authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Instructions for AT&T TDR for Gov customers:

AT&T Cybersecurity has already provisioned the AT&T Threat Detection and Response for Government (AT&T TDR for Gov) instance for you, therefore you won't receive an authentication code for your sensor. This is true regardless if it's the first sensor or additional sensors you're deploying. However, for the first sensor, you'll receive a link to access your instance.

For every sensor you deploy, you must generate an authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Register Your Sensor

You perform this procedure after deploying the USM Anywhere Sensor within your Hyper-V environment. You can acquire the IP address when you set your management interface settings in the sensor console.

To register your sensor

1. Open a web browser and enter the IP address.

   This opens the *Welcome to USM Anywhere Sensor Setup* page, which prompts you to provide the information for registering the sensor with your new USM Anywhere instance.

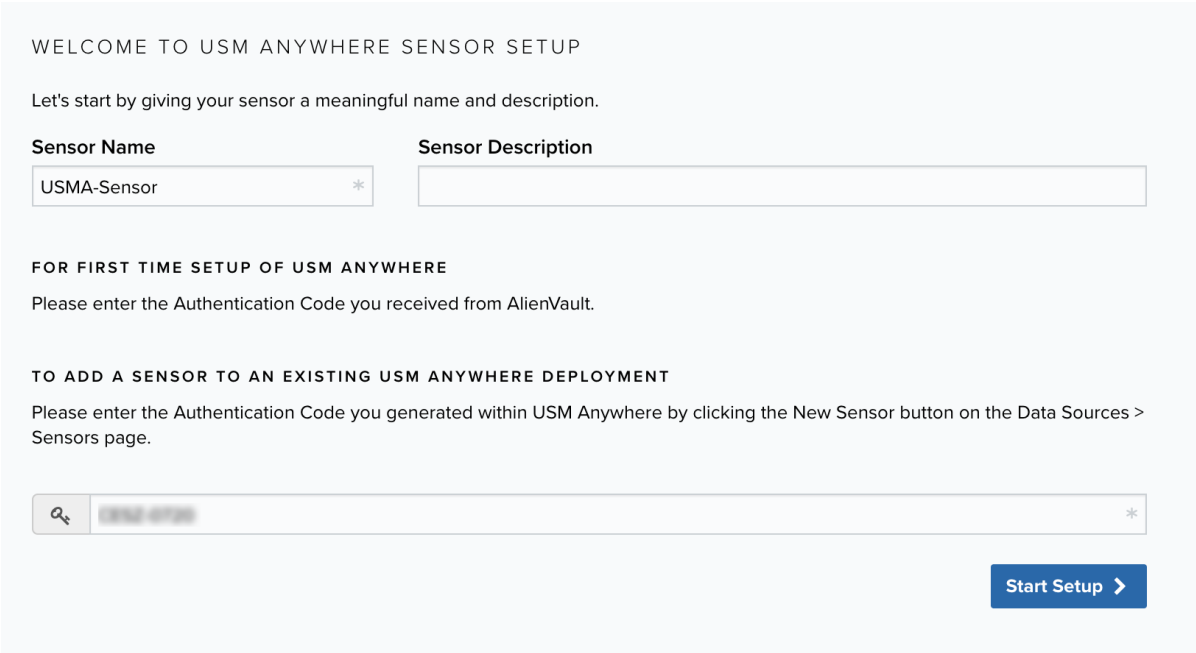

2. Enter a sensor name and sensor description.
3. Paste the authentication code into the field with the key icon ( ).
4. Click **Start Setup** to start the process of connecting the USM Anywhere Sensor.

   It takes about 20 minutes to provision your USM Anywhere instance upon registration of your initial sensor. When this instance is provisioned and running, you'll see a welcome message that provides an access link.

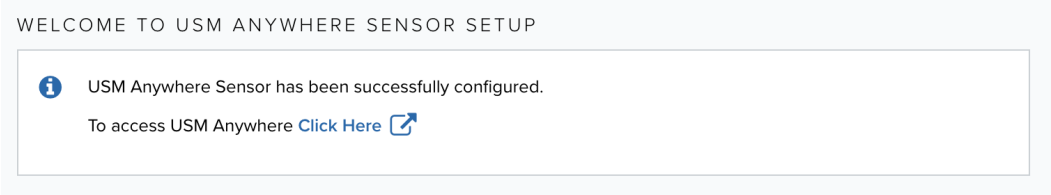

   Use this link to open the secured web console for your USM Anywhere instance. You and the other USM Anywhere users in your organization can access this console from a web

browser on any system with internet connectivity.

> **Note:** If this is your first deployment, you'll also receive an email from AT&T Cybersecurity that provides the access link to USM Anywhere.

Configure the Initial Login Credentials

When you link to a newly provisioned USM Anywhere instance, you must configure the password for the initial user account. This is the default administrator as defined in your subscription.

To configure login credentials

1. In the welcome message, click the link.

   This displays a prompt to set the password to use for the default administrator of USM Anywhere.

2. Enter the password, and then enter it again to confirm.

   Keep in mind these points when you are logging in:

   - The login credentials that you set will apply to any USM Anywhere™ and USM Central™ you have access to.
   - USM Anywhere requires all passwords to have a minimum length of 8 characters and a maximum length of 128 characters.
   - The password must contain numerical digits (0-9).
   - The password must contain uppercase letters (A-Z).
   - The password must contain lowercase letters (a-z).
   - The password must contain special characters, such as hyphen (-) and underscore (_).

> **Note:** USM Anywhere passwords expire after 90 days. When your password expires, USM Anywhere enforces a password change when you next log in. A new password must be different from the previous four passwords.After 45 days of inactivity, your user account will be locked. Manager users can unlock inactive accounts.

3. Click **Save & Continue**.
4. When the login page opens, enter the password you just set and click **Login**.

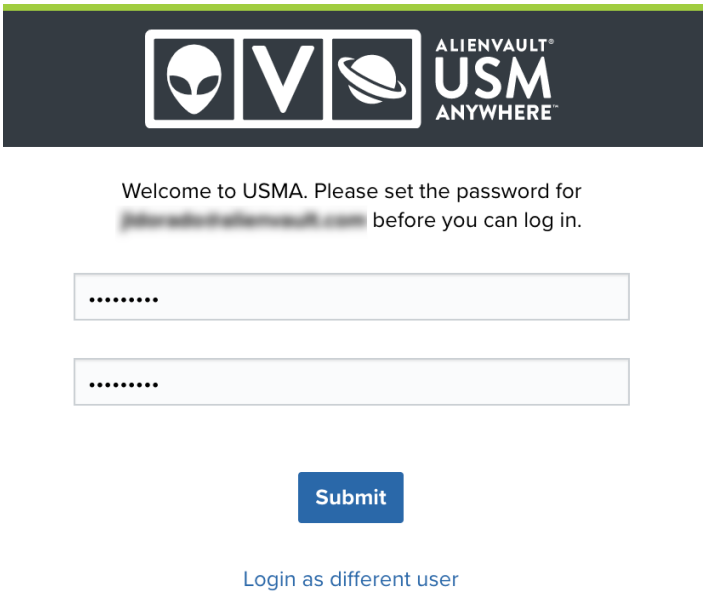

Verify That Your Sensor Is Running

It's a good idea to verify that the USM Anywhere Sensor is running. It also gives you the chance to watch the sensor actively working to find all of your assets and to record events from the start.

> **Note:** Verify that the sensor is running before performing the configuration. You can keep one web browser tab with the Welcome to USM Anywhere page in the background while you perform the verification on a different tab.

To verify that your new sensor is running

1. In USM Anywhere, go to **Data Sources > Sensors**.

   You should now see your sensor in the page. See in the *USM Anywhere User Guide* for more information.

   After a few minutes, USM Anywhere locates your assets and starts generating events.

2. You can review the activity in two locations:
   - From the primary task bar, select **Environment > Assets**.
   - From the primary task bar, select **Activity > Events**.

> **Note:** It could take up to six minutes before events appear. Make sure to refresh your browser from time to time to display the current data.

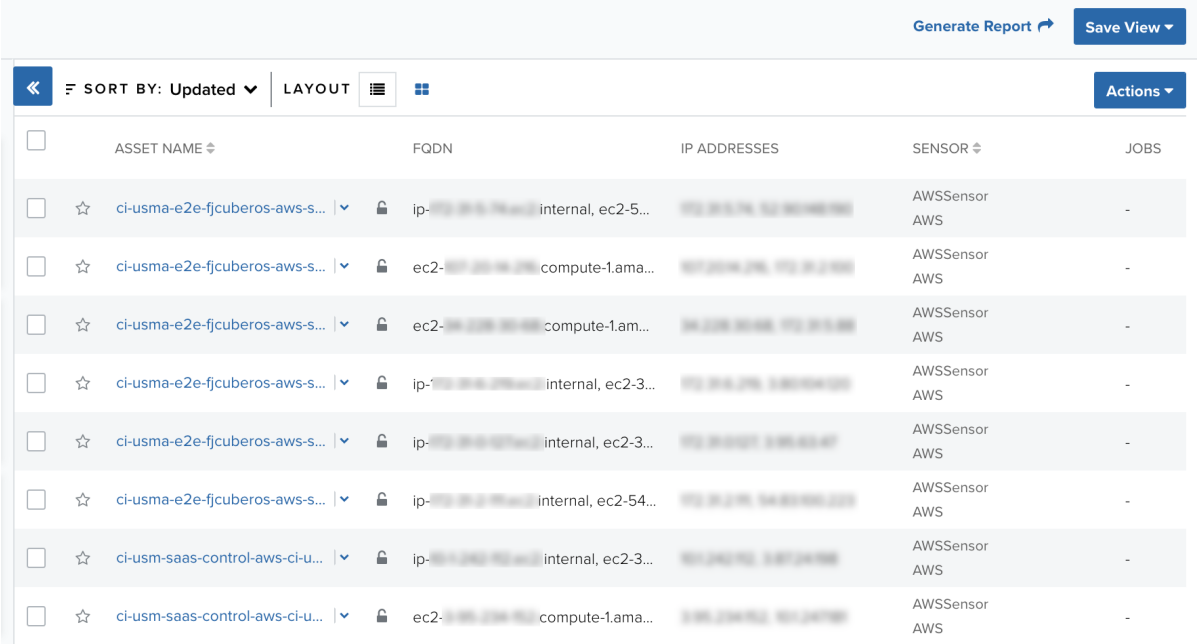

See the *USM Anywhere User Guide* for more information about using the Assets and Events pages in USM Anywhere.

Complete the Hyper-V Sensor Setup

Role Availability	Read-Only	Investigator	Analyst	Manager

After you initialize a new USM Anywhere Sensor, you must configure it in the Setup Wizard. As you complete the Hyper-V Sensor configuration, USM Anywhere performs specific actions, like running an asset discovery scan and collecting logs.

Accessing the Setup Wizard

The Setup Wizard is accessible under the following circumstances:

- After you first log in to the USM Anywhere web user interface (UI) and see the Welcome to USM Anywhere page, click **Get Started** to launch the Setup Wizard.
- If you have already registered one USM Anywhere Sensor but did not complete the setup before logging out, the USM Anywhere Sensor Configuration page launches automatically at your next login to remind you to finalize configuration of the sensor. From that page, you click **Configure** to launch the Setup Wizard and complete the sensor configuration.
- If you registered an additional USM Anywhere Sensor, but did not complete the setup, the Sensors page displays an error (⊗) in the Configured column. See in the *USM Anywhere User Guide* for more information.

  Go to **Data Sources > Sensors**, and then click the sensor name to complete the sensor configuration. See in the *USM Anywhere User Guide* for more information.

Configuring the Sensor in the Setup Wizard

The first time you log in from the Welcome to USM Anywhere web page, the Setup Wizard prompts you to complete the configuration of the first deployed sensor. Thereafter, you can use the Sensors page to configure an additional sensor or to change the configuration options for a deployed sensor. See in the *USM Anywhere User Guide* for more information.

Note: You must have already configured your network interfaces for Hyper-V. See Set Up USM Anywhere on the Hyper-V Virtual Machine for more information.

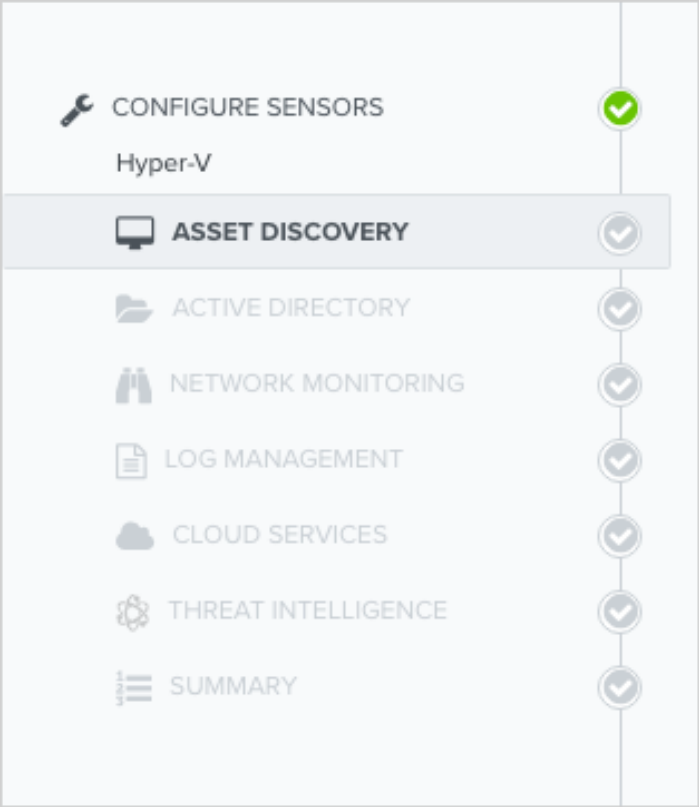

Within the Setup Wizard, complete the configuration on each page.

Asset Discovery

When you move forward to Asset Discovery, a dialog box automatically opens and prompts you to allow asset scanning. USM Anywhere must discover your assets to enable security monitoring on them.

To complete the asset discovery task

1. Click **Yes** to start the automatic asset discovery.

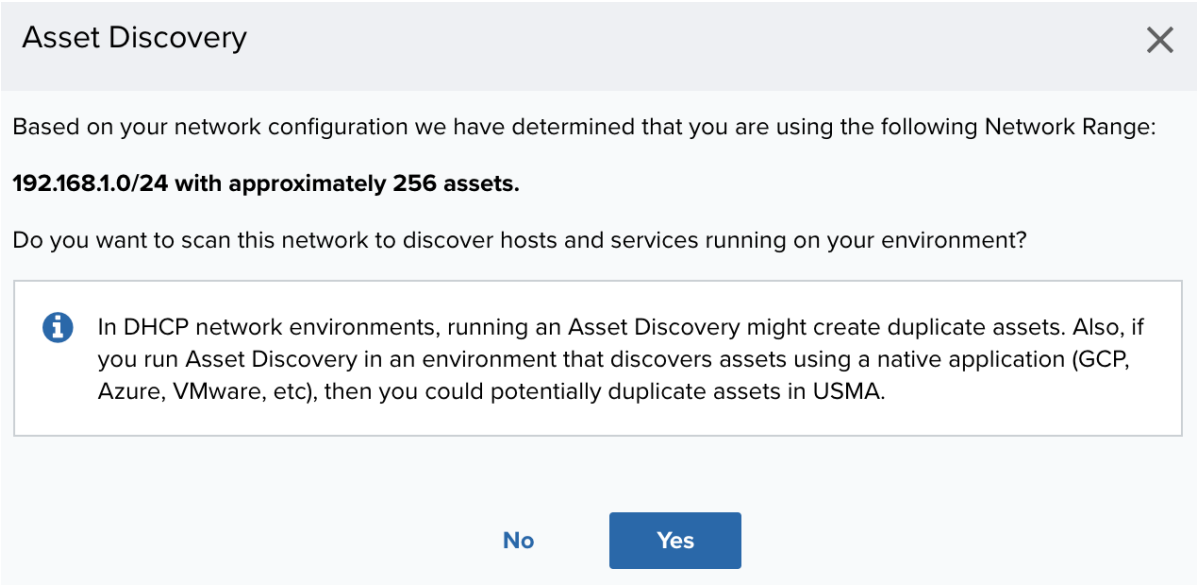

Or if you prefer to add the assets manually or scan another network, click **No** and skip to the next step.

During the automated scan, the Scan Networks status bar opens and displays the number of assets detected in your network range.

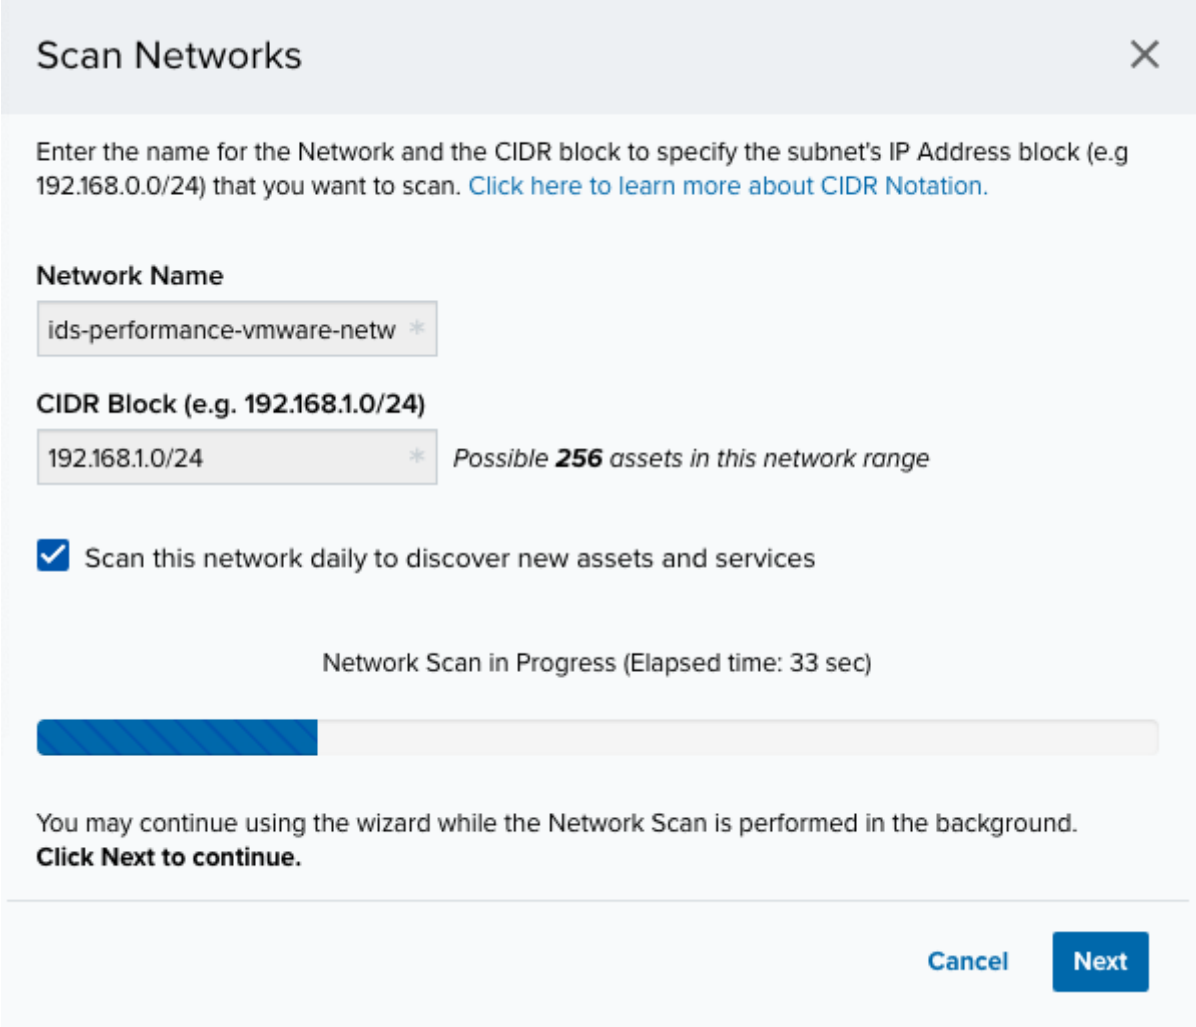

When the scan stops, you have these options:

- Click **Scan Another** to scan a different set of assets
- Click **Next** to continue with asset discovery setup options

When the initial asset scan dialog box closes, the Asset Discovery page displays status information for an ongoing scan or any discovered assets for completed scans.

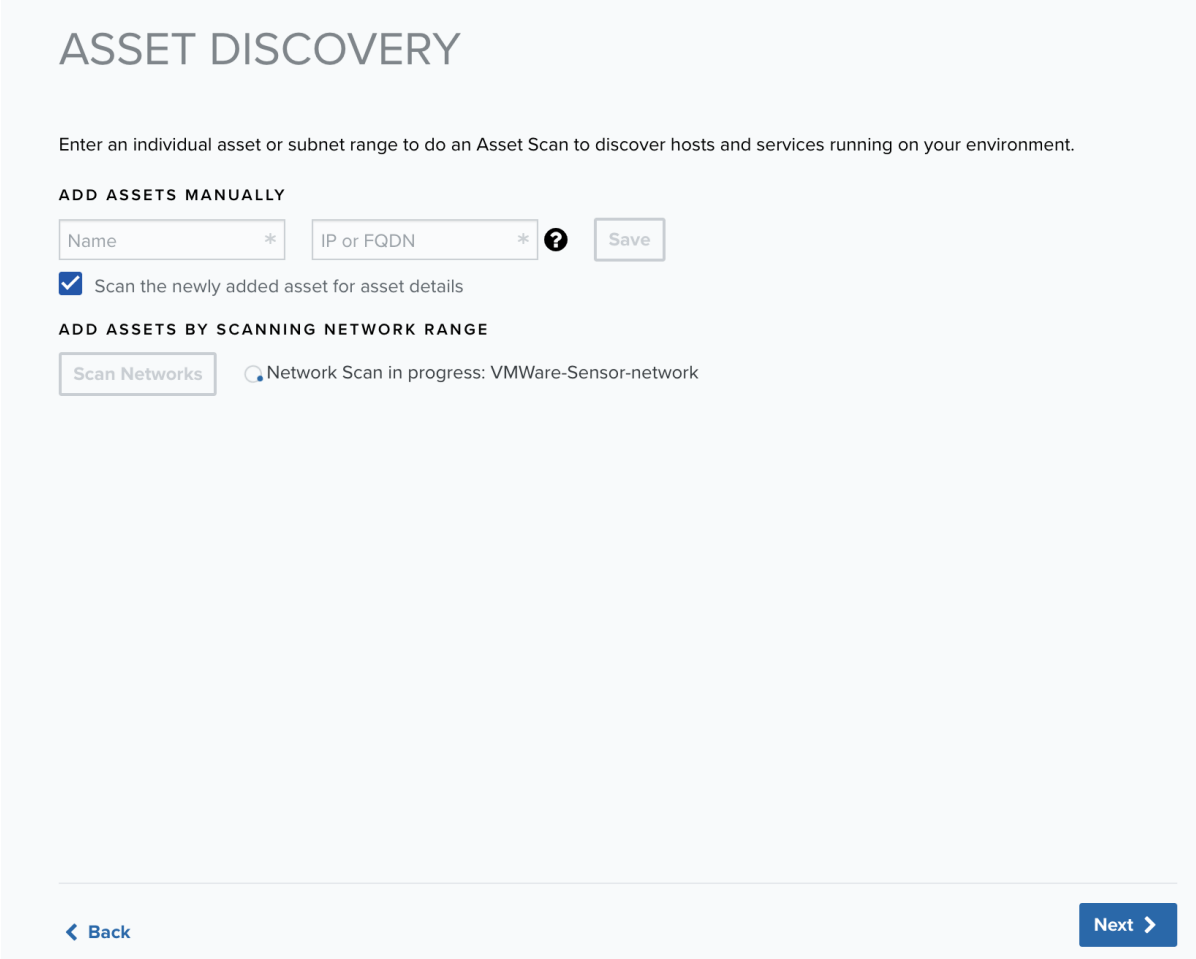

2. (Optional.) **Add assets manually**

   Enter the name and IP address or fully qualified domain name (FQDN) to specify an asset for discovery. The scan option is selected by default. Click **Save** to add the asset.

   You can repeat this for each individual asset you want to add.

3. (Optional.) **Add assets by scanning network range**

   Click **Scan Networks** to scan a network range that you specify. This runs asset discovery to scan hosts and services running on the specified network range.

4. When all the needed assets are discovered, click **Next** at the bottom of the page.

   The wizard opens the next page in the setup process, Active Directory.

Active Directory

The optional Active Directory (AD) setup page configures USM Anywhere to collect information from your AD account. To monitor Microsoft Windows systems effectively, USM Anywhere needs access to the AD server to collect inventory information.

Note: This configuration is only for one AD server. If you want to scan different AD servers, you must create an AD scan job for each of them. See Scheduling Active Directory Scans from the Job Scheduler Page for more information.

AT&T Cybersecurity recommends that you create a dedicated AD account with membership in the Domain Admins group to be used by USM Anywhere to log in to the Windows systems. You also need to activate Microsoft Windows Remote Management (WinRM) in the domain controller and in all the hosts that you want to scan. You can do this by using a group policy for all the systems in your AD.

Important: Before this feature is fully functional, you must configure access to the USM Anywhere Sensor on the AD server. See Granting Access to Active Directory for USM Anywhere for more information.

To complete the AD access configuration

1. Provide the AD credentials for USM Anywhere:
   - **Active Directory IP Address**: Enter the IP address for the AD server.
   - **Username**: Enter your username as admin of the account.
   - **Password**: Enter your admin's password.
   - **Domain**: Enter the domain for the AD instance.

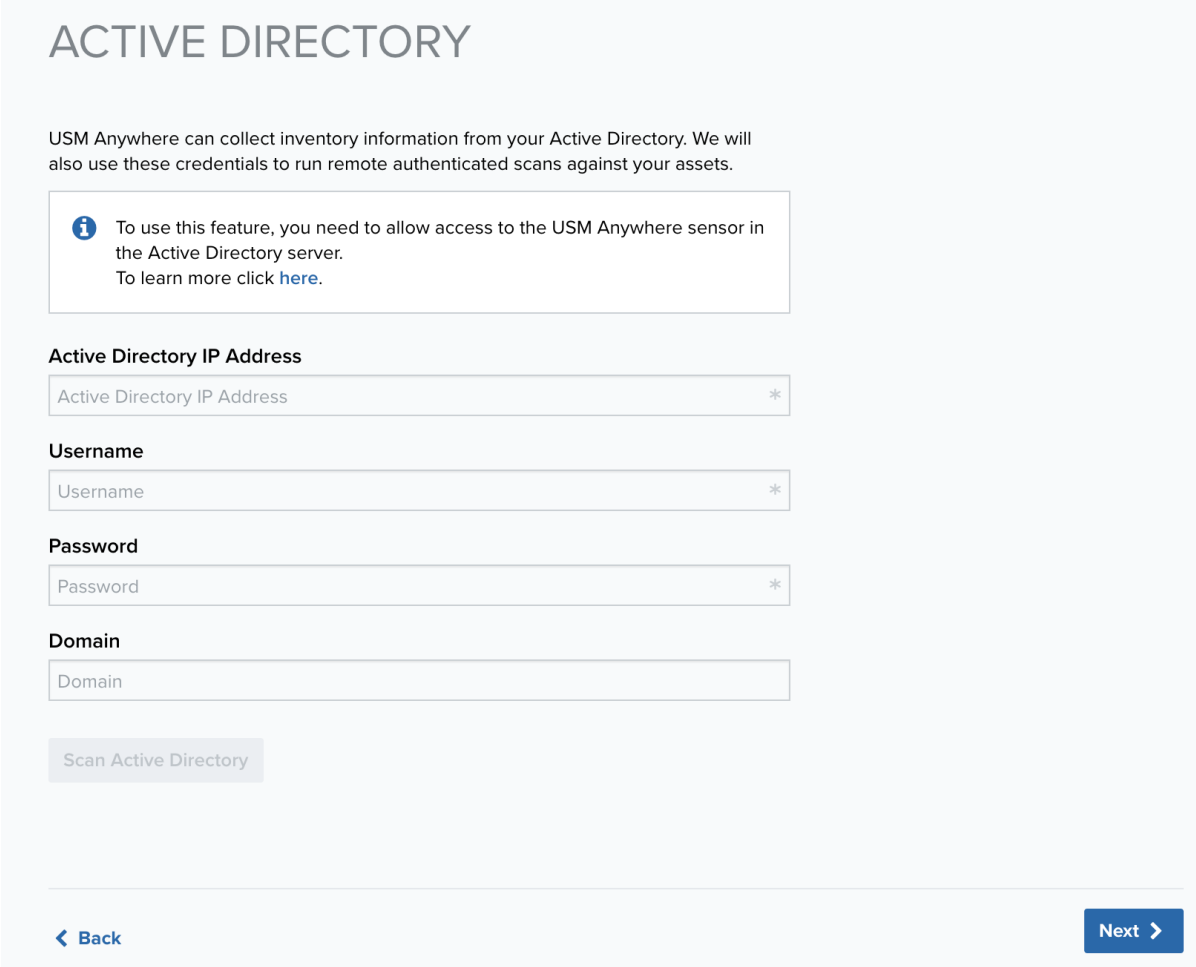

2. Click **Scan Active Directory**.

   After a successful launch of the scan, a confirmation dialog box opens.

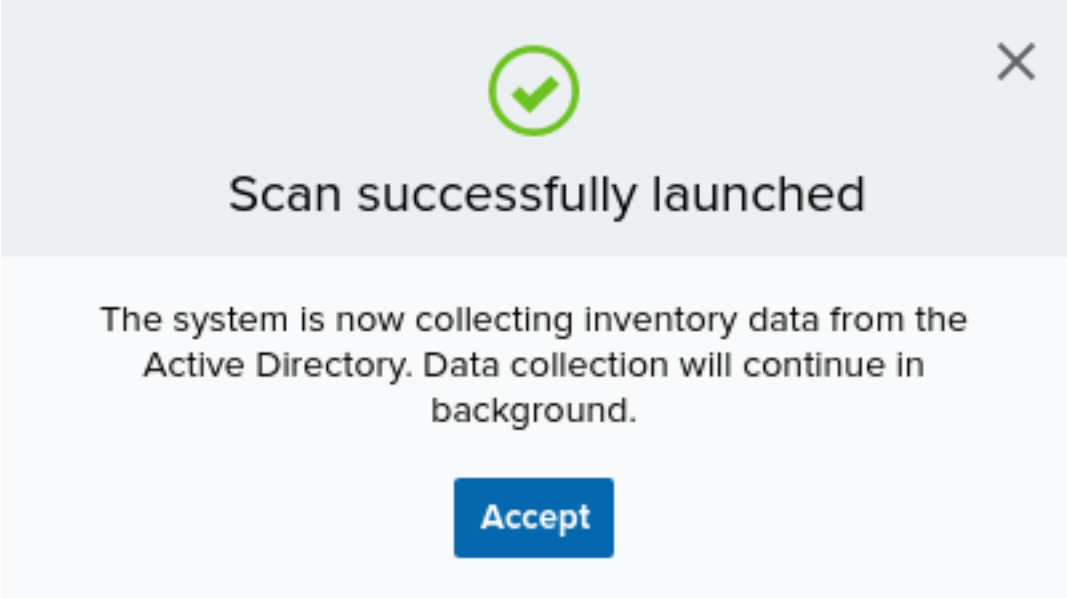

3. Click **Accept**.

   The scan continues in the background.

   Upon completion, another dialog box opens and provides information about the number of assets USM Anywhere discovered. It also prompts you to decide if you want to scan for hosts and services running in your environment.

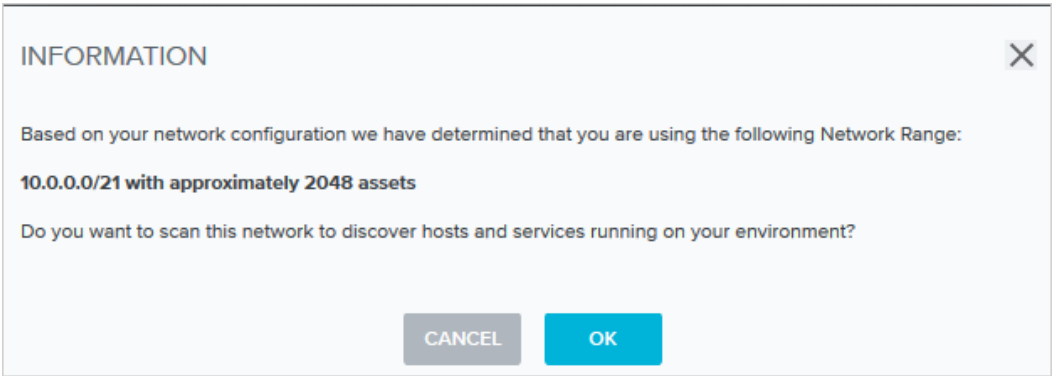

   Click **Cancel** to opt out of this scan.

4. (Optional.) If you want to scan for other hosts and services, click **OK**.

5. Click **Next** after the scan ends.

   The wizard opens the next page in the setup process, Network Security Monitoring.

Network Security Monitoring

The Network Security Monitoring page shows the status of the network interfaces monitored by the sensor (it could take a few moments to load the interfaces). All network adapters are configured for network monitoring by default.

You must manually enable port mirroring or port spanning, promiscuous mode, or both in a virtual switch to send a copy of the network traffic you want to analyze to these interfaces. This page provides links to documentation about how to configure your networking to allow for the interfaces to see the network traffic and perform network intrusion detection.

> **Note:** You must have already configured your network interfaces for Hyper-V. See Set Up USM Anywhere on the Hyper-V Virtual Machine for more information.

USM Anywhere connectivity and communications are handled by the first network interface connection on the Network Security Monitoring page. This is the primary network that provides asset scanning and log collection for the particular network.

You can connect additional interfaces to other networks for monitoring, or connect them to individual vSwitch port groups for virtual networks. Each interface should be connected to a vSwitch that mirrors a different subnet within your network.

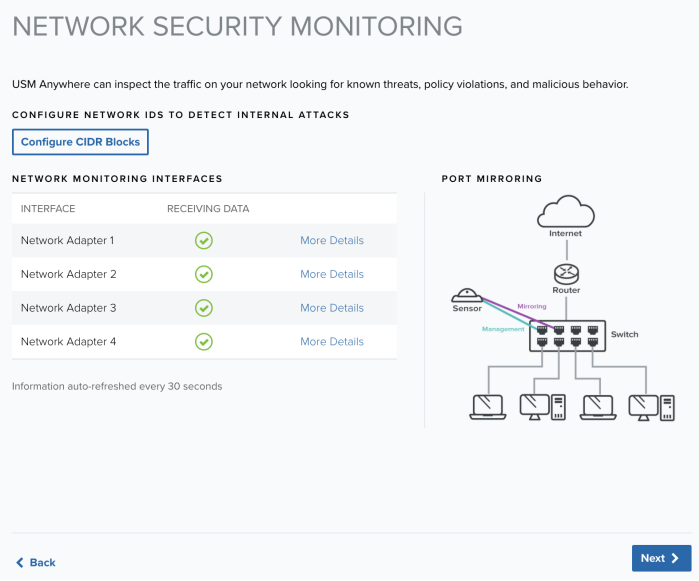

Use this page to verify that USM Anywhere can monitor your network traffic for security events.

> **Note:** You can see red X icons next to the interfaces if the port mirroring or promiscuous mode is not configured. You might also see these icons if the network interfaces have not seen any traffic in the past 30 seconds.

To access detailed information about port mirroring set up

1. Click **How do I set up port mirroring**.

   This opens a dialog box that can direct you to specific information about your VM.

   If you have not yet set up port mirroring, see Direct Traffic from Your Physical Network to the Hyper-V Sensor for more information.

2. Click **Next**.

Log Management

On the Log Management page are syslog port numbers. (These ports are the same for all USM Anywhere Sensors.)

USM Anywhere collects third-party device, system, and application data through syslog over UDP on port 514 and over TCP on ports 601 or 602 by default. It collects Transport Layer Security (TLS)-encrypted data through TCP on ports 6514 or 6515 by default. These ports support the RFC 3164 and RFC 5424 formats. To configure any third-party devices to send data to USM Anywhere, you must provide the IP address and the port number of your USM Anywhere Sensor.

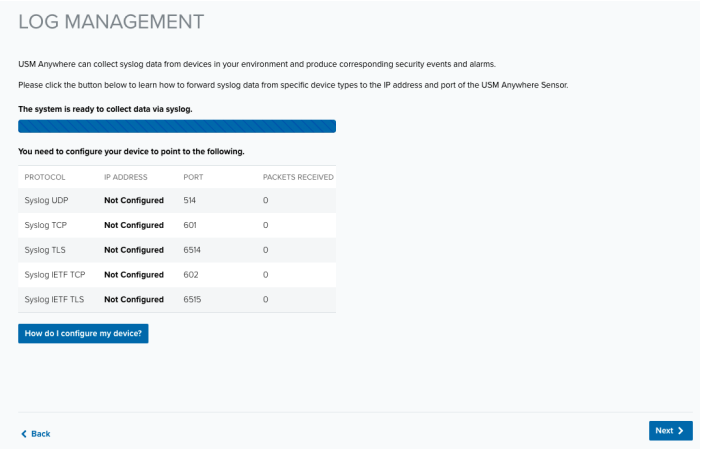

To enable log collection and configure your log management

1. Make sure that you have granted the necessary permissions for your OS to allow USM Anywhere to access its logs. You can also integrate a wide variety of data sources to send log data over syslog to the USM Anywhere Sensor.

   To learn how to configure your operating systems and supported third-party devices to forward syslog log data, see the following related topics:

   - **The Syslog Server Sensor App**: Log collection (UDP, TCP, and TLS-encrypted TCP) from rsyslog
   - **Collecting Linux System Logs**: Log collection from a Linux system
   - Collecting Windows System Logs: Log collection from a Windows system
   - Go to the specific AlienApp in USM Anywhere for instructions about syslog forwarding

   > **Note:** Because the log scan can take some time, you might not see all of the automatically discovered log sources immediately after deploying the first USM Anywhere Sensor.

2. When you have finished the log collection setup and integrated any needed plugins, verify that the data transfer is occurring.
3. Click **Next** when this step is complete.

Cloud Services

Use the **Cloud Services** page to configure USM Anywhere to collect information from any of the supported cloud services apps. These apps allow you to monitor and detect threats against your cloud services accounts, such as G Suite (Google Apps) and Office 365, directly from USM Anywhere. When configured , the AlienApp collects log data from via the cloud service API and analyzes that data against our built-in threat intelligence to look for anomalies and intrusions.

To setup cloud services in USM Anywhere

1. Select the AlienApp for each cloud service where you want to monitor activity.

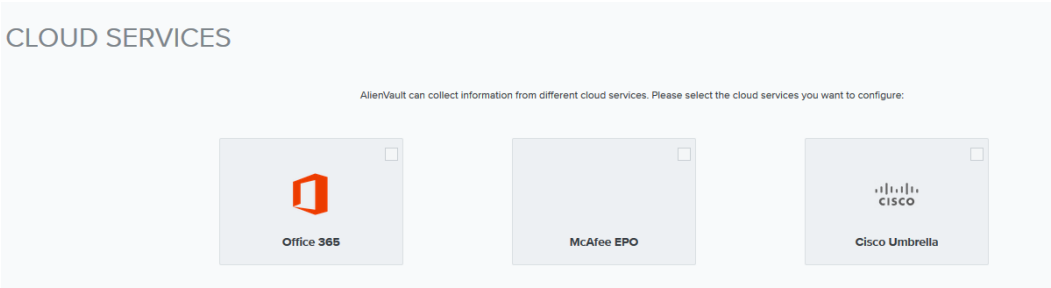

2. Click **Next**.

   USM Anywhere displays a configuration page for each cloud service you selected.

3. Complete the form for its configuration and click **Save Credentials**.

4. Click **Next.** when this step is complete.

OTX

AT&T Alien Labs™ Open Threat Exchange® (OTX™) is an open information-sharing and analysis network providing users with the ability to collaborate, research, and receive alerts on emerging threats and indicators of compromise (IoCs) such as IP addresses, file hashes, and domains.

You must have an OTX account to receive alerts based on threats identified in OTX. This account is separate from your USM Anywhere account. Go to The World's First Truly Open Threat Intelligence Community to create an OTX account.

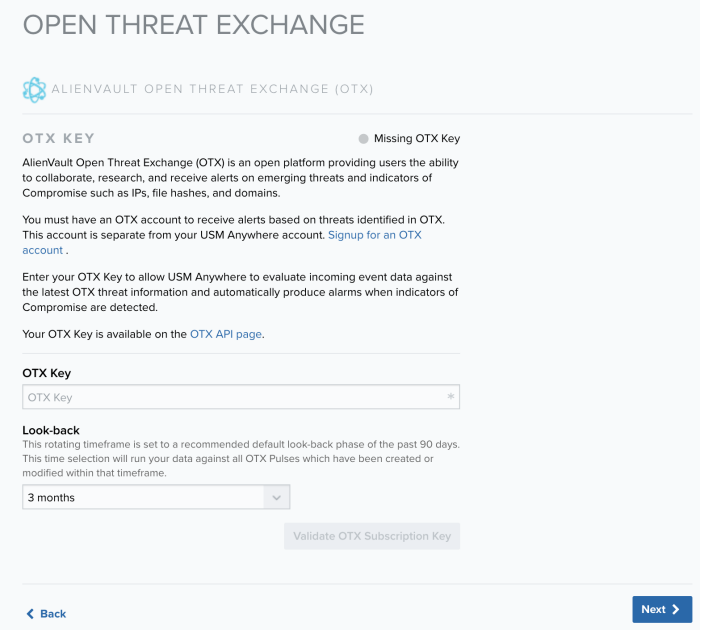

Note: If you do not already have an OTX account, click the **Sign up** link. This opens another browser tab or window that displays the OTX signup page. After you confirm your email address, you can log in to OTX and retrieve the unique API key for your account.

See Open Threat Exchange® and USM Anywhere in the *USM Anywhere User Guide* for more information about OTX integration in USM Anywhere.

To enable USM Anywhere to evaluate event data against the latest OTX intelligence

1. Log in to OTX and open the API page (https://otx.alienvault.com/api).
2. In the DirectConnect API Usage pane, click the icon to copy your unique OTX connection key.

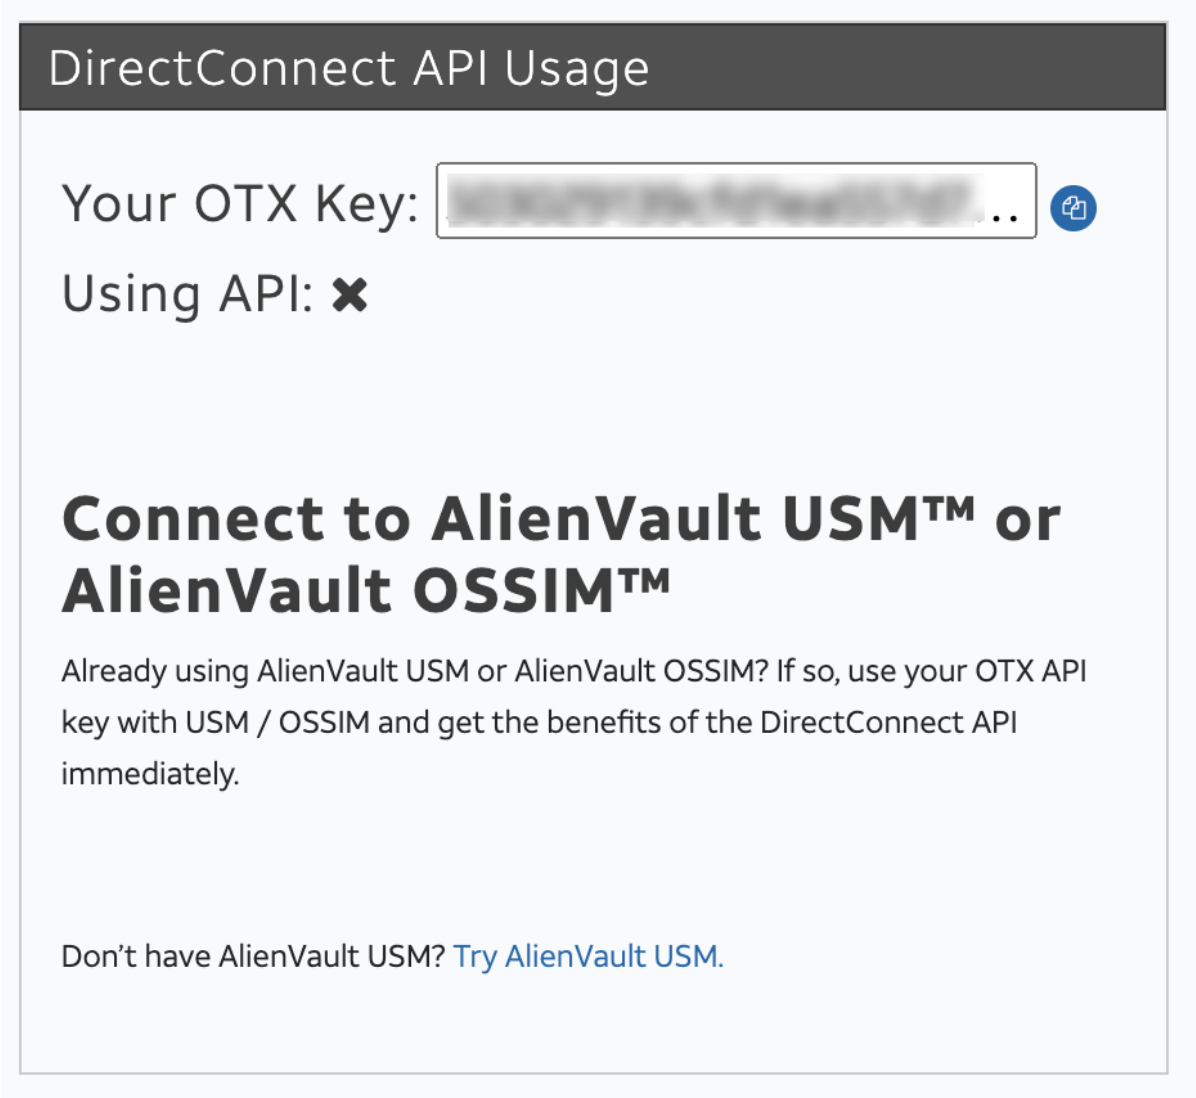

3. Return to the Open Threat Exchange (OTX) page of the USM Anywhere Sensor Setup Wizard and paste the value in the OTX Key text box.

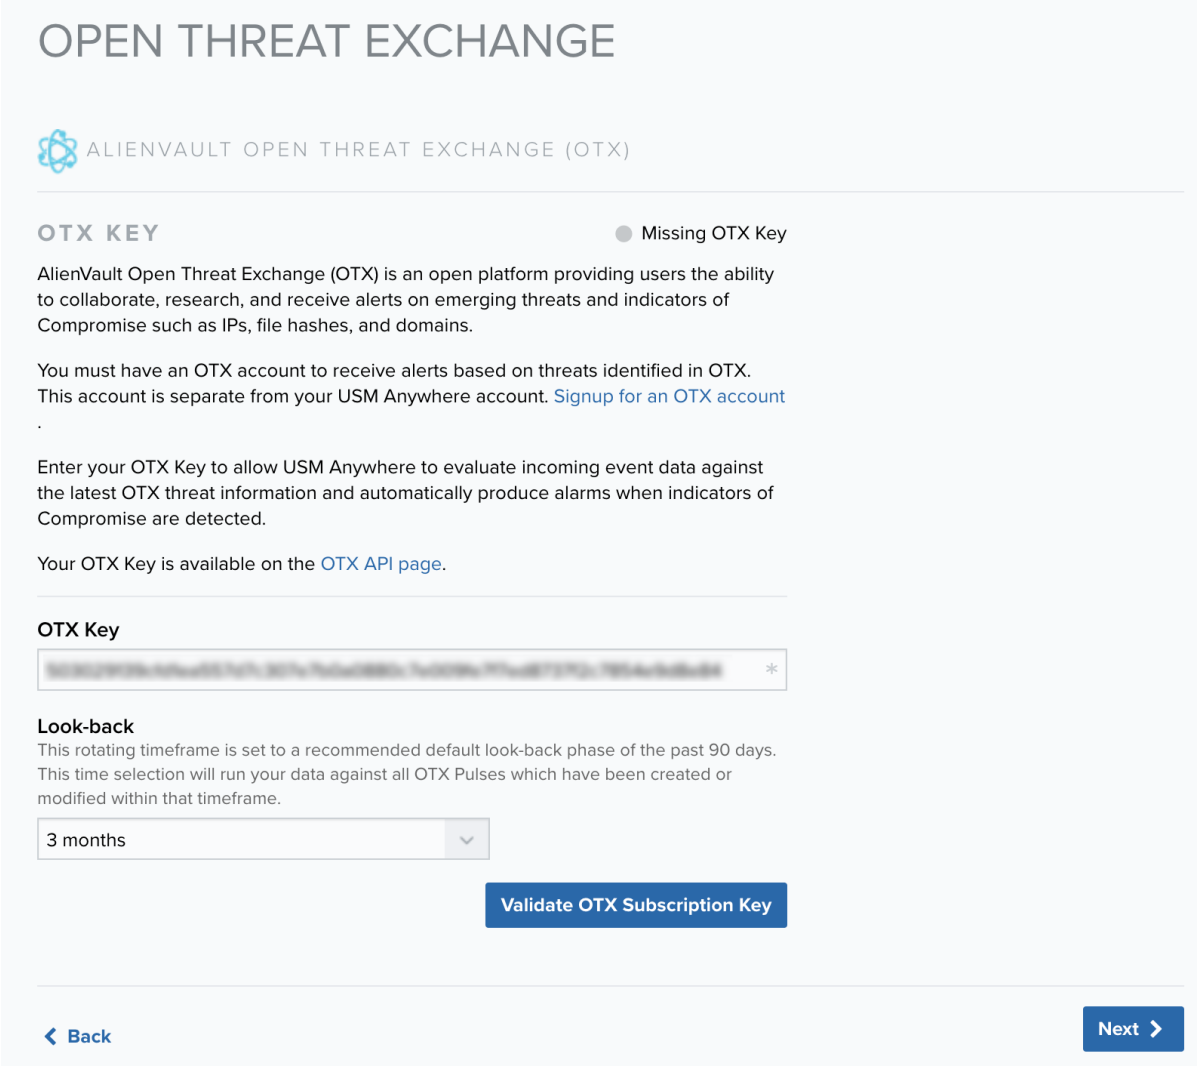

4. Click **Validate OTX Subscription Key**.

   With a successful validation of the key, the status at the top of the page changes to "Valid OTX key".

5. Click **Next** when this task is complete.

Setup Complete

The Congratulations page summarizes the status of your configuration.

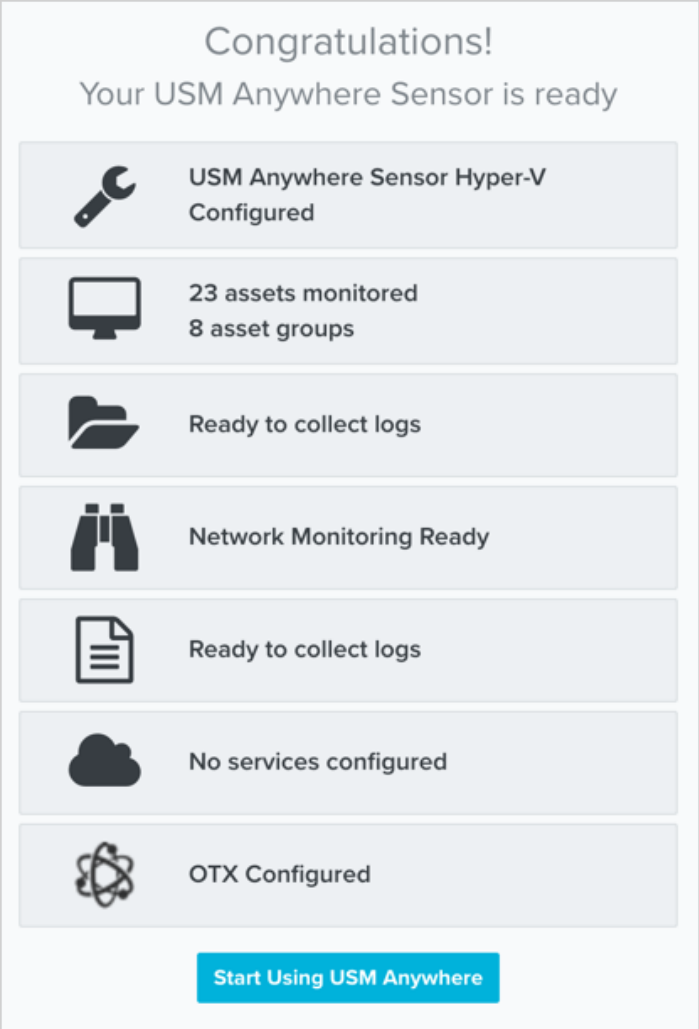

Click **Start Using USM Anywhere**, which takes you to the Overview dashboard.

Next...

Now is a great time to run a vulnerability scan. See Vulnerability Assessment in the *USM Anywhere User Guide* for detailed information about running a vulnerability scan.

USM Anywhere Sensor Deployment on AWS

The USM Anywhere Sensor provides operational visibility into the security of your Amazon Web Services (AWS) environment. Based on the collected log information, USM Anywhere analyzes the data generated by your AWS environment and provides real-time alerting to identify malicious activity. When the sensor is deployed into your AWS environment, it provides ultimate control over the installation and the data contained within it, and also prevents any external access to your environment.

This section includes the following topics:

About AWS Sensor Deployment 115

Requirements for AWS Sensor Deployment 116

Deploy the AWS Sensor 129

Connect the AWS Sensor to USM Anywhere 133

Complete the AWS Sensor Setup 137

Enable Connections in an AWS VPC 148

VPC Traffic Mirroring with an AWS Sensor 150

Collect Logs from Amazon S3 Buckets with KMS Encryption 158

AWS Log Discovery and Collection in USM Anywhere 159

About AWS Sensor Deployment

All USM Anywhere Sensors allow for authenticated scans of assets by leveraging stored credentials that you define in USM Anywhere. This enables USM Anywhere to detect potential vulnerabilities, installed software packages, and running processes and services. Unlike the other USM Anywhere Sensors, the Amazon Web Services (AWS) Sensor queries AWS directly to discover assets using an AWS API.

Log Collection and Scans

The AWS Sensor collects AWS logs and system logs, and generates asset scans and vulnerability assessments, consisting of the following:

- AWS CloudTrail logs
- AWS Elastic Load Balancing (ELB) logs
- Amazon Simple Storage Service (S3) access logs
- Amazon CloudWatch log collection
- Amazon S3 log collection
- Operational logs for critical software packages deployed, such as HTTP servers and database servers
- Asset scans on your virtual machines (VMs) to inventory installed software packages, running processes, and services
- Periodic vulnerability assessments

Log Analysis

USM Anywhere analyzes these logs in these stages:

Stage 1: Collects logs from systems and software running in your environment

Stage 2: Configures log line processing and generates events

- Includes IP addresses and timestamps culled from extracted log-line data
- Adds other data to the event, such as security context and environmental information

Stage 3: Analyzes events and stores them

Deployment Overview

AT&T Cybersecurity distributes the AWS Sensor as a CloudFormation Template in a virtual private cloud (VPC).

The deployment process for an initial USM Anywhere Sensor in your AWS environment consists of these primary tasks:

1. Review requirements for an AWS Sensor deployment.
2. Deploy the USM Anywhere Sensor within your AWS environment.
3. Register the sensor with your sensor authentication code to provision the USM Anywhere instance and connect the deployed sensor.
4. Complete your AWS Sensor configuration, including initial asset discovery.

Requirements for AWS Sensor Deployment

USM Anywhere deploys the Amazon Web Services (AWS) Sensor in the Amazon Elastic Compute Cloud (EC2) platform through the Amazon Virtual Private Cloud (VPC).

This table includes the requirements for the AWS Sensor deployment.

AWS Sensor Deployment Requirements

Requirement	Description
m5.large instance	An m5.large instance in an Amazon VPC.
100 GB EBS `/data` volume	The Amazon Elastic Block Store (EBS) provides short-term storage for your data as it is processed. A 100 GB Amazon EBS `/data` volume is designated as the default size for optimal performance.

AWS Sensor Deployment Requirements (Continued)

Requirement	Description
Internet connection to the USM Anywhere Secure Cloud	Review the Sensor Ports and Connectivity for more information.
Zonal SSD persistent disks	Persistent disk storage offers reliable network storage that your instances can access like physical disks. For optimal performance, 50 GB and 128 GB volumes are designated as the default size for the root and data partitions.

Important: Because the needs of a sensor differ based on the varying demands of different deployment environments and the complexity of events being processed, the number of events per second (EPS) a sensor can process varies.

Depending on your environment, you may need to deploy additional sensors to ensure that all events are processed.

AWS Sensor Deployment Regions

The AWS Sensor is deployed in one of the AWS endpoint regions based on your location. The following table lists the code and name of each region.

Important: The Update Server and the AlienVault Agent always use the 3.235.189.112/28 range no matter where your USM Anywhere is deployed. The AT&T TDR for Gov Update Server uses the 3.32.190.224/28 range.

AWS Regions for AWS Sensor Deployment

Code	Name
ap-northeast-1	Asia Pacific (Tokyo)
ap-northeast-2	Asia Pacific (Seoul)
ap-northeast-3	Asia Pacific (Osaka)
ap-south-1	Asia Pacific (Mumbai)

AWS Regions for AWS Sensor Deployment (Continued)

Code	Name
ap-southeast-1	Asia Pacific (Singapore)
ap-southeast-2	Asia Pacific (Sydney)
ca-central-1	Canada (Central)
eu-central-1	Europe (Frankfurt)
eu-north-1	Europe (Stockholm)
eu-west-1	Europe (Ireland)
eu-west-2	Europe (London)
eu-west-3	Europe (Paris)
me-central-1	Middle East (UAE)
sa-east-1	South America (São Paulo)
us-east-1	US East (N. Virginia)
us-east-2	US East (Ohio)
us-west-1	US West (N. California)
us-west-2	US West (Oregon)

Note: Though your sensor is deployed to one region it will monitor all regions, pulling assets, logs, and other data regardless of region.

Application Service Dependencies

With the AWS CloudFormation Template provided by AT&T Cybersecurity, you can automatically deploy USM Anywhere as a service into your environment. Review the following tables for information about the outbound and inbound IP addresses, ports, and services used by USM Anywhere.

Note: To launch the USM Anywhere Sensor web UI during the initial setup, you need to allow inbound traffic to the sensor IP address through TCP port 80. You can remove access to this port after the sensor successfully connects to USM Anywhere. You do not need to allow inbound traffic to this port from the Internet.

The following tables list the inbound and outbound ports.

Sensor Ports and Connectivity (Outbound Ports)

Type	Ports	Endpoints	Purpose
TCP	443	update.alienvault.cloud	Communication with AT&T Cybersecurity for initial setup and future updates of the sensor.
TCP	443	reputation.alienvault.com	Ongoing communication with AT&T Alien Labs™ Open Threat Exchange® (OTX™).
TCP	443	otx.alienvault.com	Ongoing communication with OTX to retrieve vulnerability scores. Connecting to otx.alienvault.com is not required but highly recommended. OTX uses the AWS CloudFront services. Refer to the AWS IP address ranges page when you deploy a new sensor. This page contains the current IP address ranges for the service and instructions on how to filter the addresses.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
SSL	443	storage-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send and retrieve backups.
SSL	443	metrics-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send metrics and messages.
SSL/TCP	443	api-parameters-<REGION>-prod.alienvault.cloud[3] api-message-proxy-<REGION>-prod.alienvault.cloud api.message-proxy.<REGION>.prod.alienvault.cloud	Ongoing communication with USM Anywhere. It is only necessary to allowlist the address that corresponds to the region where your USM Anywhere instance is hosted.
SSL/TCP	7100	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
UDP	53	DNS Servers (Google Default)	Ongoing communication with USM Anywhere.
UDP	123	169.254.169.123	Sync with network time protocol (NTP) services.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	22 and 443	prod-usm-saas-tractorbeam.alienvault.cloud prod-gov-usm-saas-tractorbeam.gov.alienvault.us (for AT&T TDR for Gov)	SSH communications with the USM Anywhere remote support server. See Troubleshooting and Remote Sensor Support for more information about remote technical support through the USM Anywhere Sensor console.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	geoip-ap-northeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-south-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-2-prod.alienvault.cloud/geo-ip/sensor geoip-ca-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-me-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-sa-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-us-gov-west-1-prod-gov.gov.alienvault.us/geo-ip/sensor (for AT&T TDR for Gov)	Allows resolution of IP addresses for geolocation services. It is only necessary to allowlist the GeoIP address that corresponds to the region where your USMA instance is hosted.

1

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

2

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

3

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

4

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

5

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

Important: A USM Anywhere Sensor deployed in AWS might require outbound access to specific AWS resources based on the sensor app in use. For example, the AWS Sensor app must have the ability to connect to the AWS API (port 443). However, the actual API endpoint might differ depending on the service (such as Amazon Simple Storage Service [S3] or Amazon CloudWatch).

USM Anywhere normally gives systems explicit access to the AWS API.

Sensor Ports and Connectivity (Inbound Ports)

Type	Ports	Purpose
SSH	22	Inbound method for secure remote login from a computer to USM Anywhere.
HTTP	80	Inbound communication for HTTP traffic.
UDP (RFC 3164)	514	USM Anywhere collects data through syslog over UDP on port 514 by default.

Sensor Ports and Connectivity (Inbound Ports) (Continued)

Type	Ports	Purpose
TCP (RFC 3164)	601	Inbound communication for reliable syslog service. USM Anywhere collects data through syslog over TCP on port 601 by default.
TCP (RFC 5424)	602	USM Anywhere collects data through syslog over TCP on port 602 by default.
Traffic Mirroring	4789	Inbound communication for virtual extensible local area network (VXLAN).
WSMANS	5987	Inbound WBEM WS-Management HTTP over Secure Sockets Layer/Transport Layer Security (SSL/TLS) (NXLog).
TLS/TCP (RFC 3164)	6514	USM Anywhere collects TLS-encrypted data through syslog over TCP on port 6514 by default.
TLS (RFC 5424)	6515	USM Anywhere collects data through syslog over TLS on port 6515 by default.
TCP	9000	Inbound communication used internally for HTTP sensor traffic.
Graylog	12201	Inbound communication for Graylog Extended Log Format (GELF).

Security Groups in Your AWS VPC

For sensor deployment in an AWS VPC, the AWS CloudFormation template automatically creates the security groups needed for network connectivity between the instances within the VPC. However, this does require that you manually assign the USMServicesSG security group to your hosts to enable access to the UDP port 514 so that the sensor can receive syslog packet transmissions.

See Enable Connections in an AWS VPC for more detailed information about these security groups.

AWS Services

USM Anywhere uses the following AWS services:

- Amazon CloudWatch
- AWS CloudTrail
- AWS Elastic Load Balancing (ELB)
- Amazon Simple Storage Service (S3)

- Amazon EC2
- AWS Identity and Access Management (IAM)
- Amazon GuardDuty
- Amazon Relational Database Service (RDS)

See IAM Roles and Permissions Required by Your AWS Sensor for a full description of the IAM roles and permissions that your AWS Sensor requires for these AWS services.

Note: USM Anywhere uses *us-east-1* as a default region in the amazon-aws app. As a result, you might want to verify whether your sensors are communicating with *us-east-1*, even if you have never deployed to that region.

USM Anywhere IP Addresses for Allowlisting

Your sensor is connected to a USM Anywhere instance deployed in one of the Amazon Web Services (AWS) endpoint regions based on your location. If you need to configure your firewall to allow communication between the sensor and the USM Anywhere instance, refer to the following table with the reserved IP address ranges for each region.

Important: The Update Server and the AlienVault Agent always use the 3.235.189.112/28 range no matter where your USM Anywhere is deployed. The AT&T TDR for Gov Update Server uses the 3.32.190.224/28 range.

Note: The regional IP ranges listed in this table are limited to the control nodes (subdomain). You must also meet all requirements provided in the Sensor Ports and Connectivity (Outbound Ports) table.

AWS Regions Where USM Anywhere Instance Is Available

Code	Name	Reserved Static IP Address Ranges
ap-northeast-1	Asia Pacific (Tokyo)	18.177.156.144/28 3.235.189.112/28 44.210.246.48/28
ap-south-1	Asia Pacific (Mumbai)	3.7.161.32/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-1	Asia Pacific (Singapore)	18.143.203.80/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-2	Asia Pacific (Sydney)	3.25.47.48/28 3.235.189.112/28 44.210.246.48/28
ca-central-1	Canada (Central)	3.96.2.80/28 3.235.189.112/28 44.210.246.48/28
eu-central-1	Europe (Frankfurt)	18.156.18.32/28 3.235.189.112/28 44.210.246.48/28
eu-west-1	Europe (Ireland)	3.250.207.0/28 3.235.189.112/28 44.210.246.48/28

AWS Regions Where USM Anywhere Instance Is Available (Continued)

Code	Name	Reserved Static IP Address Ranges
eu-west-2	Europe (London)	18.130.91.160/28 3.235.189.112/28 44.210.246.48/28
me-central-1	Middle East (UAE)	3.29.147.0/28 3.235.189.112/28 44.210.246.48/28
sa-east-1	South America (São Paulo)	18.230.160.128/28 3.235.189.112/28 44.210.246.48/28
us-east-1	US East (N. Virginia)	3.235.189.112/28 44.210.246.48/28
us-west-2	US West (Oregon)	44.234.73.192/28 3.235.189.112/28 44.210.246.48/28
us-gov-west-1	AWS GovCloud (US-West)	3.32.190.224/28

Installation Prerequisites

Before you install the AWS Sensor, make sure you have the following prerequisites available.

Installation Prerequisites

Prerequisites	Description
AWS CloudFormation template provided by AT&T Cybersecurity	The AWS CloudFormation template automatically creates all required AWS resources for deployment, including an IAM role and instance profile for use by the USM Anywhere Sensor instance. URLs for these templates are included in the Deploy the AWS Sensor instructions.
Privileged user account on AWS	To deploy the AWS CloudFormation template, you must have a privileged user account in AWS with permissions to create IAM resources.

Multiple AWS Accounts or Amazon VPCs

If you have multiple AWS accounts, you must deploy the AWS Sensor in each AWS account that you want to monitor.

Amazon VPC enables you to launch AWS resources into a virtual network that you have defined. A single sensor can monitor an entire AWS account, even when it contains multiple Amazon VPCs.

Note: If you intend to use the USM Anywhere vulnerability scanner with the AWS Sensor, you must allow traffic from the sensor and the target instance you are scanning. You can usually accomplish this by using Amazon VPC peering (see the AWS documentation for more information).

IAM Roles and Permissions Required by Your AWS Sensor

The roles and permissions detailed below are required by the AWS services listed, on which your AWS Sensor relies.

During deployment, the AWS CloudFormation template provided by AT&T Cybersecurity automatically manages and assigns these as needed by your sensor.

The following table shows the IAM roles and permissions required by your AWS Sensor.

Warning: The sensor's capacity to extract the information will be endangered if you disable the below services. The sensor won't have permission to perform the disabled function.

IAM Roles and Permissions Required by Your AWS Sensor

Prerequisites	Description
Amazon CloudWatch	• "cloudwatch:Describe*" • "cloudwatch:Get*" • "cloudwatch:List*" • "logs:Describe*" • "logs:Get*" • "logs:TestMetricFilter*"
AWS CloudTrail	• "cloudtrail:Describe*" • "cloudtrail:Get*" • "cloudtrail:List*"
AWS Elastic Load Balancing (ELB)	• "elasticloadbalancing:Describe*"
Amazon Simple Storage Service (S3)	• "s3:Get*" • "s3:List*"
Amazon EC2	• "ec2:Describe*"
AWS IAM	• "iam:List*" • "iam:Get*"
Amazon GuardDuty	• "guardduty:Get*" • "guardduty:List*"
Amazon Relational Database Service (RDS)	• "rds:Describe*"

Deploy the AWS Sensor

Role Availability	Read-Only	Investigator	Analyst	Manager

After you review the requirements and make sure that your Amazon Web Services (AWS) environment is configured as needed, you can deploy the AWS Sensor. Using the AWS CloudFormation Template provided by AT&T Cybersecurity, you automatically deploy USM Anywhere as a service into your environment.

The following procedure describes how to launch the AWS Sensor when provisioning the USM Anywhere service for the first time. In this process, you launch the USM Anywhere product from the AWS Management Console using the AWS CloudFormation template.

Important: If you are using these instructions to redeploy an existing AWS Sensor, your IP address will not be the same as for your previous sensor. After these steps are complete, you must also update any syslog or NXLog log collection, and any port mirroring to use the new IP address.

To create a new sensor in the AWS Management Console

1. Log in to the AWS Management Console.
2. Under Find Services, enter a name, keyword, or acronym to launch the AWS CloudFormation service page.
3. In the upper right corner, click **Create stack**, and then select **With new resources (standard)**.

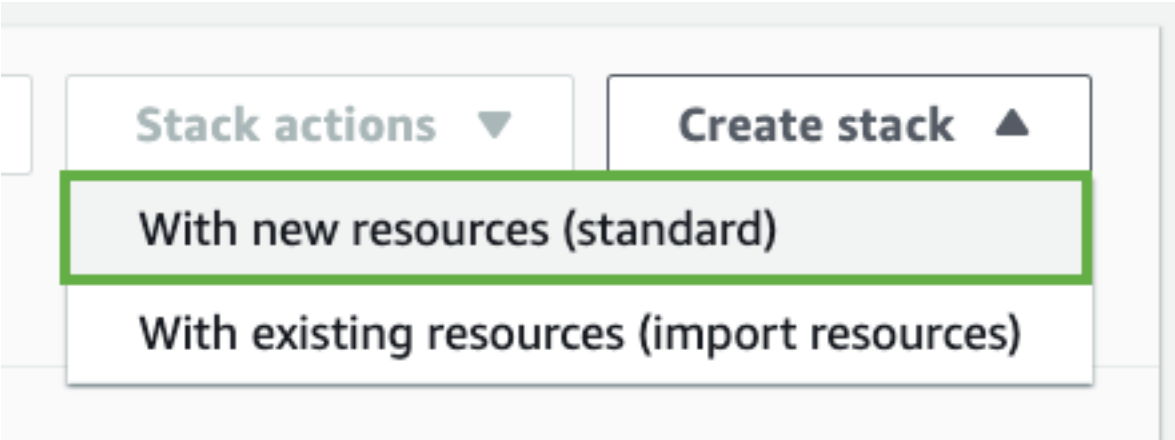

4. Go to the USM Anywhere Sensor Downloads page, click the icon of your specific sensor, and copy the URL.
5. Use the copied URL in the Amazon Simple Storage Service (S3) URL field.
6. Click **Next**, and then click **Next** again to continue.

7. On the Specify stack details page, in the Stack name text box, enter a name to identify the stack.

   The name must be one word. Use hyphens if desired. For example, you could call the stack "USM-sensor-1".

8. Set parameters for the AWS Sensor:

   Note: The volume size should be prefilled. You can leave this setting at the default value.

   - In the USM Anywhere Sensor Name text box, enter a name for the sensor. This is usually the same as the stack name.
   - In the Key Name list, select the key pair that allows SSH connections to the sensor. See AWS documentation, Create or import a key pair, for more information.
   - In the Traffic Mirroring Mode list, select **Yes** to deploy a sensor ready for VPC traffic mirroring, or select **No** to deploy a sensor without those additional considerations.

     Note: See Enabling VPC Traffic Mirroring for more information on this feature.

   - In the HTTP Access Range text box, specify the IP address range that allows HTTP access to the sensor.
   - In the SSH Access Range text box, specify the IP address range that allows SSH access to the sensor.

9. Click **Next**.

10. Select the appropriate VPC ID and subnet ID, specify whether to use a public or private IP address, and then click **Next**.

    Important: If you choose to deploy your sensor with a public IP address, the subnet you select must have *Auto-assign public IPv4 address* enabled.

11. (Optional.) On the Configure stack options page, set tags for the instance, and then click **Next**.

12. On the Review page, select the checkbox at the bottom of the page next to the statement "I acknowledge that AWS CloudFormation might create IAM resources."

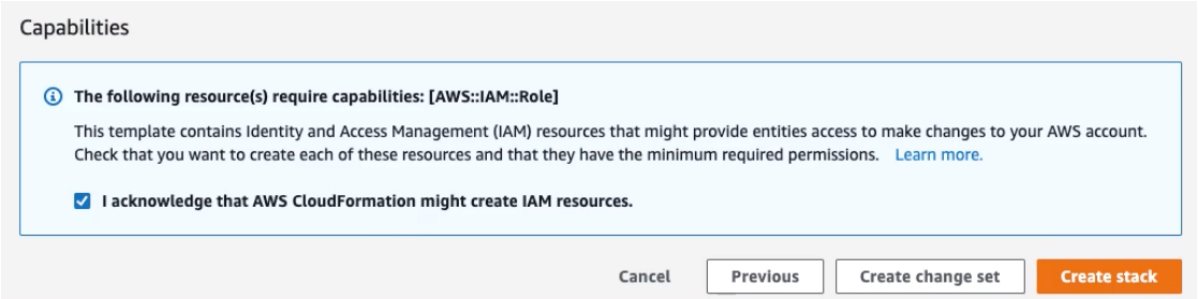

13. Click **Create stack**.

14. In the Stacks page, confirm that your newly-created stack status reads like this:

```
CREATE_IN_PROGRESS
```

Stack creation typically takes about 15 minutes. When the stack build is complete, you see the following confirmation:

```
CREATE_COMPLETE
```

> **Note:** See the Troubleshooting CloudFormation page for more information about the possible errors with your AWS CloudFormation stack.

15. After your new stack is complete, click the **Outputs** tab and locate the URL.

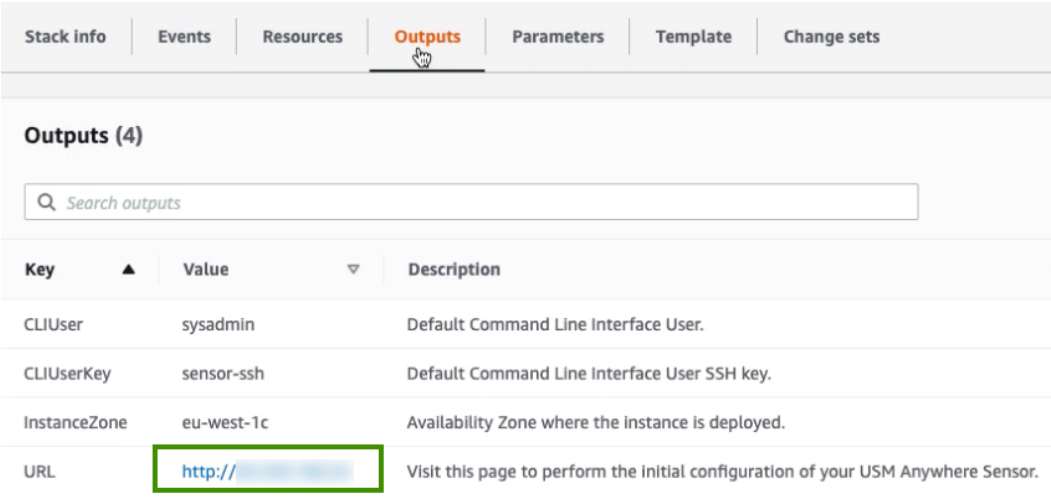

This URL is based on the public IPv4 address of your deployed sensor (http://<ip-address>). Make note of this address so that you have it for configuring your data sources to send data to the AWS Sensor.

See the AWS documentation for more information about managing public IPv4 addresses.

16. Click the URL link to launch the USM Anywhere Sensor Setup page.

Connect the AWS Sensor to USM Anywhere

Role Availability	Read-Only	Investigator	Analyst	Manager

After deploying the Amazon Web Service (AWS) Sensor, you must connect it to USM Anywhere through registration.

Obtain the Authentication Code

You must enter an authentication code when registering the USM Anywhere Sensor. How to obtain the authentication code depends on your USM Anywhere instance and whether this is the first sensor you're deploying.

Instructions for USM Anywhere customers:

If this is your first USM Anywhere Sensor, you must register the sensor using the initial authentication code (starts with a "C") received from AT&T Cybersecurity. With this code, the registration process provisions a new USM Anywhere instance and defines its attributes, such as how many sensors to allow for connection, how much storage to provide, and what email address to use for the initial user account. After registration, you will gain access to the sensor through the USM Anywhere web user interface (UI), where you can complete the sensor setup.

If you are deploying additional sensors, you must generate the authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Instructions for AT&T TDR for Gov customers:

AT&T Cybersecurity has already provisioned the AT&T Threat Detection and Response for Government (AT&T TDR for Gov) instance for you, therefore you won't receive an authentication code for your sensor. This is true regardless if it's the first sensor or additional

sensors you're deploying. However, for the first sensor, you'll receive a link to access your instance.

For every sensor you deploy, you must generate an authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Register Your Sensor

You perform this procedure after deploying the USM Anywhere Sensor within your AWS account. The URL link is displayed after you create the USM Anywhere Sensor stack and the instance is running in your AWS account.

To register your sensor

1. Click the URL displayed for the running stack in the AWS console.

   This opens the *Welcome to USM Anywhere Sensor Setup* page, which prompts you to provide the information for registering the sensor with your new USM Anywhere instance.

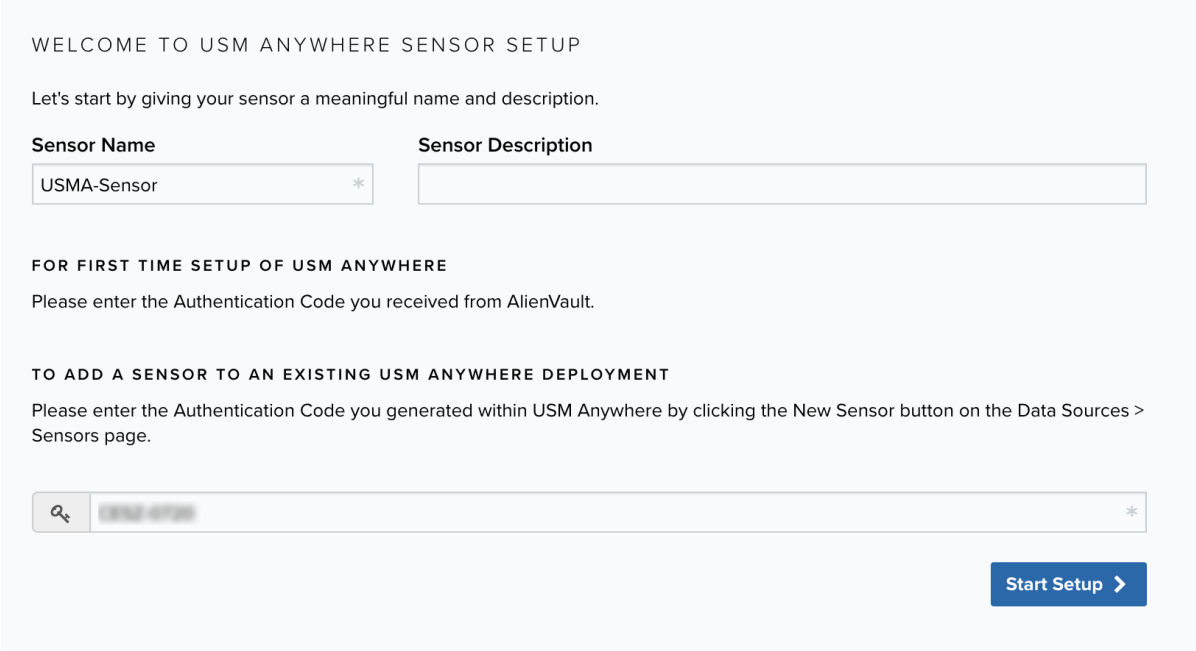

2. Enter a sensor name and sensor description.
3. Paste the authentication code into the field with the key icon ( ).
4. Click **Start Setup** to start the process of connecting the USM Anywhere Sensor.

It takes about 20 minutes to provision your USM Anywhere instance upon registration of your initial sensor. When this instance is provisioned and running, you'll see a welcome message that provides an access link.

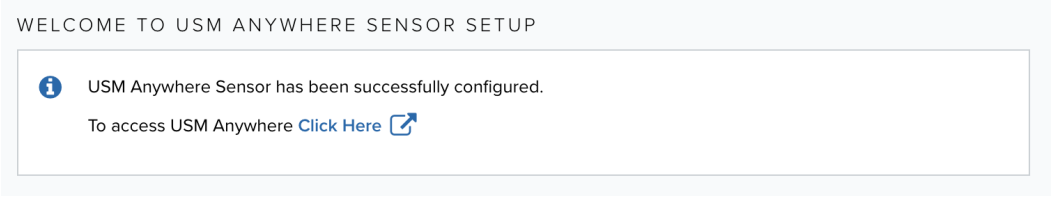

Use this link to open the secured web console for your USM Anywhere instance. You and the other USM Anywhere users in your organization can access this console from a web browser on any system with internet connectivity.

Note: If this is your first deployment, you'll also receive an email from AT&T Cybersecurity that provides the access link to USM Anywhere.

Configure the Initial Login Credentials

When you link to a newly provisioned USM Anywhere instance, you must configure the password for the initial user account. This is the default administrator as defined in your subscription.

To configure login credentials

1. In the welcome message, click the link.

   This displays a prompt to set the password to use for the default administrator of USM Anywhere.

2. Enter the password, and then enter it again to confirm.

   Keep in mind these points when you are logging in:

   - The login credentials that you set will apply to any USM Anywhere™ and USM Central™ you have access to.
   - USM Anywhere requires all passwords to have a minimum length of 8 characters and a maximum length of 128 characters.
   - The password must contain numerical digits (0-9).

- The password must contain uppercase letters (A-Z).
- The password must contain lowercase letters (a-z).
- The password must contain special characters, such as hyphen (-) and underscore (_).

> **Note:** USM Anywhere passwords expire after 90 days. When your password expires, USM Anywhere enforces a password change when you next log in. A new password must be different from the previous four passwords.After 45 days of inactivity, your user account will be locked. Manager users can unlock inactive accounts.

3. Click **Save & Continue**.
4. When the login page opens, enter the password you just set and click **Login**.

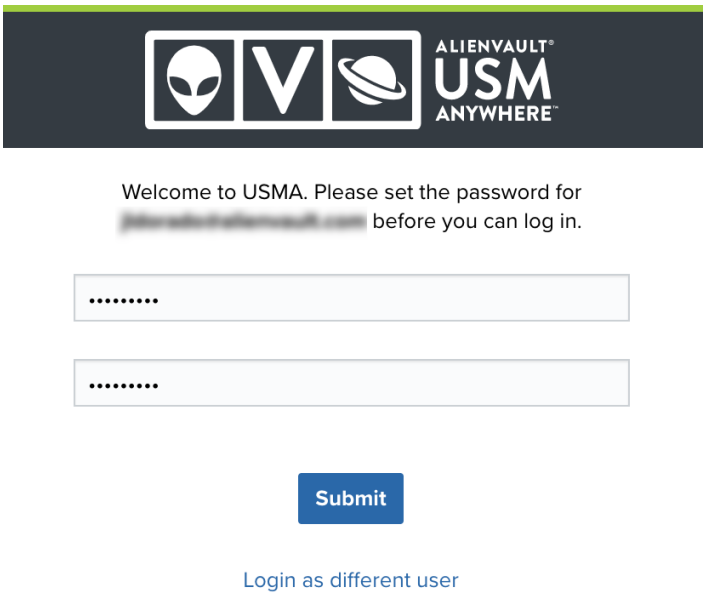

Verify That Your Sensor Is Running

It's a good idea to verify that the USM Anywhere Sensor is running. It also gives you the chance to watch the sensor actively working to find all of your assets and to record events from the start.

> **Note:** Verify that the sensor is running before performing the configuration. You can keep one web browser tab with the Welcome to USM Anywhere page in the background while you perform the verification on a different tab.

To verify that your new sensor is running

1. In USM Anywhere, go to **Data Sources > Sensors**.

   You should now see your sensor in the page. See in the *USM Anywhere User Guide* for more information.

   After a few minutes, USM Anywhere locates your assets and starts generating events.

2. You can review the activity in two locations:

   - From the primary task bar, select **Environment > Assets**.
   - From the primary task bar, select **Activity > Events**.

   Note: It could take up to six minutes before events appear. Make sure to refresh your browser from time to time to display the current data.

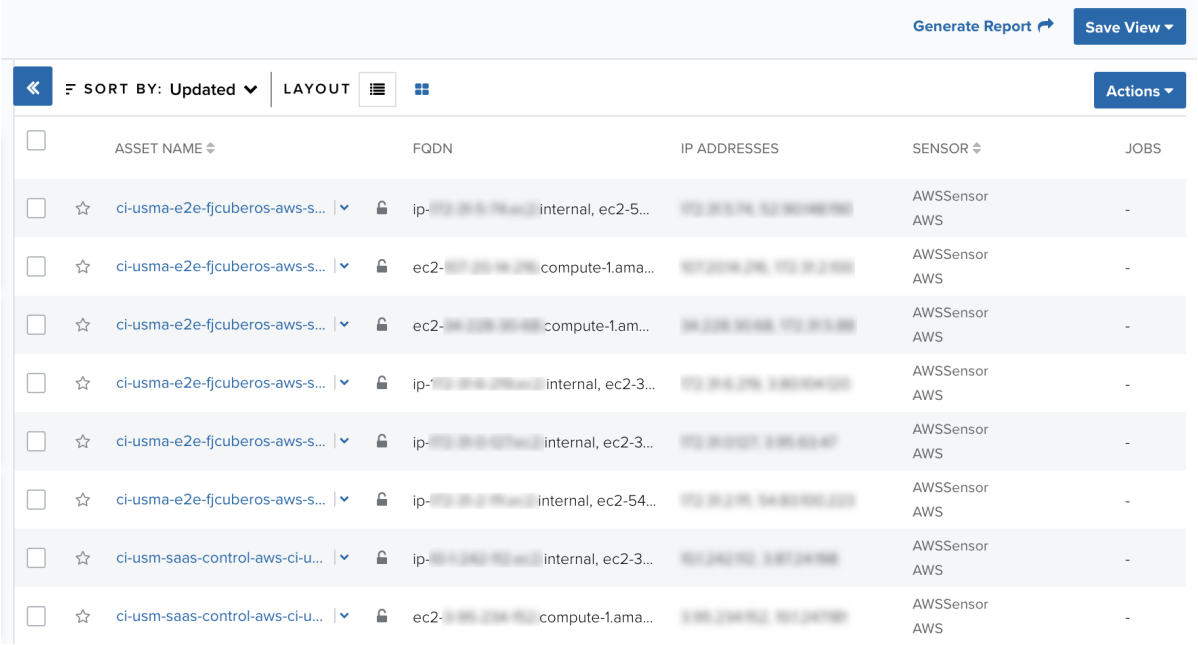

   See the *USM Anywhere User Guide* for more information about using the Assets and Events pages in USM Anywhere.

Complete the AWS Sensor Setup

Role Availability	Read-Only	Investigator	Analyst	Manager
	✗	✗	✗	✓

After you initialize a new USM Anywhere Sensor, you must configure it in the Setup Wizard. As you configure the sensor, you can enable USM Anywhere to perform specific actions through scheduled jobs, such as running an asset discovery scan or collecting security events from a predefined cloud storage location.

About Accessing the Setup Wizard

The Setup Wizard is accessible under the following circumstances:

- After you first log in to the USM Anywhere web user interface (UI) and see the Welcome to USM Anywhere page, click **Get Started** to launch the Setup Wizard.
- If you have already registered one USM Anywhere Sensor but did not complete the setup before logging out, the USM Anywhere Sensor Configuration page launches automatically at your next login to remind you to finalize configuration of the sensor. From that page, you click **Configure** to launch the Setup Wizard and complete the sensor configuration.
- If you registered an additional USM Anywhere Sensor, but did not complete the setup, the Sensors page displays an error (⊗) in the Configured column. See in the *USM Anywhere User Guide* for more information.

  Go to **Data Sources > Sensors**, and then click the sensor name to complete the sensor configuration. See in the *USM Anywhere User Guide* for more information.

Configuring the Sensor in the Setup Wizard

The first time you log in from the Welcome to USM Anywhere web page, the Setup Wizard prompts you to complete the configuration of the first deployed sensor. Thereafter, you can use the Sensors page to configure an additional sensor or to change the configuration options for a deployed sensor. See in the *USM Anywhere User Guide* for more information.

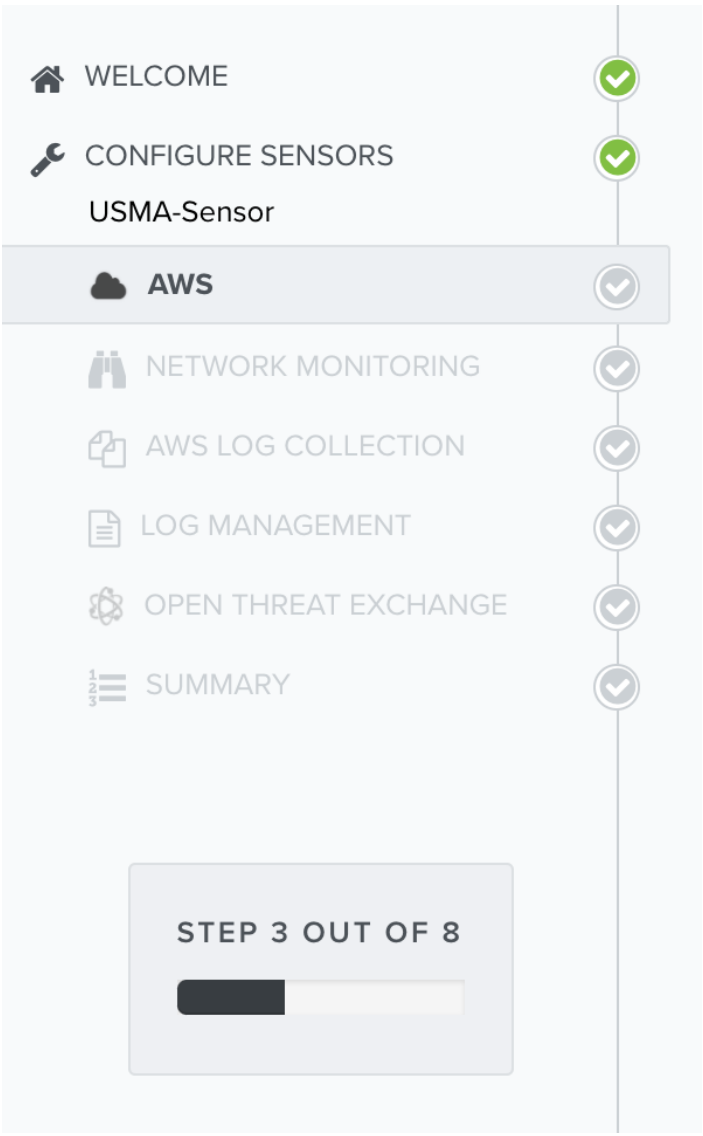

The Amazon Web Services Configuration page provides information about the asset discovery that occurs upon the initial deployment of the USM Anywhere Sensor, summarizing the number of instances, instance types, and regions in your environment.

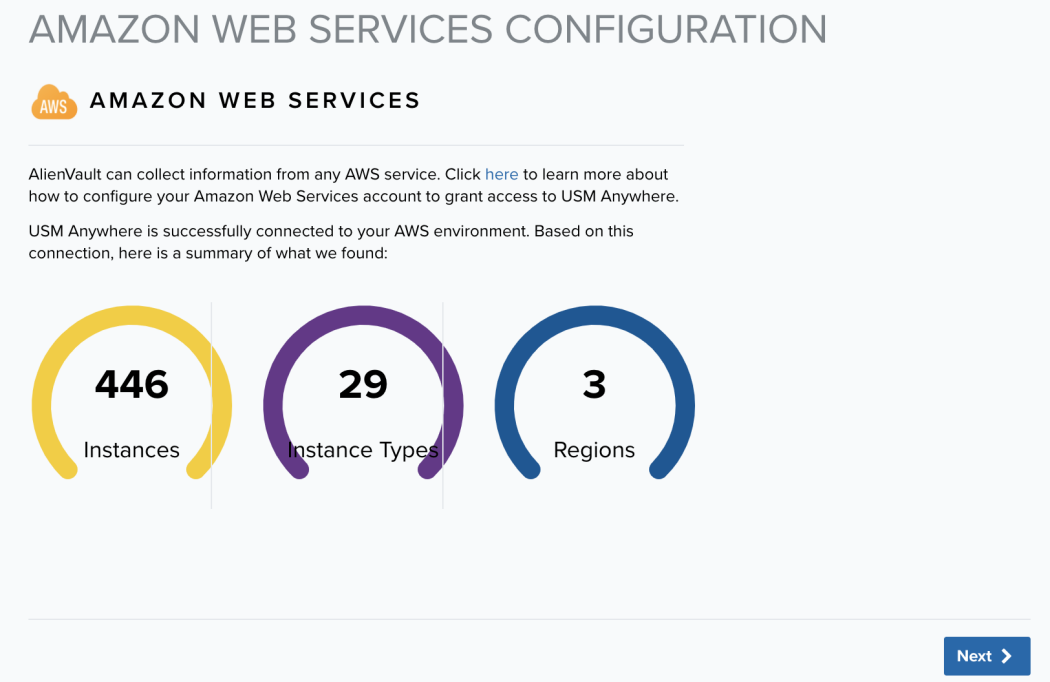

Click **Next** to proceed with the Setup Wizard and complete additional configuration on each page.

Network Security Monitoring

The Network Security Monitoring page shows the status of the network interfaces monitored by the sensor (it could take a few moments to load the interfaces). All network adapters are configured for network monitoring by default.

You must manually enable port mirroring or port spanning, promiscuous mode, or both in a virtual switch to send a copy of the network traffic you want to analyze to these interfaces. This page provides links to documentation about how to configure your networking to allow for the interfaces to see the network traffic and perform network intrusion detection.

USM Anywhere connectivity and communications are handled by the first network interface connection on the Network Security Monitoring page. This is the primary network that provides asset scanning and log collection for the particular network.

You can connect additional interfaces to other networks for monitoring, or connect them to individual vSwitch port groups for virtual networks. Each interface should be connected to a vSwitch that mirrors a different subnet within your network.

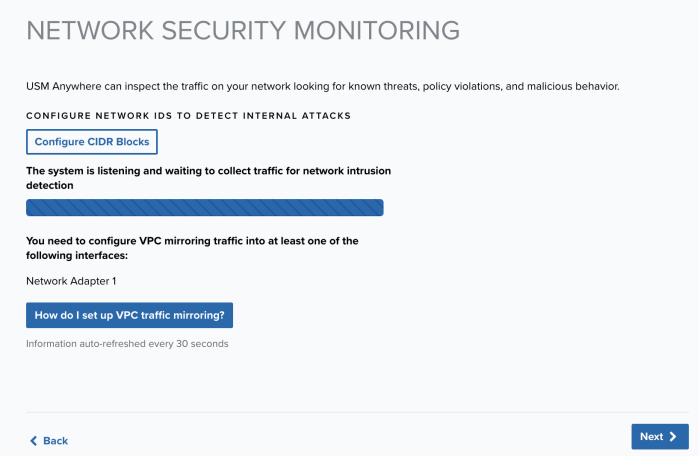

Use this page to verify that USM Anywhere can monitor your network traffic for security events.

Note: You can see red X icons next to the interfaces if the port mirroring or promiscuous mode is not configured. You might also see these icons if the network interfaces have not seen any traffic in the past 30 seconds.

To access detailed information about virtual private cloud (VPC) traffic mirroring

1. Click **How do I set up VPC traffic mirroring?**

   This opens a dialog box.

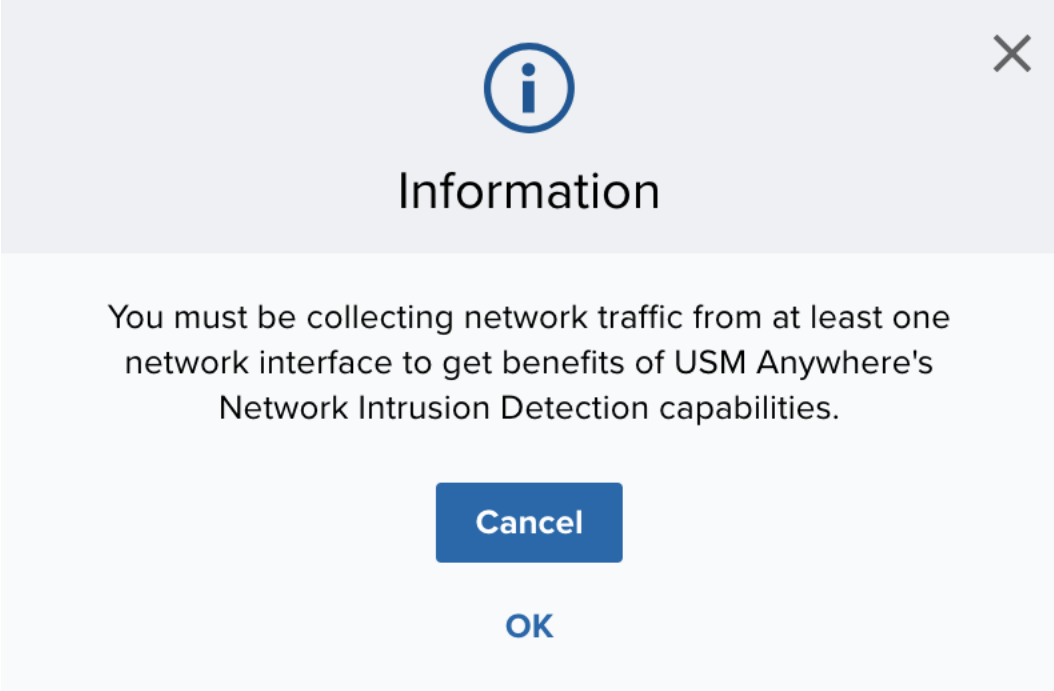

   If you have not yet set up VPC traffic mirroring, see VPC Traffic Mirroring with an AWS Sensor for more information.

2. Click **Next**.

AWS Log Collection

USM Anywhere automatically discovers several of out-of-box logs as long as you have enabled them within your Amazon Web Services (AWS) subscription. See AWS Log Discovery and Collection in USM Anywhere for more information about these logs and how they function within the AWS environment.

To enable the Amazon Simple Storage Service (S3) and Amazon CloudWatch out-of-box log collection jobs

1. Locate the jobs you want to enable and click the icon.

   This turns the icon green ( ).

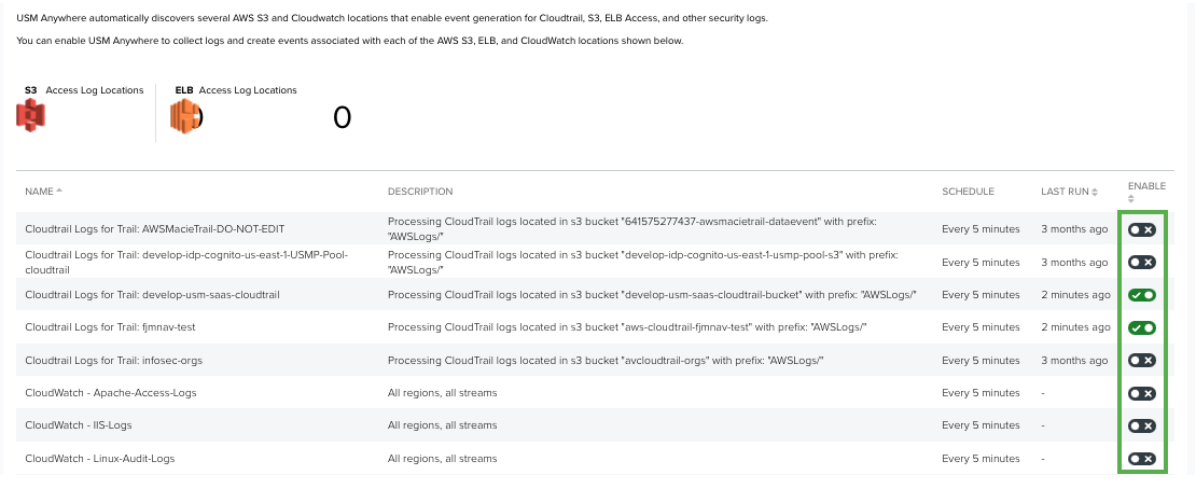

> **Note:** You can also enable AWS CloudTrail logs, Amazon Elastic Load Balancing (ELB) access logs, and other security logs. However, make sure you've enabled these first on your AWS account.

2. Click **Next**.

Log Management

On the Log Management page are syslog port numbers. (These ports are the same for all USM Anywhere Sensors.)

USM Anywhere collects third-party device, system, and application data through syslog over UDP on port 514 and over TCP on ports 601 or 602 by default. It collects Transport Layer Security (TLS)-encrypted data through TCP on ports 6514 or 6515 by default. These ports support the RFC 3164 and RFC 5424 formats. To configure any third-party devices to send data to USM Anywhere, you must provide the IP address and the port number of your USM Anywhere Sensor.

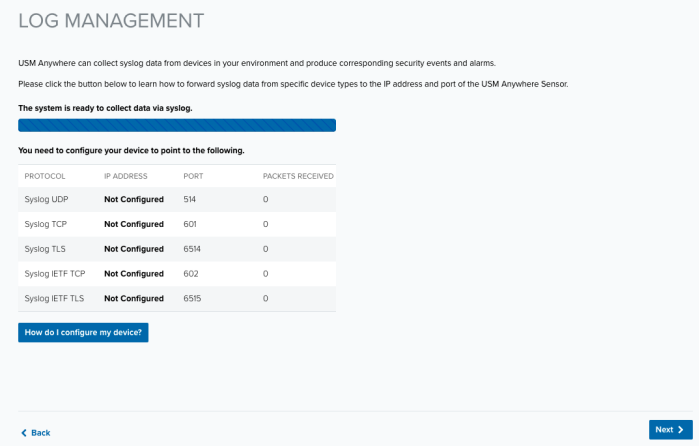

To enable log collection and configure your log management

1. Make sure that you have granted the necessary permissions for your OS to allow USM Anywhere to access its logs. You can also integrate a wide variety of data sources to send log data over syslog to the USM Anywhere Sensor.

   To learn how to configure your operating systems and supported third-party devices to forward syslog log data, see the following related topics:

   - **The Syslog Server Sensor App**: Log collection (UDP, TCP, and TLS-encrypted TCP) from rsyslog
   - **Collecting Linux System Logs**: Log collection from a Linux system
   - Collecting Windows System Logs: Log collection from a Windows system
   - Go to the specific AlienApp in USM Anywhere for instructions about syslog forwarding

   > **Note:** Because the log scan can take some time, you might not see all of the automatically discovered log sources immediately after deploying the first USM Anywhere Sensor.

2. When you have finished the log collection setup and integrated any needed plugins, verify that the data transfer is occurring.
3. Click **Next** when this step is complete.

OTX

AT&T Alien Labs™ Open Threat Exchange® (OTX™) is an open information-sharing and analysis network providing users with the ability to collaborate, research, and receive alerts on emerging threats and indicators of compromise (IoCs) such as IP addresses, file hashes, and domains.

You must have an OTX account to receive alerts based on threats identified in OTX. This account is separate from your USM Anywhere account. Go to The World's First Truly Open Threat Intelligence Community to create an OTX account.

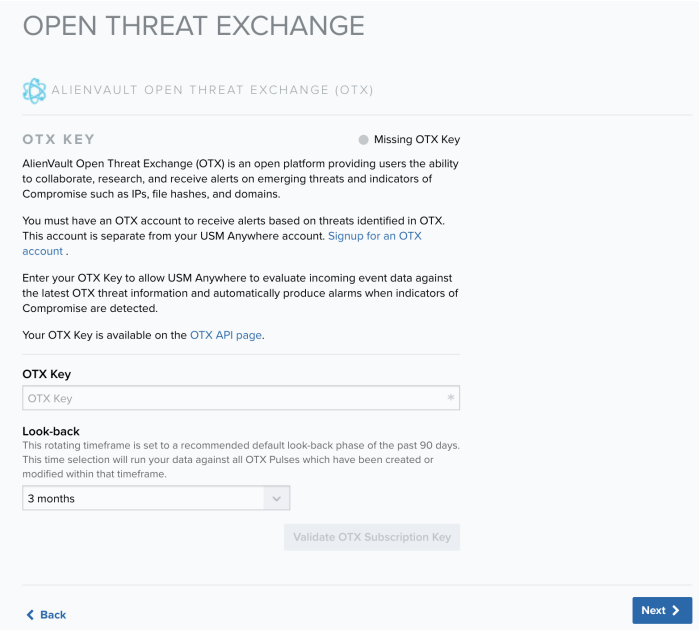

Note: If you do not already have an OTX account, click the **Sign up** link. This opens another browser tab or window that displays the OTX signup page. After you confirm your email address, you can log in to OTX and retrieve the unique API key for your account.

See Open Threat Exchange® and USM Anywhere in the *USM Anywhere User Guide* for more information about OTX integration in USM Anywhere.

To enable USM Anywhere to evaluate event data against the latest OTX intelligence

1. Log in to OTX and open the API page (https://otx.alienvault.com/api).
2. In the DirectConnect API Usage pane, click the icon to copy your unique OTX connection key.

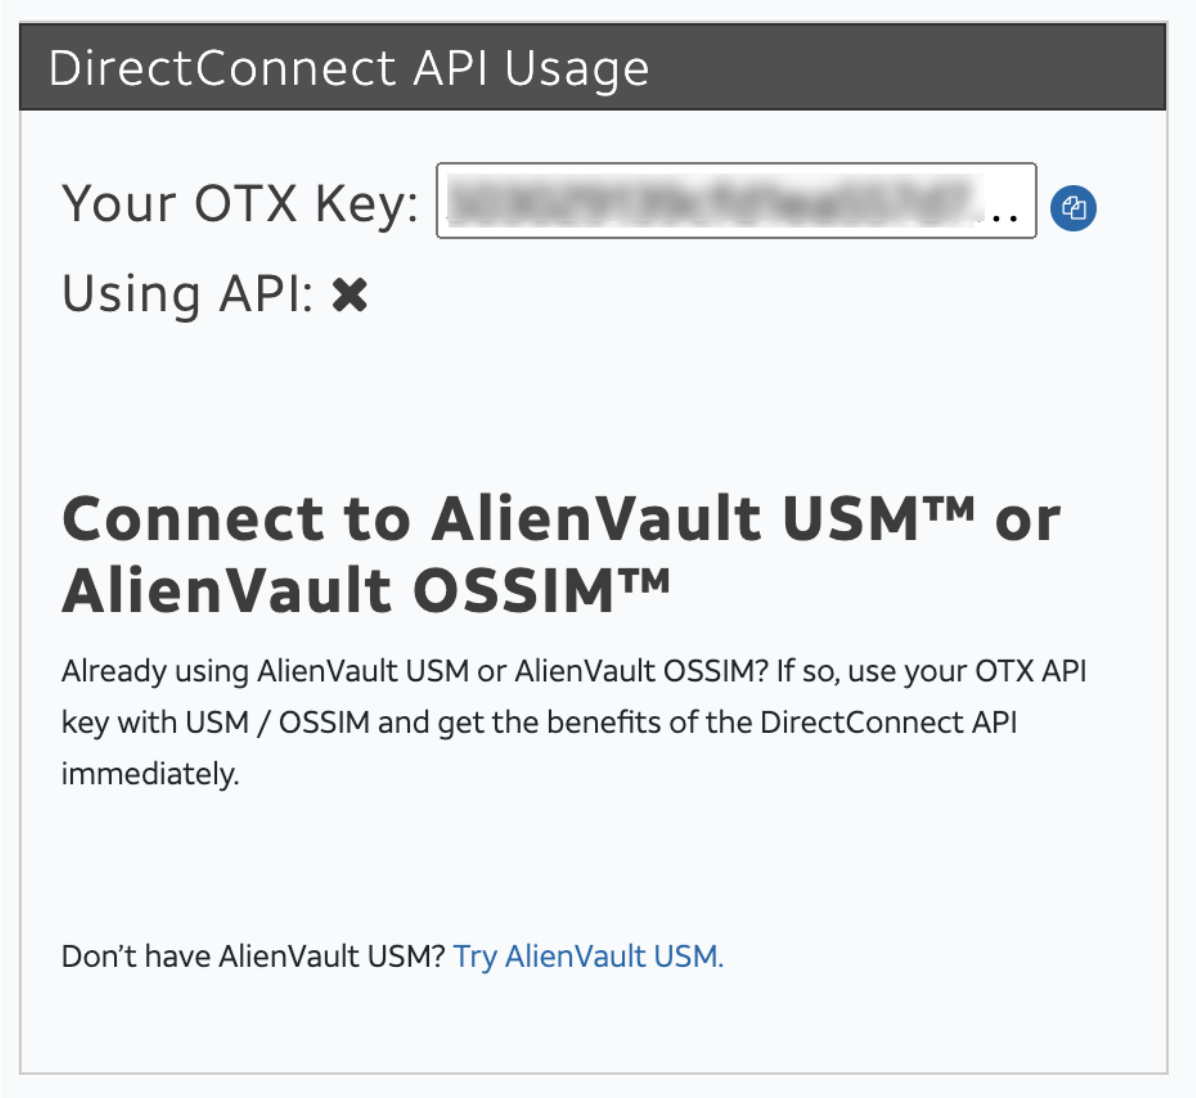

3. Return to the Open Threat Exchange (OTX) page of the USM Anywhere Sensor Setup Wizard and paste the value in the OTX Key text box.

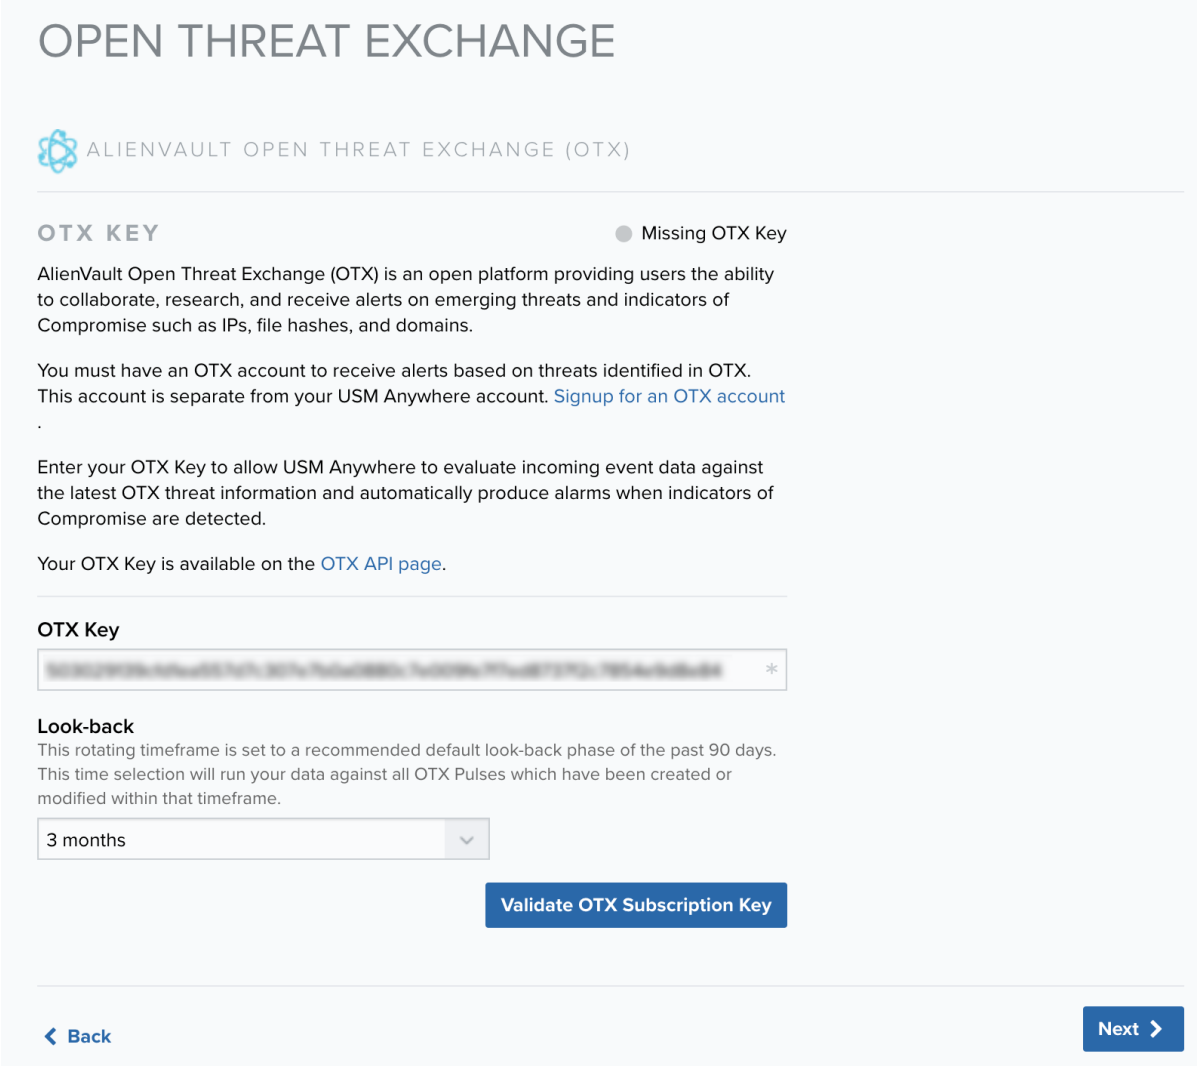

4. Click **Validate OTX Subscription Key**.

   With a successful validation of the key, the status at the top of the page changes to "Valid OTX key".

5. Click **Next** when this task is complete.

Setup Complete

The Congratulations page summarizes the status of your configuration.

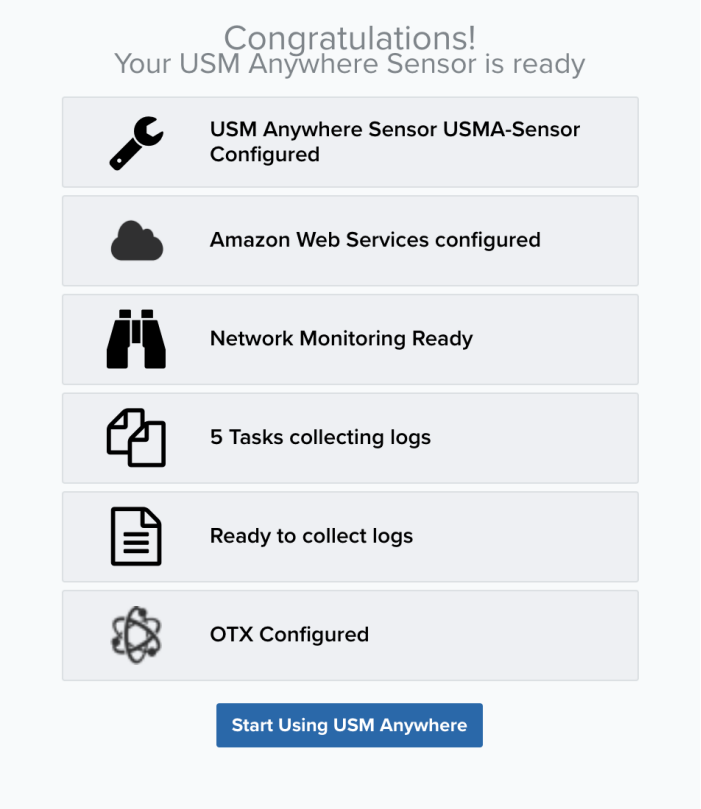

Click **Start Using USM Anywhere**, which takes you to the Overview dashboard.

Next...

Now is a great time to run a vulnerability scan. See Vulnerability Assessment in the *USM Anywhere User Guide* for detailed information about running a vulnerability scan.

Enable Connections in an AWS VPC

A USM Anywhere Sensor deployed in Amazon Web Services (AWS) to a virtual private cloud (VPC) automatically listens for syslog packets on UDP port 514, but you must enable access to it. This allows the other hosts in your network to send data to the sensor. You enable this access by opening this port using the AWS security groups that were created by the AWS CloudFormation template that you used to deploy the sensor.

The AWS Security Groups

There are five AWS security groups that help control network connectivity between the instances:

- **USMConnectionSG**: Accepts incoming HTTP, HTTPS, and SSH connections from the Classless Inter-Domain Routing (CIDR) block you specified when you completed the AWS CloudFormation template parameters.

  These connections are only required to enable remote sensor management, and to connect to the web user interface (UI) during deployment and setup.

- **USMLogServicesSG**: Accepts incoming UDP connections on port 514 from any virtual machine (VM) instance in the USMEnableLogServicesSG. It also enables syslog TCP on port 601, syslog Transport Layer Security (TLS) on port 6514, and Graylog UDP on port 12201.
- **USMEnableLogServicesSG**: Does not have inbound or outbound rules, nor is it assigned to the sensor.

  It exists solely as a convenience, so that you can assign it to VMs for connection to UDP over port 514 on the sensor as specified in the USMLogServicesSG. This also enables syslog TCP on port 601 and 602, syslog TLS on port 6514 and 6515, and Graylog UDP on port 12201 on that sensor.

- **USMTrafficInterfacesSG**: Enables USM traffic mirroring connectivity on your USM Sensor traffic network interface.
- **USMEnableTrafficMirroringSG**: Does not have inbound nor outbound rules, nor is it assigned to the sensor, but allows virtual extensible local area network (VXLAN) traffic over UDP port 4789.

  It exists solely as a convenience, so that you can assign it to VMs for connection to UDP over port 514 on the sensor as specified in the USMEnableTrafficMirroringSG. This enables traffic mirroring traffic.

UDP Port 514

You can open UDP port 514 to receive syslog packet transmissions from the AWS console using any *one* of the following methods:

- Assign the USMBaseSG security group to the selected VMs by navigating to **Networking > Change Security Groups action**. (You can also do this through the AWS command-line interface [CLI].)
- Add the default security group from your VPC to the USMLogServicesSG. This allows all the VMs in that security group to send to port 514 UDP.
- Put the AWS Sensor in the default security group from your VPC. This gives all of the VMs in the local VPC full access to all ports on the sensor.

VPC Traffic Mirroring with an AWS Sensor

With Amazon Web Services (AWS) Virtual Private Cloud (VPC) Traffic Mirroring, network traffic from your AWS environment can be mirrored and sent directly to your sensor for monitoring, bringing network intrusion detection system (NIDS) functionality to your AWS Sensor.

Important: AWS has a limit of 10 mirror sources per interface. In situations where you need to accommodate a higher number of mirror sources, you will need to deploy additional AWS sensors to mirror or create a NLB as VPC Target. See (Optional.) Create NLB as VPC Target for more information.

See Amazon's documentation on traffic mirroring for more information regarding mirroring.

Prepare Your Environment for VPC Traffic Mirroring

To enable VPC Traffic Mirroring, you must first ensure the following:

- Your sensor is at least an m5.xlarge
- Your sensor has a second network interface configured, called the traffic interface
- Your firewall rules allow virtual extensible local area network (VXLAN) traffic through your traffic interface (inbound UDP 4789 is applied)

Note: If you select **Enable Traffic Mirroring** when you first deploy (or redeploy) the sensor, these requirements are automatically set up for you. Otherwise, you must make all three of these changes manually.

Configuring VPC Traffic Mirroring

To run VPC Traffic Mirroring in your environment, you must create and configure the following:

- **Create a security group:** The security group configures the UDP protocol, and the inbound and outbound traffic.
- **Create a network interface:** The network interface specifies the subnet you want to monitor.
- **Attach the sensor:** Attach the sensor to the configured information.
- **Create a target:** This target serves as the destination for your mirrored traffic.
- **Define a mirror filter:** The mirror's filter specifies which traffic is mirrored for your AWS Sensor.
- **Create the mirror session:** This session configures precisely how your traffic is mirrored.

To create a security group

1. Go to **Network & Security > Security Groups** in your AWS Management Console.
2. Click **Create security group**.
3. Enter the following:
   - **Security group name:** A name for this security group
   - **Description:** A description of this security group
   - **VPC:** A VPC for this security group
   - **UDP Protocol:** Enter the 4789 port
   - **Inbound:** The traffic must be admitted
   - **Outbound:** Select all traffic, all protocol, all port range, and 0.0.0.0/0 as the destination

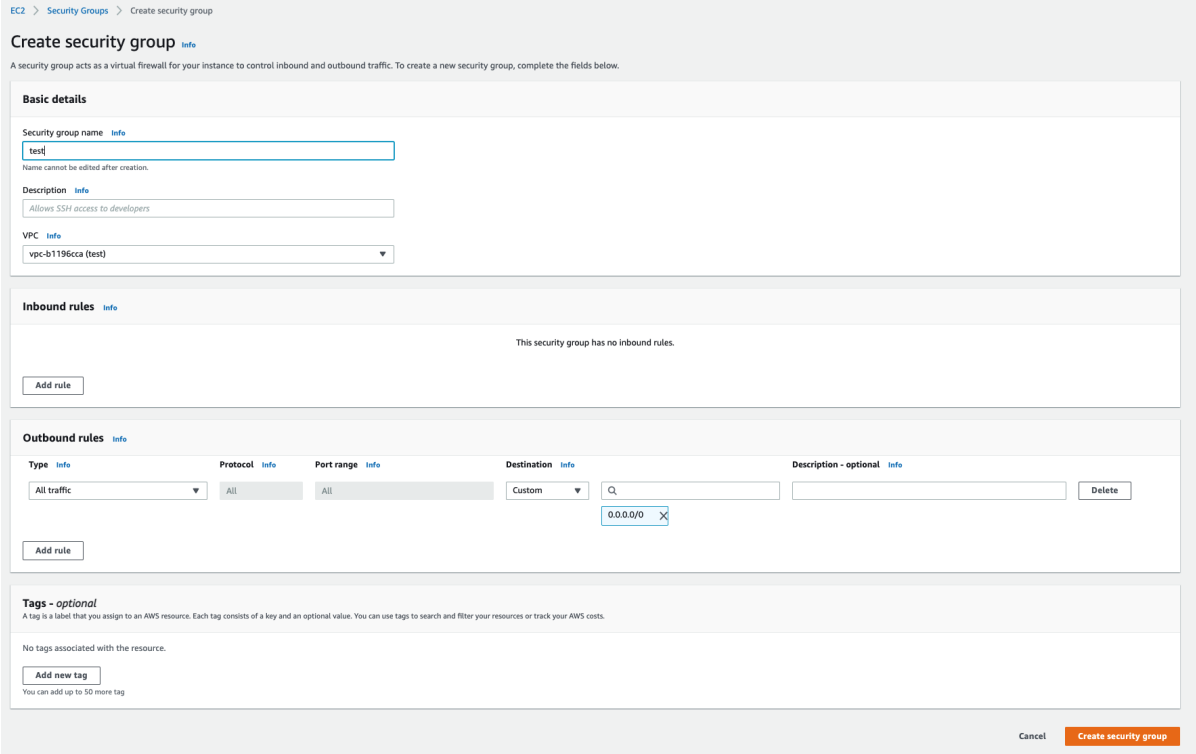

4. Click **Create security group**.

To create a network interface

1. Go to **Network & Security > Network Interfaces** in your AWS Management Console.
2. Click **Create Network Interface**.
3. Enter the following:
   - (Optional.) **Description:** A description of this network interface
   - (Optional.) **Subnet:** Select the subnet you want to monitor
   - **Security group:** Select the security group you've created

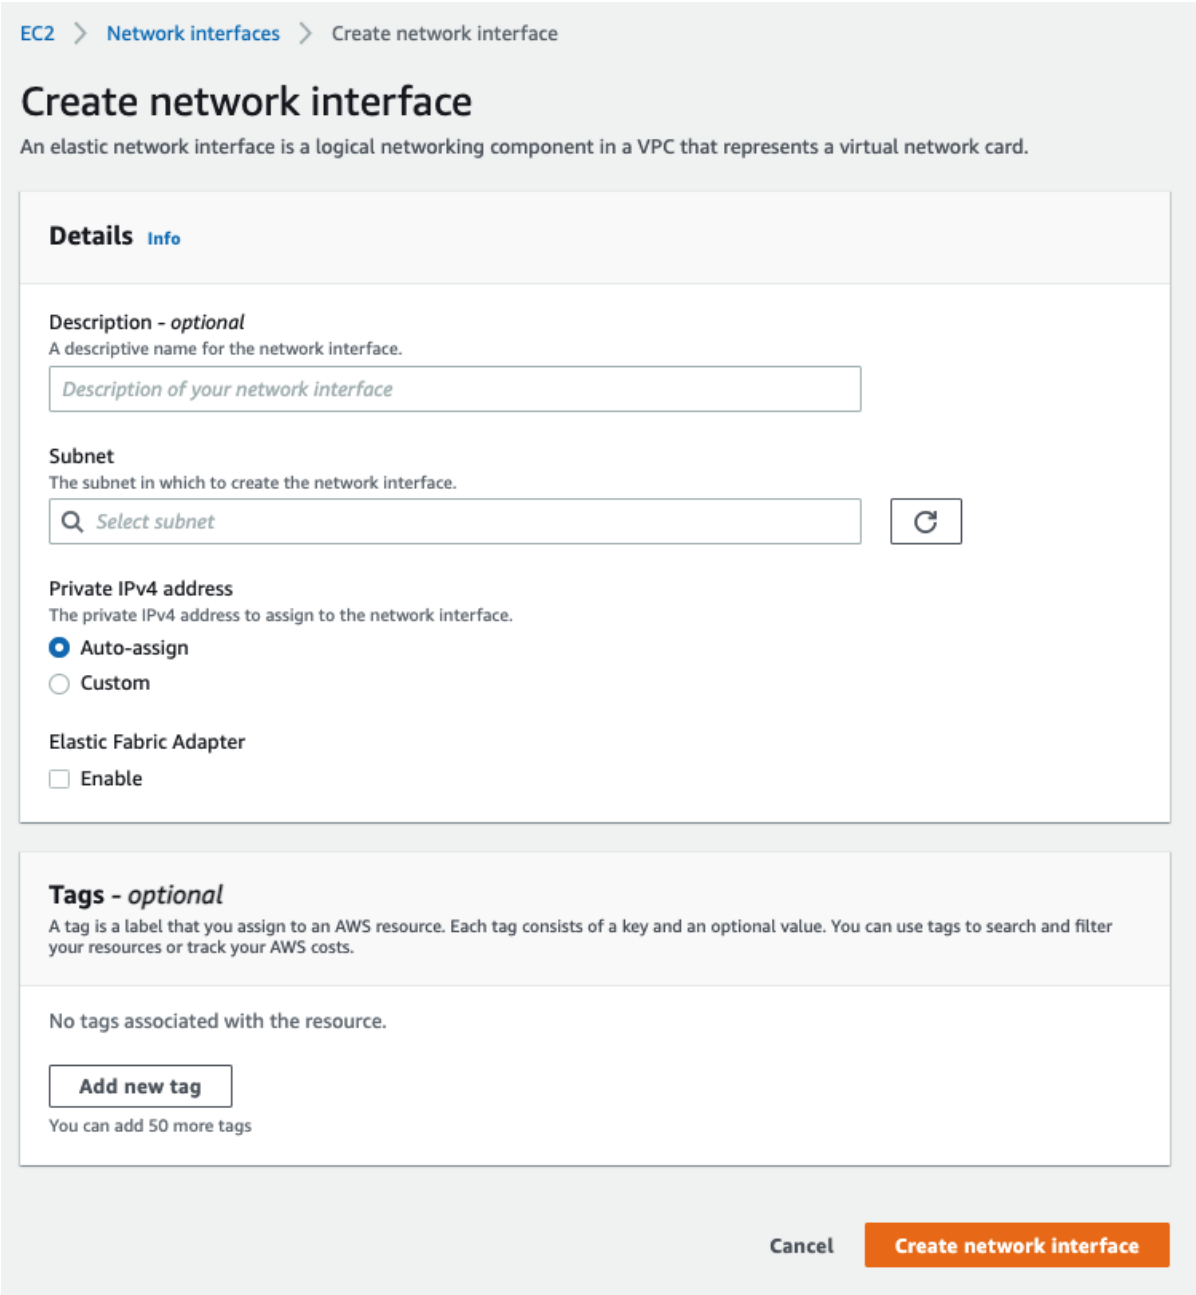

4. Click **Create Network Interface**.

To attach the sensor

1. Go to **Network & Security > Network Interfaces** in your AWS Management Console.
2. Click **Actions > Attach**.

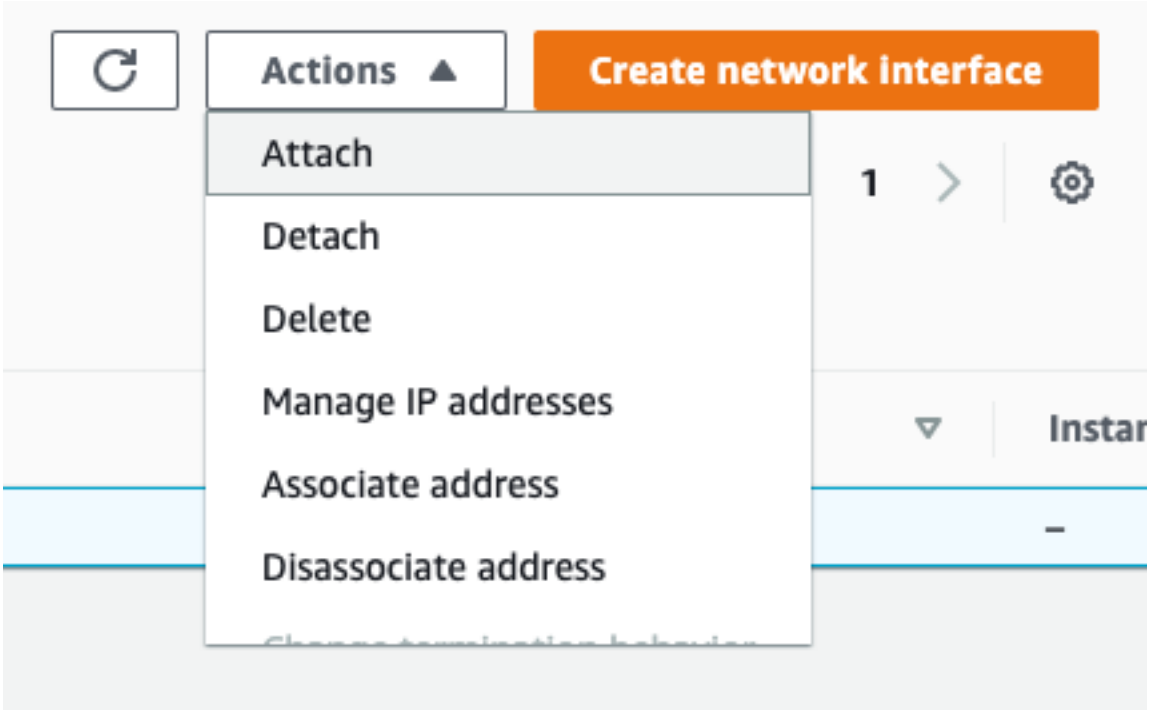

3. Select your instance.
4. Click **Attach**.
5. Reload the sensor.

To create a target

1. Go to **Networking & Content Delivery > VPC** in your AWS Management Console and click **Mirror Targets** under Traffic Mirroring.
2. Click **Create traffic mirror target**.
3. Enter the following:
   - (Optional.) **Name tag:** A name for this target
   - (Optional.) **Description:** A description of this target
   - **Target Type:** This value must be **Network Interface**
   - **Target:** The identification (ID) of your instance's traffic interface

4. (Optional.) If you use tags to organize your AWS resources, you can add tags to this target.
5. Click **Create**.

> **Important:** Each elastic network interface (ENI) has a three session limit per ENI. If you use a Network Load Balancer (NBL) as your mirror target, there is no session limit. See (Optional.) Create NLB as VPC Target for more information.

To define the mirror filter

This filter defines what traffic is mirrored to your AWS Sensor. You can specify inbound and outbound filters, as well as applying filters for Amazon network services.

1. Go to **Networking & Content Delivery > VPC** in your AWS Management Console and click **Mirror Filters** under Traffic Mirroring.
2. Click **Create traffic mirror filter**.
3. (Optional.) Enter the following information:
   - **Name tag:** A name for your traffic mirror filter
   - **Description:** A description for your traffic mirror filter
   - **Network services:** Select this checkbox to have your filter mirror network services

data

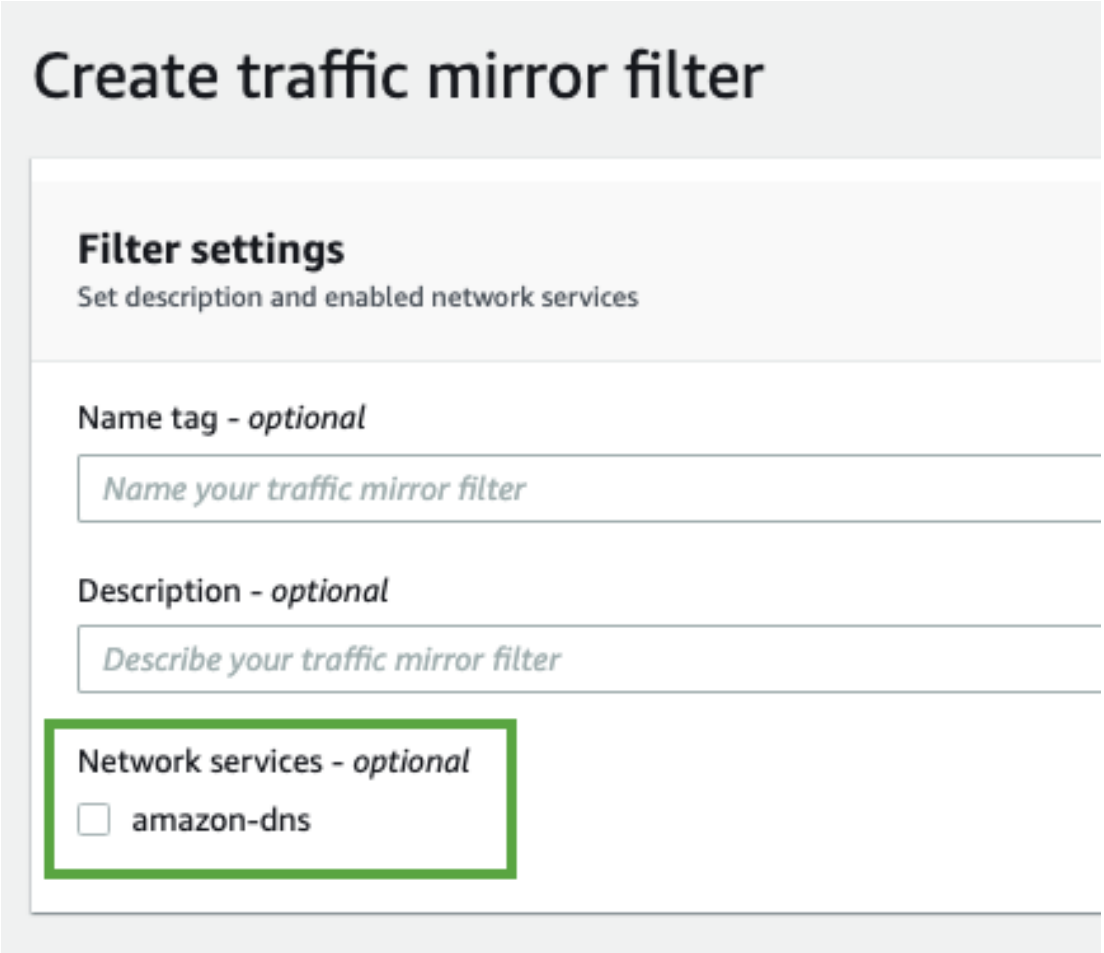

Note: If you want your filter to capture Domain Name System (DNS) traffic, you must select the **amazon-dns** checkbox.

4. Configure your inbound and outbound filtering rules.
   a. Click **Add rule**.
   b. Use the options provided to define these rules:
      - **Number:** Priority settings to order which rules are evaluated before others
      - **Rule Action:** Specify the action (accept or reject) to take for the filtered packet
      - **Protocol:** Specify one protocol to collect specific traffic, or select **All protocols** to collect all traffic
      - (Optional.) **Source Port Range:** Specify the source port range you want to filter

- (Optional.) **Destination Port Range:** Specify the destination port range you want to filter
- **Source CIDR Block:** Specify the source IP ranges you want to filter
- **Destination CIDR Block:** Specify the destination IP ranges you want to filter
- (Optional.) **Description:** A description for this filtering rule

5. (Optional.) If you use tags to organize your AWS resources, you can add tags to this filter.
6. Click **Create**.

To create a session

1. Go to **Networking & Content Delivery > VPC** in your AWS Management Console and click **Mirror Sessions** under Traffic Mirroring.
2. Click **Create traffic mirror session**.
3. Enter the following information:
   - (Optional.) **Name tag:** A name for your traffic mirror session
   - (Optional.) **Description:** A description for your traffic mirror session
   - **Mirror Source:** The network interface ID of the instance you want to monitor
   - **Mirror Target:** The ID of your instance's traffic interface
   - **Session Number:** Priority settings to order which sessions are evaluated before others
   - **VNI:** Set this to 1169
   - (Optional.) **Packet length:** The number of bytes from each packet to mirror

     (AT&T Cybersecurity recommend leaving this blank to mirror the entire packet)
   - **Filter:** The filter you have created for your VPC Traffic Mirroring session

   Important: You must select a supported instance type to successfully create a mirror session. See Amazon's documentation on traffic mirroring limitations for a list of supported and unsupported instance types.
4. (Optional.) If you use tags to organize your AWS resources, you can add tags to this session.
5. Click **Create**.

Important: You must create one mirror session per device.

(Optional.) Create NLB as VPC Target

You can create a Network Load Balancer (NLB) to use as your VPC target. To create an NLB, follow Amazon's Create a Network Load Balancer instructions, and for the Target Group and Health Check portions of the setup, use the settings provided below.

When you configure the Target Group section, use the following settings:

- **Target Type:** IP
- **Protocol:** UDP
- **Port:** 4789

When you configure the Health Checks section, use the following settings:

- **Protocol:** HTTP
- **Path:** /api/2.0/status

Advance Health Check Settings:

- **Override:** Port - 80

Collect Logs from Amazon S3 Buckets with KMS Encryption

Role Availability	Read-Only	Investigator	Analyst	Manager

If you are using a key management service (KMS) key to encrypt the Amazon S3 buckets where your logs are stored, you need to perform the following steps to enable your USM Anywhere Sensor to decrypt those buckets.

> **Note:** To do this, you first need to know the bucket that is encrypted, the KMS key used for the encryption, and the Identity and Access Management (IAM) role created for your sensor.

To enable your sensor to decrypt KMS-encrypted buckets

1. Log in to the AWS Management Console and navigate to the Key Management Service (KMS) page.
2. Open the Customer Managed Keys page and locate the KMS key you are using.
3. Scroll down to the *Key Users* section.

4. Click **Add**.
5. Use the list or the search bar to select the IAM role created for your sensor.

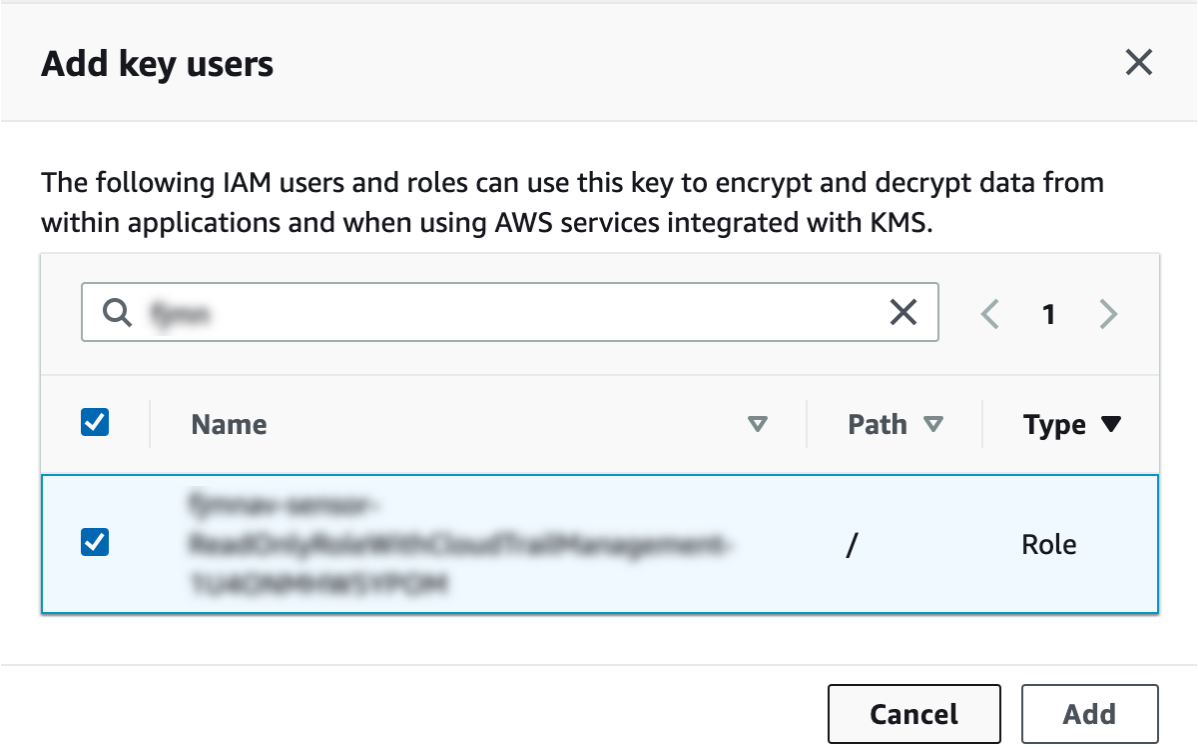

6. Click **Add**.

AWS Log Discovery and Collection in USM Anywhere

Amazon Web Services (AWS) customers have access to service-specific log files to gain insight into how each AWS service is operating. In addition, applications running in AWS also generate various log files in different formats. With a deployed AWS Sensor, USM Anywhere can collect both logs from AWS, but the procedures are slightly different:

- Use a predefined scheduler job

  USM Anywhere automatically discovers the AWS CloudTrail logs, the Amazon Simple Storage Service (S3) access logs, and some Amazon CloudWatch logs when they are enabled within your AWS account. There are predefined scheduler jobs in USM Anywhere to collect these logs but they are disabled by default. Go to **Settings > Scheduler > Log**

Collection for the full list. You need to enable each job based on which log you want to collect. See Collect AWS CloudTrail Logs on an AWS Sensor, Collect Amazon S3 Access Logs and Collect ELB Access Logs for more information.

- Use a customer-defined scheduler job

  If none of the predefined jobs collect from your log location, you can create a new job under **Settings > Scheduler > Log Collection**. Depending on where your logs are stored, USM Anywhere provides two ways to collect them:

  - **Amazon CloudWatch Logs:** If you choose to use Amazon CloudWatch Logs in your AWS environment, USM Anywhere can collect CloudWatch logs directly. See Collect AWS CloudTrail Logs on an AWS Sensor for more information. For example, you can collect the Amazon Virtual Private Cloud (VPC) flow logs using this method.
  - **Amazon S3 bucket:** If you choose to store logs in an Amazon S3 bucket instead, USM Anywhere can also collect logs directly from an Amazon S3 bucket. See Collect Other Logs from an Amazon S3 Bucket for more information.

Configure Amazon GuardDuty for the AWS Sensor

Role Availability	Read-Only	Investigator	Analyst	Manager

You can leverage your Amazon GuardDuty service within the AWS Sensor to translate the raw log data into normalized events for analysis.

Amazon GuardDuty service is automatically detected when a new AWS Sensor is deployed. However, it still needs to be enabled for USM Anywhere to receive information from it.

To enable Amazon GuardDuty for your AWS Sensor

1. Go to **Settings > Scheduler**.
2. Search for **GuardDuty** in the Job Scheduler **Filter By** field.
3. In the row for the GuardDuty job, click icon.

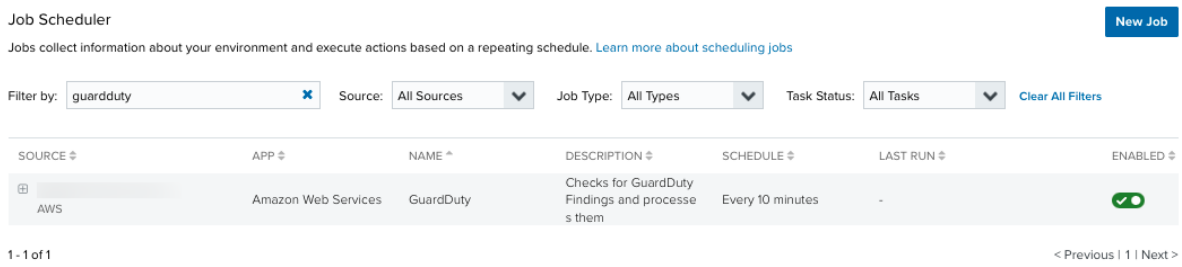

Collect AWS CloudTrail Logs on an AWS Sensor

Amazon Web Services (AWS) CloudTrail provides a complete audit log for all actions taken with the Amazon API, either through the web user interface (UI), the AWS Command Line Interface (CLI), or an AWS software development kit (SDK). Ongoing monitoring of this log gives you visibility of end user and automated actions in your environment. This helps you quickly detect abuse cases and security incidents, such as a user trying to make changes to an AWS account that are inconsistent with their privileges.

USM Anywhere automatically detects AWS CloudTrail and retrieves your AWS CloudTrail logs across all regions within a single AWS account. USM Anywhere also provides you the credentials to securely access your AWS CloudTrail logs. When a new trail is detected, a new log collection job is automatically created and enabled to capture the logs in that trail. Similarly, if a trail is deleted, the existing job that was created for it is automatically deleted.

As the AWS Sensor collects this raw log data, USM Anywhere uses its AWS CloudTrail data source to normalize the data and generate meaningful events. Depending on the size and activity in your AWS account, this log collection can produce an excessive number of events. See Managing Collected CloudTrail Event Logs for a list of possible CloudTrail events. Similarly, if your AWS instance includes organizations, you may create a trail that will log all events for any AWS accounts assigned to an organization.

Note: If you choose not to enable AWS CloudTrail, USM Anywhere processes all stored logs at initial startup. See the Amazon documentation for information about enabling AWS CloudTrail. After that initial processing, log collection jobs run every five minutes to ensure that logs are captured and can generate meaningful events in a timely manner.

Note: Sometimes you may see that the CloudTrail events in USM Anywhere display a different username compared to the raw log. This is because CloudTrail provides different types of user identities, one of which is *AssumedRole*. When the user identity type is set to AssumedRole, it means that the user credential is temporary and the username you see in the raw log is not the actual username. See Amazon documentation for more information.

To enable AWS CloudTrail for your AWS Sensor

1. Go to **Settings > Scheduler**.
2. Search for **CloudTrail**in the Job Scheduler **Filter By** field.
3. In the row for the CloudTrail job, click the icon to enable the AWS CloudTrail jobs.

   This turns the icon green.

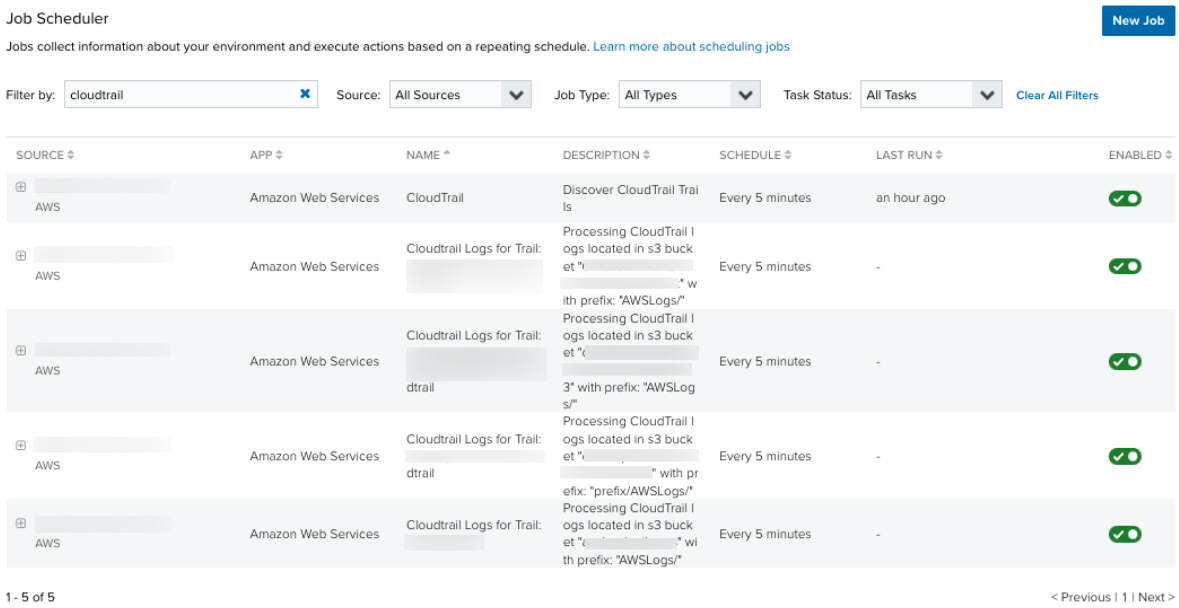

Collect Amazon CloudWatch Logs

Role Availability ✗ Read-Only ✗ Investigator ✓ Analyst ✓ Manager

Amazon CloudWatch Logs monitors applications and systems using log data, aggregating and storing application logs. CloudWatch Logs is useful because you can easily configure it to process additional metadata with the log files. Visit the AWS documentation to learn more about VPC flow log collection.

Important: If you choose to enable CloudWatch Logs in your Amazon Web Services (AWS) environment, you should make sure that you are not collecting more data than you need because this service incurs AWS costs based upon usage. See the CloudWatch pricing information to plan and configure your usage.

If not already done, install and configure the Amazon CloudWatch agent to collect logs from Amazon Elastic Compute Cloud (EC2) instances. See Amazon documentation for instructions.

USM Anywhere provides some CloudWatch log collection jobs out of the box, but they are disabled by default. You can enable them under **Settings > Scheduler**. When enabled, these jobs monitor certain log groups and collect logs from CloudWatch every five minutes. You must configure your CloudWatch agent to use these log group names and to keep the log types the same within a given log group.

USM Anywhere Log Collection Jobs and CloudWatch Log Groups

USM Anywhere Log Collection Job Name	CloudWatch Log Group Name	Default File Path	Date Format
CloudWatch - Apache-Access-Logs	Apache-Access-Logs	/var/log/apache2/access.log	%d/%b/%Y:%H:%M:%S
CloudWatch - Linux-Audit-Logs	Linux-Audit-Logs	/var/log/audit/audit.log	Use the default

USM Anywhere Log Collection Jobs and CloudWatch Log Groups (Continued)

USM Anywhere Log Collection Job Name	CloudWatch Log Group Name	Default File Path	Date Format
CloudWatch - Linux-Auth-Logs	Linux-Auth-Logs	/var/log/auth.log	%b %d %H:%M:%S
CloudWatch - Osquery-Logs	OSQuery-Logs	/var/log/osquery/osqueryd.results.log	Use the default

If you want to collect logs from other log groups, ensure that all streams in the same group are of the same type so that USM Anywhere can use a designated data source to parse the collected raw log data. You can then set up a CloudWatch log collection job for each log group.

To create a new CloudWatch log collection job

1. Go to **Settings > Scheduler**.
2. In the left navigation menu, click **Log Collection**.

> **Note:** You can use the Sensor filter at the top of the list to review the available log collection jobs on your AWS Sensor.

3. Click **Create Log Collection Job**.

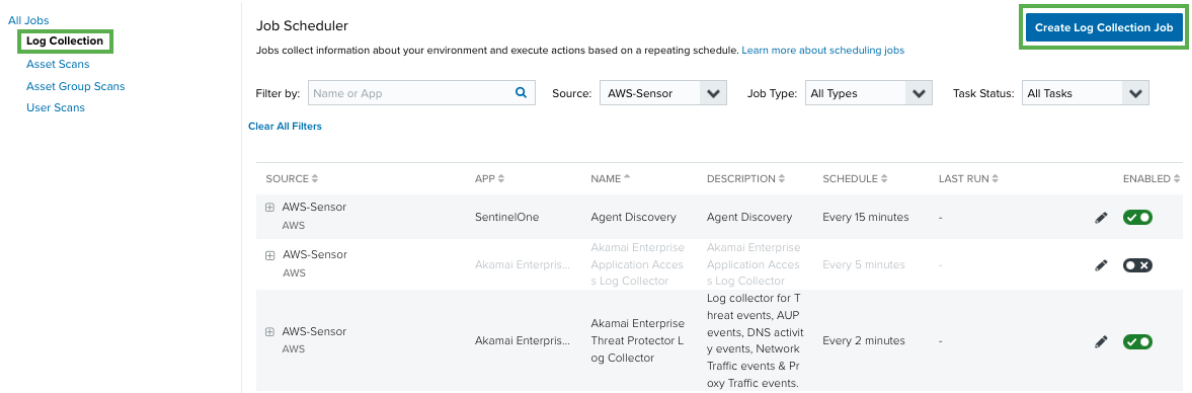

Note: If you have recently deployed a new USM Anywhere Sensor, it can take up to 20 minutes for USM Anywhere to discover the various log sources. After it discovers the logs, you must manually enable the AWS log collection jobs you want before the system collects the log data.

The Schedule New Job dialog box opens.

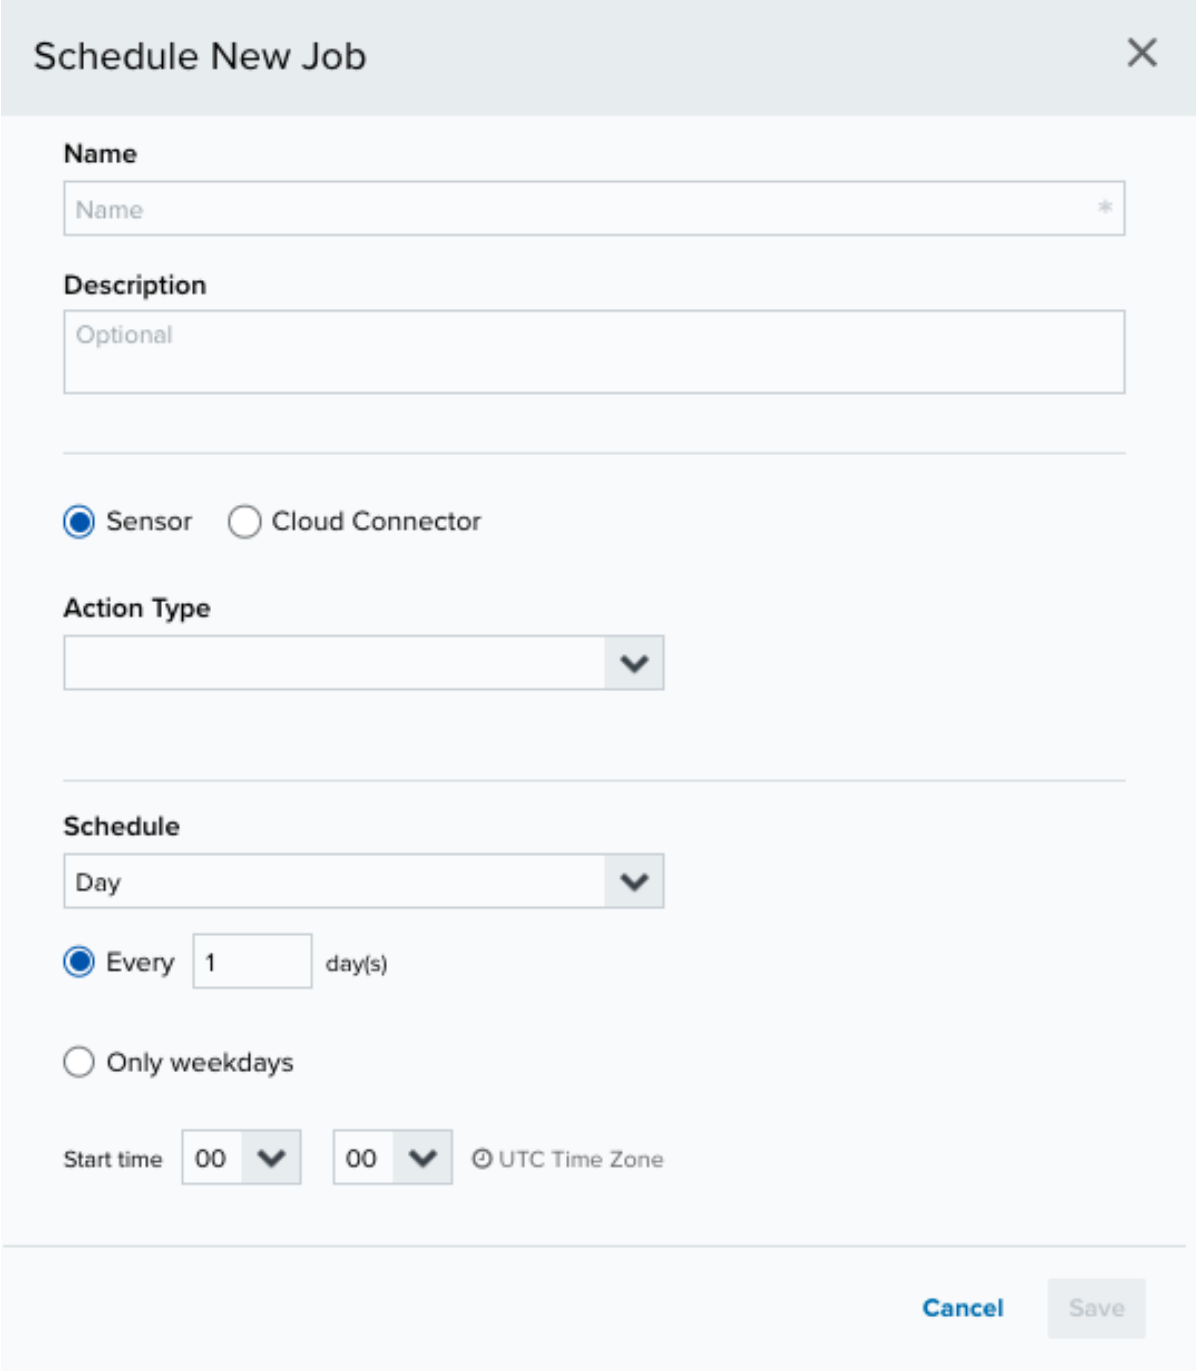

1. Enter the name and description for the job.

   The description is optional, but it is a best practice to provide this information so that others can easily understand what it does.

2. Select **Sensor** as the source for your new job.
3. In the Action Type drop-down list, select **Amazon Web Services**.

4. In the App Action drop-down list, select **Monitor CloudWatch**.

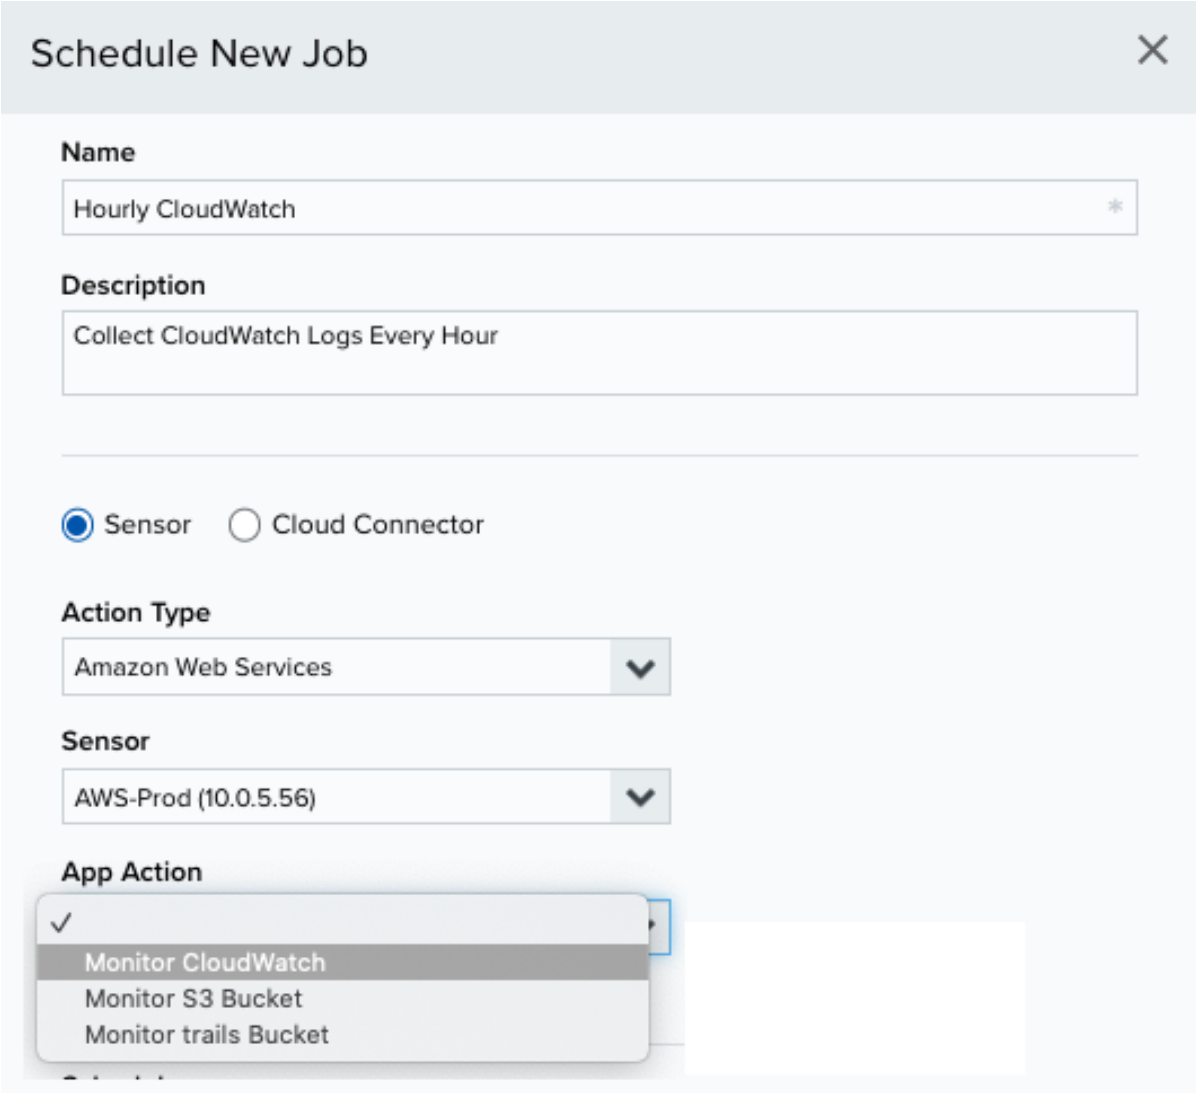

5. Enter the **Region Name**, **Group Name**, and **Stream Name** information for your AWS account. Region name can be an asterisk (*****) to monitor all regions for a given group.

6. In Source Format, select either of the following log formats:

   - **Syslog:** All messages transmitted to USM Anywhere are processed with the assumption that they are syslog formatted.

     When you choose syslog as the source format, the data source selection is bypassed and USM Anywhere uses the auto-detect hints from the data sources to match the incoming messages to the correct data source.

- **Raw:** Use for non-syslog formatted data.

  If you select this option, you must choose the data source that USM Anywhere will use to parse all of the streams in the group. For example, to collect Amazon Virtual Private Cloud (VPC) flow logs, select the **VPC Flow Logs** data source.

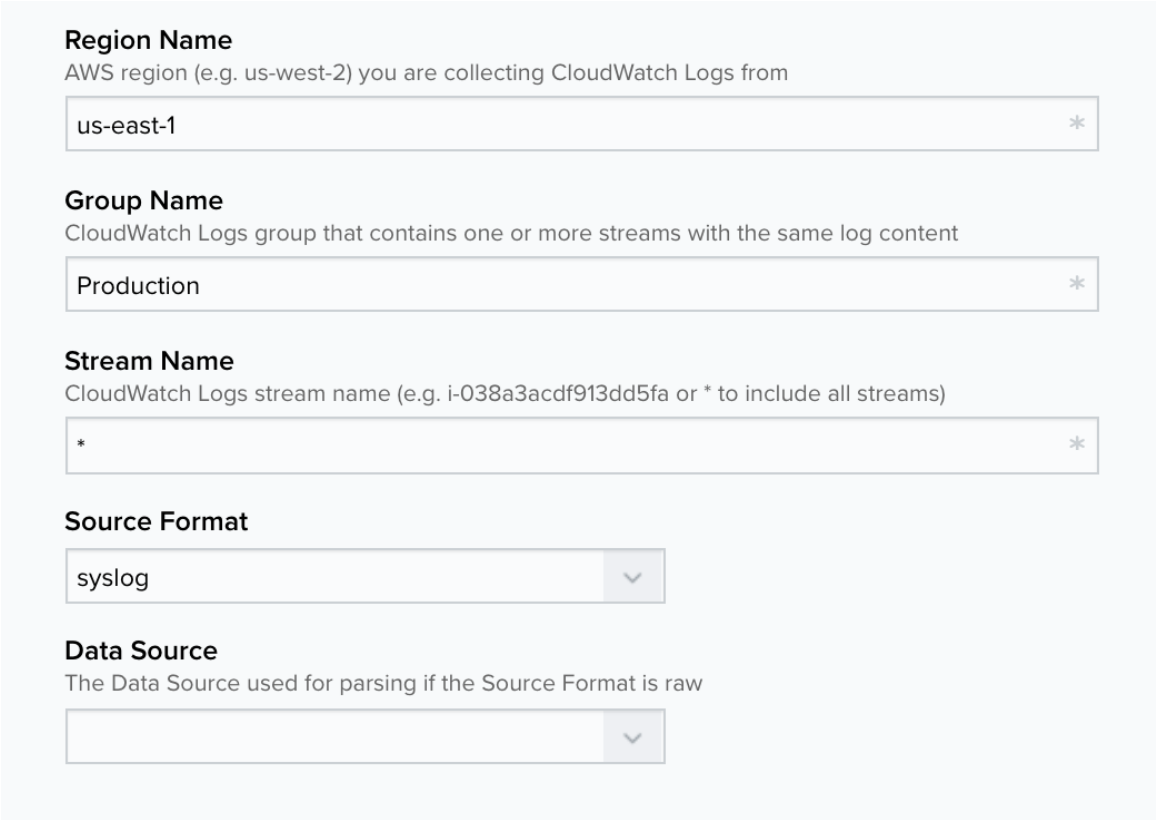

Important: If a group contains streams of mixed log formats, USM Anywhere parses all of them with the data source that you chose, which produces undesired results. In this case, you need to configure CloudWatch to separate the streams into different groups so that each contains only a single log type that can be mapped to the correct data source.

7. In the Schedule section, specify when USM Anywhere runs the job:

   a. Select the increment as **Minute**, **Hour**, **Day**, **Week**, **Month**, or **Year**.

   Warning: After a frequency change, monitor the system to check its performance. For example, you can check the system load and CPU. See USM Anywhere System Monitor for more information.

b. Set the interval options for the increment.

The selected increment determines the available options. For example, on a weekly increment, you can select the days of the week to run the job.

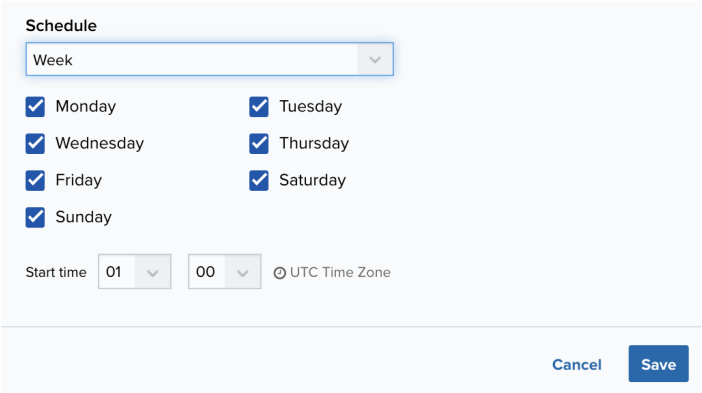

Or on a monthly increment, you can specify a date or a day of the week that occurs within the month.

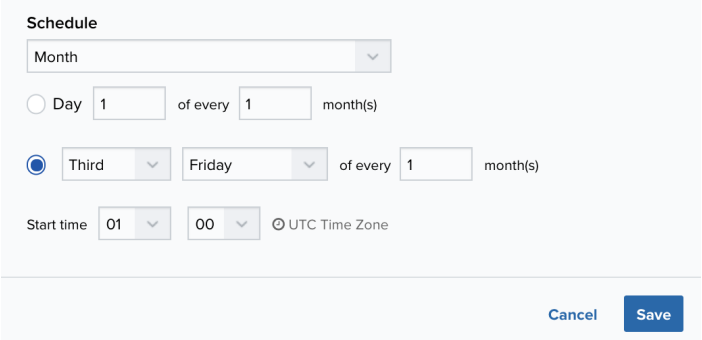

Important: USM Anywhere restarts the schedule on the first day of the month if the option "Every x days" is selected.

c. Set the start time.

This is the time that the job starts at the specified interval. It uses the time zone configured for your USM Anywhere instance (the default is Coordinated Universal Time [UTC]).

8. Click **Save**.

USM Anywhere detects any enabled jobs with the same configuration and asks you to confirm before continuing. This is because having two jobs with the same configuration generates duplicate events and alarms.

Collect Amazon S3 Access Logs

Role Availability	Read-Only	Investigator	Analyst	Manager

Amazon Simple Storage Service (S3) is object storage with a simple web service interface that you can use to store and retrieve any amount of data from anywhere on the web. Organizations running an Amazon Web Services (AWS) environment typically use it as the primary storage for their cloud-native applications, as a bulk repository, as a target for backup and recovery, and as a long-term archive location.

When enabled, Amazon S3 can provide complete access logs for all actions taken in an Amazon S3 bucket. This gives you insight into who is accessing the data, and what actions are being taken. See Amazon's documentation to learn how to enable S3 access logging.

Note: In AWS, you must enable Amazon S3 access logging in every Amazon S3 bucket that you want to monitor.

With a deployed AWS Sensor, USM Anywhere automatically discovers the Amazon S3 access logs when you have enabled them within your AWS account. All you need to do is to enable the log collection job in USM Anywhere.

To enable Amazon S3 access logs collection in USM Anywhere

1. Go to **Settings > Scheduler**.
2. In the left navigation pane, click **Log Collection**.
3. Locate the **Discover S3 buckets** job and click the icon.

   This turns the icon green ( ). To disable an already-enabled job, toggle the icon to its

original status.

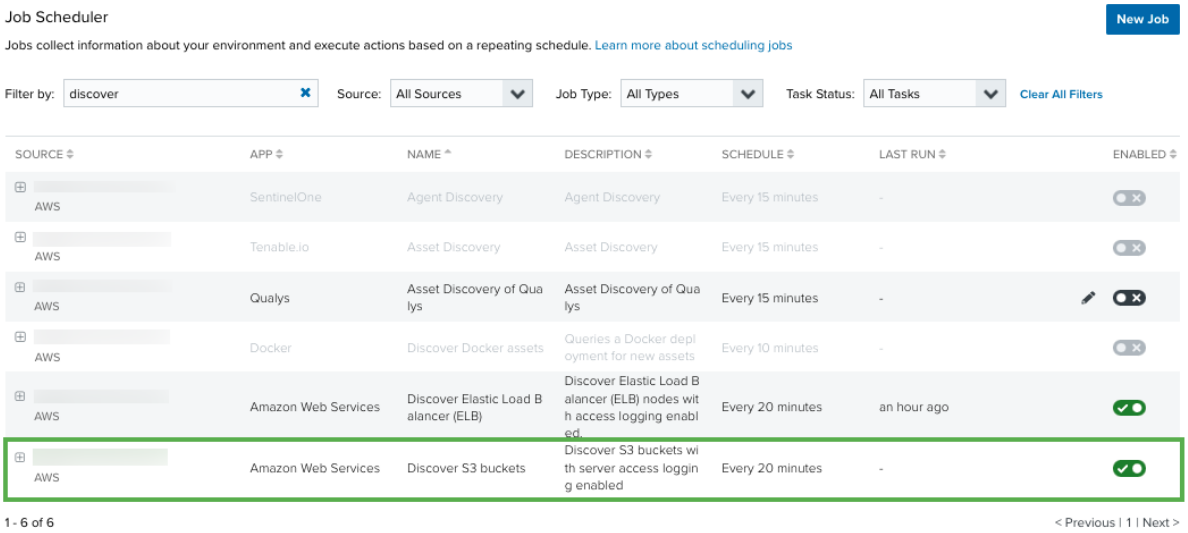

After you have enabled log collection, USM Anywhere automatically discovers your Amazon S3 access logs every 20 minutes. They will now begin generating events and you can see them in the Amazon S3 Dashboard.

Collect ELB Access Logs

Role Availability ✗ Read-Only ✗ Investigator ✓ Analyst ✓ Manager

Elastic Load Balancing (ELB) is an important feature in Amazon Web Services (AWS) because it automatically distributes incoming application traffic across multiple targets. AWS ELB access logs provide insight into who is accessing your web resources. They also help you identify common abuse patterns and use of automated hacking tools such as web application scanners.

USM Anywhere supports log discovery in two types of load balancers:

- **AWS Application Load Balancer**: You must enable Application Load Balancer logs for every AWS ELB that you want to monitor. See the Amazon documentation to learn how to enable Application Load Balancer access logging in AWS.
- **AWS Classic Load Balancer**: You must enable Classic Load Balancer logs for every AWS ELB that you want to monitor. See the Amazon documentation to learn how to enable Classic Load Balancer access logging in AWS.

Collecting AWS Application Load Balancer Access Logs

Once you have enabled Application Load Balancer access logging in AWS, you must also configure a scheduled job to monitor the Amazon Simple Storage Service (S3) bucket for the AWS Application Load Balancer. Only after this has been completed will USM Anywhere be able to automatically discovery your ELB access logs.

To create an AWS Application Load Balancer access log collection in USM Anywhere

1. Go to **Settings > Scheduler**.
2. Click **New Job**.
3. Configure your new scheduled job to collect access logs
   - **Action Type**: Amazon Web Services
   - **App Action**: Monitor S3 Bucket
   - **Bucket Name**: The name of the S3 bucket you want to monitor
   - **Path**: The prefix for the path you want to monitor
   - **Source Format**: Specify whether the source is raw or syslog
   - **Data Source**: AWS Application Load Balancer
4. Set a schedule for your new scheduled job.
5. Click **Save**.

After you have enabled your new job, USM Anywhere will use this job to discover your AWS Application Load Balancer access logs on the schedule you chose. These logs will now begin generating events and you can see them in the AWS Load Balancer Dashboard.

Collecting AWS Classic Load Balancer Access Logs

The AWS Sensor automatically detects Classic Load Balancer access logs after you have enabled them in AWS. After they're enabled in AWS, all you need to do is to enable the log collection job in USM Anywhere.

To enable AWS Classic Load Balancer access log collection in USM Anywhere

1. Go to **Settings > Scheduler**.
2. In the left navigation pane, click **Log Collection**.
3. Locate the Discover Elastic Load Balancer (ELB) job and click the icon.

This turns the icon green ( ). To disable an already-enabled job, toggle the icon to its original status.

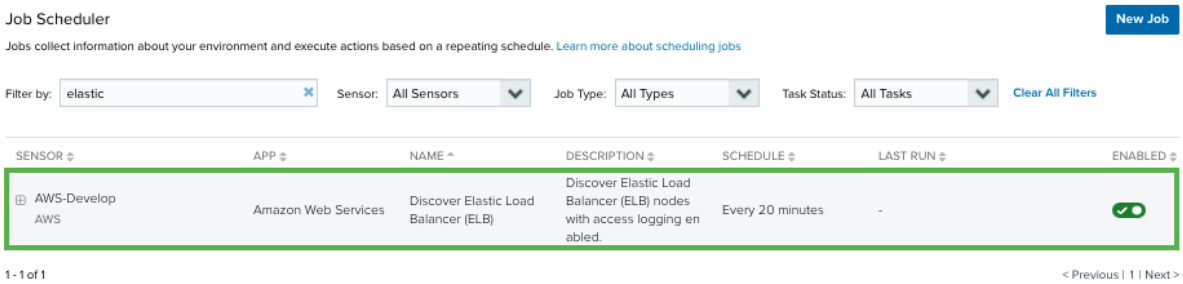

After you have enabled log collection, USM Anywhere automatically discovers your AWS Classic Load Balancer access logs every 20 minutes. They will now begin generating events and you can see them in the AWS Load Balancer dashboard.

Collect Other Logs from an Amazon S3 Bucket

Role Availability ✗ Read-Only ✗ Investigator ✓ Analyst ✓ Manager

In addition to the native service-specific logging that Amazon Web Services (AWS) provides, individual applications you run in the AWS environment often generate their own log files. You can forward these logs to an Amazon Simple Storage Service (S3) bucket and configure USM Anywhere to collect logs from that Amazon S3 bucket. USM Anywhere does not restrict the number of logs you can collect, but AWS does set limits on the number of logs it can return in each operation.

For example, to collect logs from AWS Web Application Firewall (WAF), you first need to follow AWS documentation to configure AWS WAF logging to store logs in an Amazon S3 bucket. Then configure a scheduler job in USM Anywhere to collect logs from the bucket.

Note: USM Anywhere accepts any file type when collecting log files. For compressed files, it looks for the file extension .gz, .zip, or .bz2 and uses the standard java.util or Apache Commons library to read the files. All other files are read as plain text.

To collect logs from an Amazon S3 bucket

1. Go to **Settings > Scheduler**.
2. In the left navigation menu, click **Log Collection**.

> **Note:** You can use the Sensor filter at the top of the list to review the available log collection jobs on your AWS Sensor.

3. Click **Create Log Collection Job**.

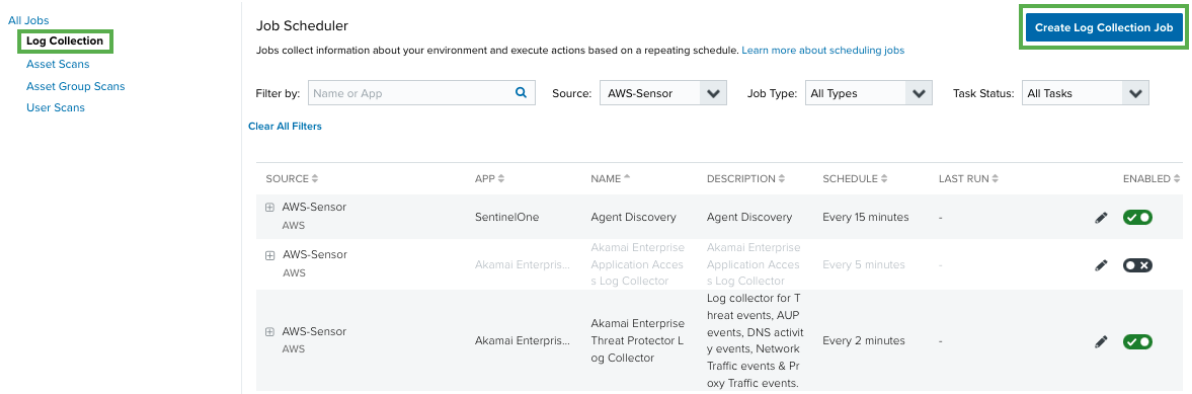

> **Note:** If you have recently deployed a new USM Anywhere Sensor, it can take up to 20 minutes for USM Anywhere to discover the various log sources. After it discovers the logs, you must manually enable the AWS log collection jobs you want before the system collects the log data.

The Schedule New Job dialog box opens.

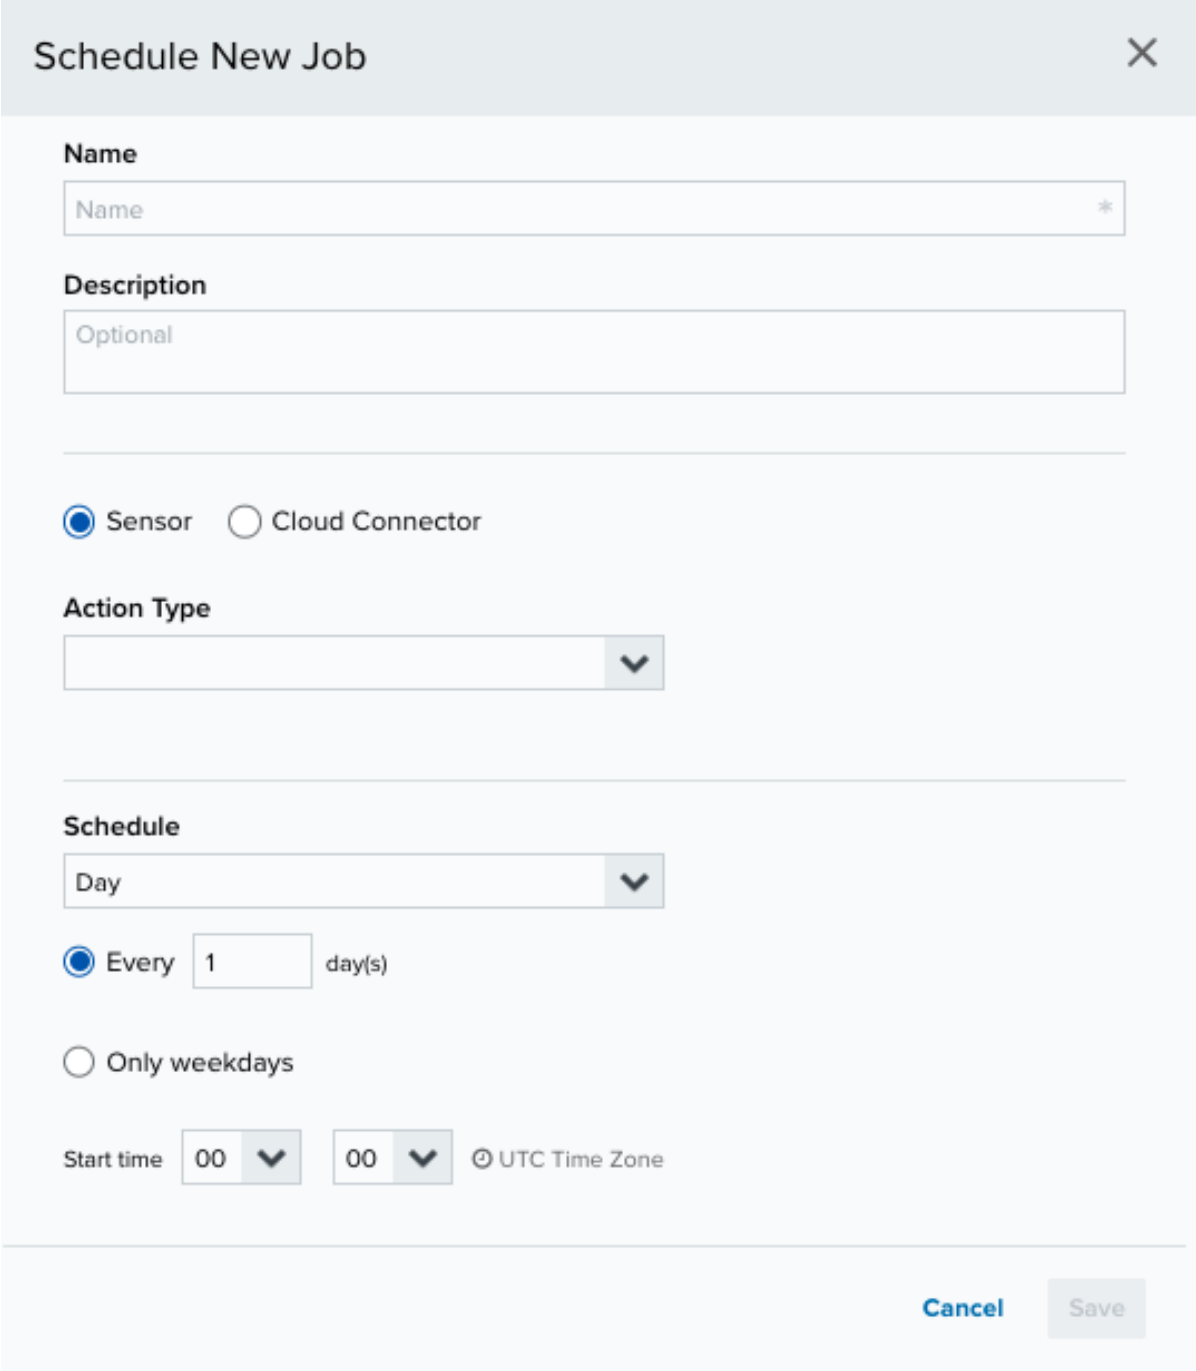

4. Enter the name and description for the job.

   The description is optional, but it is a best practice to provide this information so that others can easily understand what it does.

5. Select **Sensor** as the source for your new job.

6. In the Action Type option, select **Amazon Web Services**.
7. Select a sensor if you have more than one installed in your environment.
8. In the App Action option, select **Monitor S3 Bucket**.

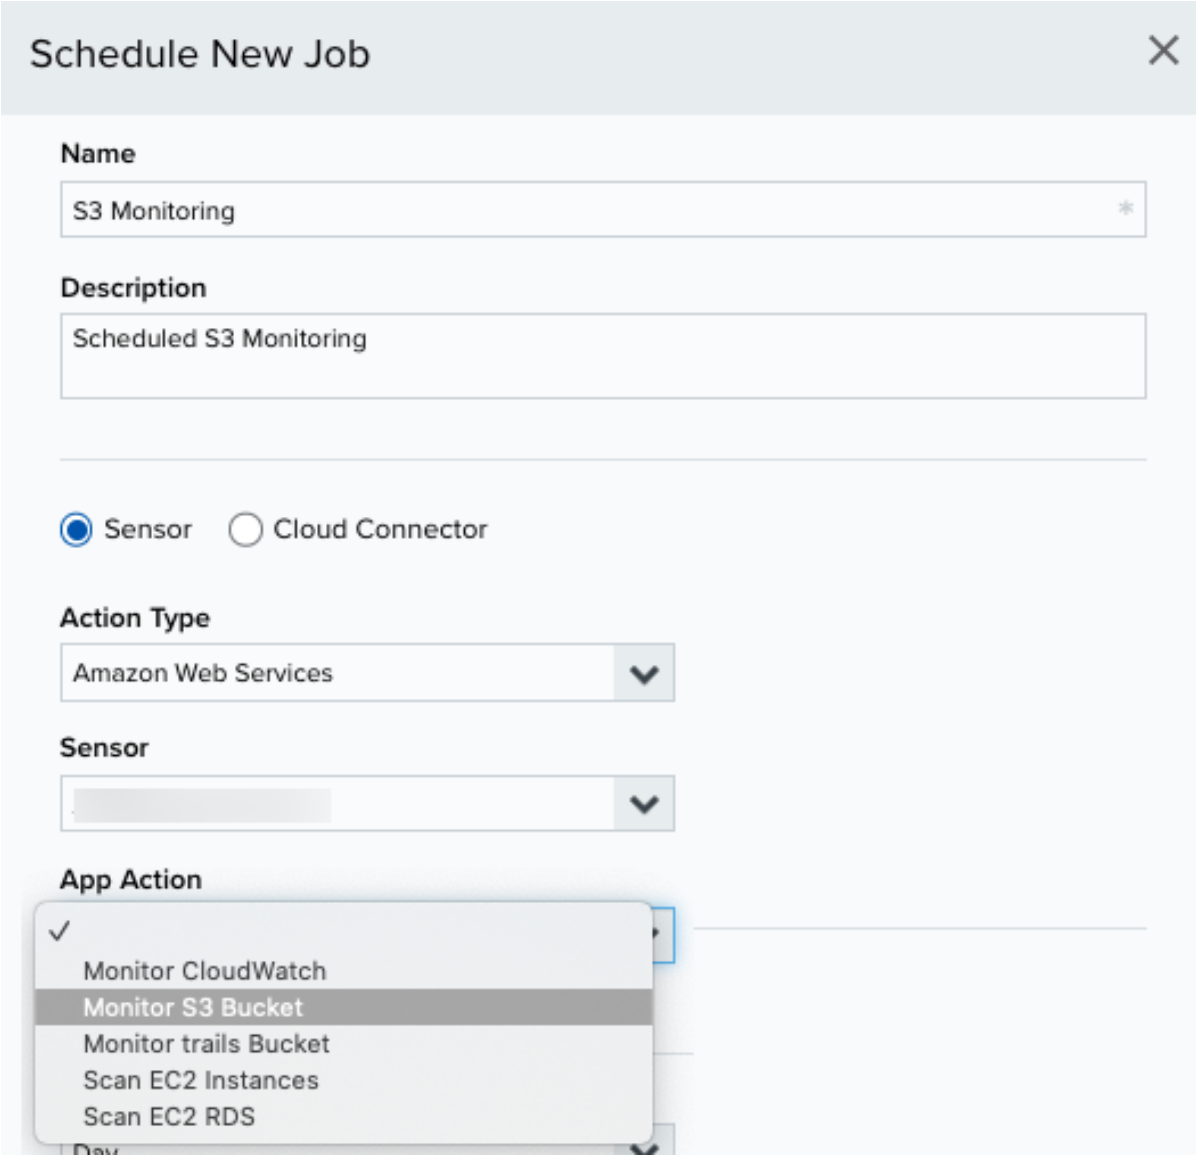

9. Enter the **Bucket Name** and **Path**.

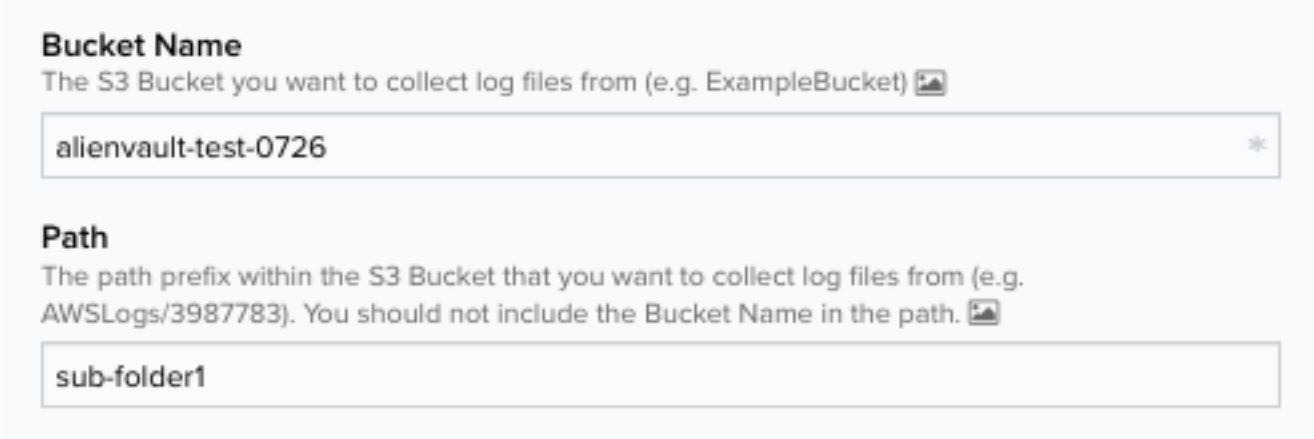

The bucket name is the name of the Amazon S3 bucket as configured in your AWS account, such as `alienvault-test-0726` in the screenshot below.

The path is the path prefix within the Amazon S3 bucket, such as `sub-folder1` in the screenshot below. This does not include the bucket name.

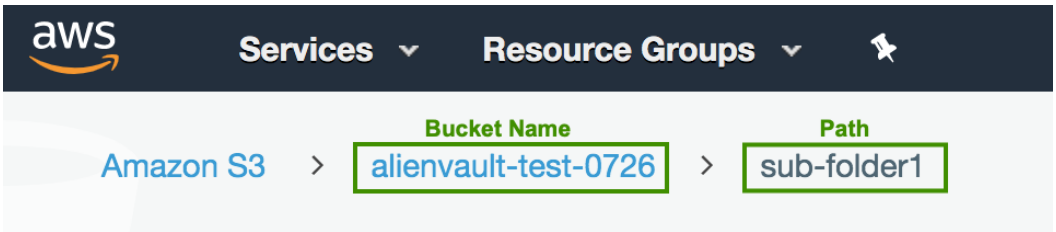

Note: Logs from the directory and its subdirectories are collected.

Important: If you have selected Elastic Load Balancer (ELB), Application Load Balancer (ALB), or Cloud Trail sources, then you need to use, inside the path field, the same prefix you have introduced in your AWS configuration. If the prefix field is empty in your AWS configuration, then you must leave the path field inside USM Anywhere empty.

10. In Source Format, select either of the following log formats:

    - **syslog**: Standard format for transmitting log data to USM Anywhere.
    - **raw**: Use for non-syslog formatted data.

11. In the Schedule section, specify when USM Anywhere runs the job:

    a. Select the increment as **Minute**, **Hour**, **Day**, **Week**, **Month**, or **Year**.

    > **Warning:** After a frequency change, monitor the system to check its performance. For example, you can check the system load and CPU. See USM Anywhere System Monitor for more information.

    b. Set the interval options for the increment.

    The selected increment determines the available options. For example, on a weekly increment, you can select the days of the week to run the job.

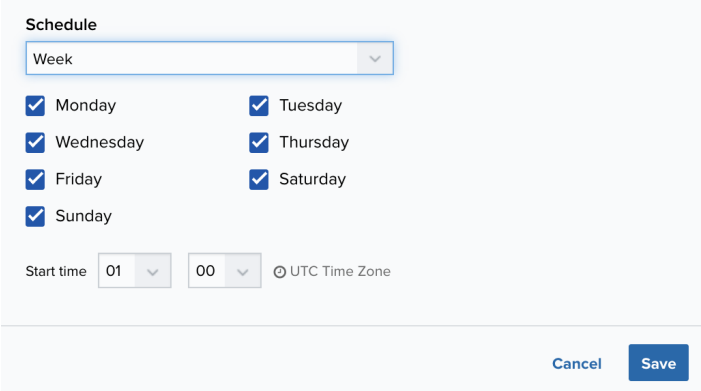

Or on a monthly increment, you can specify a date or a day of the week that occurs within the month.

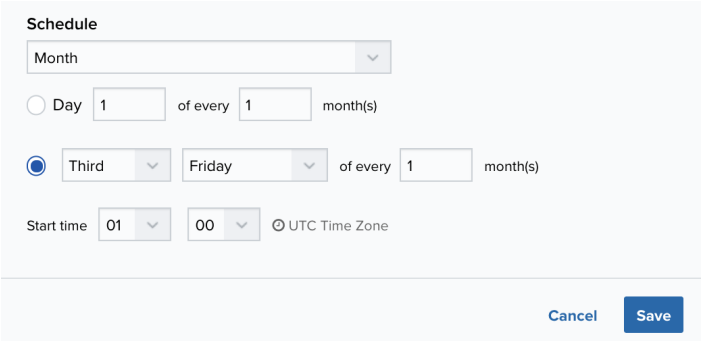

Important: USM Anywhere restarts the schedule on the first day of the month if the option "Every x days" is selected.

c. Set the start time.

   This is the time that the job starts at the specified interval. It uses the time zone configured for your USM Anywhere instance (the default is Coordinated Universal Time [UTC]).

12. Click **Save**.

    USM Anywhere detects any enabled jobs with the same configuration and asks you to confirm before continuing. This is because having two jobs with the same configuration generates duplicate events and alarms.

13. In the AWS console, restart the AWS Sensor instance so that it detects the new configuration.
    You can confirm that the scheduled job is collecting logs by going back to **Settings > Scheduler > Log Collection** and expanding the job you've created. Each log collection event will be listed under Schedule History.

Moving Logs from an Amazon EC2 Instance to an Amazon S3 Bucket

In Amazon Elastic Compute Cloud (EC2), it can be difficult to create direct network connections between isolated parts of your environment. Amazon S3 provides a convenient way to move application logs from an Amazon EC2 instance to an Amazon S3 bucket. Amazon S3 buckets are used to store objects that consist of data and metadata that describes the data. You then configure the AWS Sensor to retrieve and process the log files.

You'll want to synchronize logs from your instance with an Amazon S3 bucket. There are multiple ways to do this. The easiest method is to use the AWS Command Line Interface (CLI) as documented by Amazon. You then create a script similar to the following example and configure it to run periodically as a cron job.

```
aws s3 sync "<path_to_log>" "S3://<bucket_name>/<storage_path>/"
```

Collect Logs from Amazon S3 Buckets with KMS Encryption

Role Availability	✗ Read-Only	✗ Investigator	✗ Analyst	✓ Manager

If you are using a key management service (KMS) key to encrypt the Amazon S3 buckets where your logs are stored, you need to perform the following steps to enable your USM Anywhere Sensor to decrypt those buckets.

> **Note:** To do this, you first need to know the bucket that is encrypted, the KMS key used for the encryption, and the Identity and Access Management (IAM) role created for your sensor.

To enable your sensor to decrypt KMS-encrypted buckets

1. Log in to the AWS Management Console and navigate to the Key Management Service (KMS) page.
2. Open the Customer Managed Keys page and locate the KMS key you are using.
3. Scroll down to the *Key Users* section.
4. Click **Add**.

5. Use the list or the search bar to select the IAM role created for your sensor.

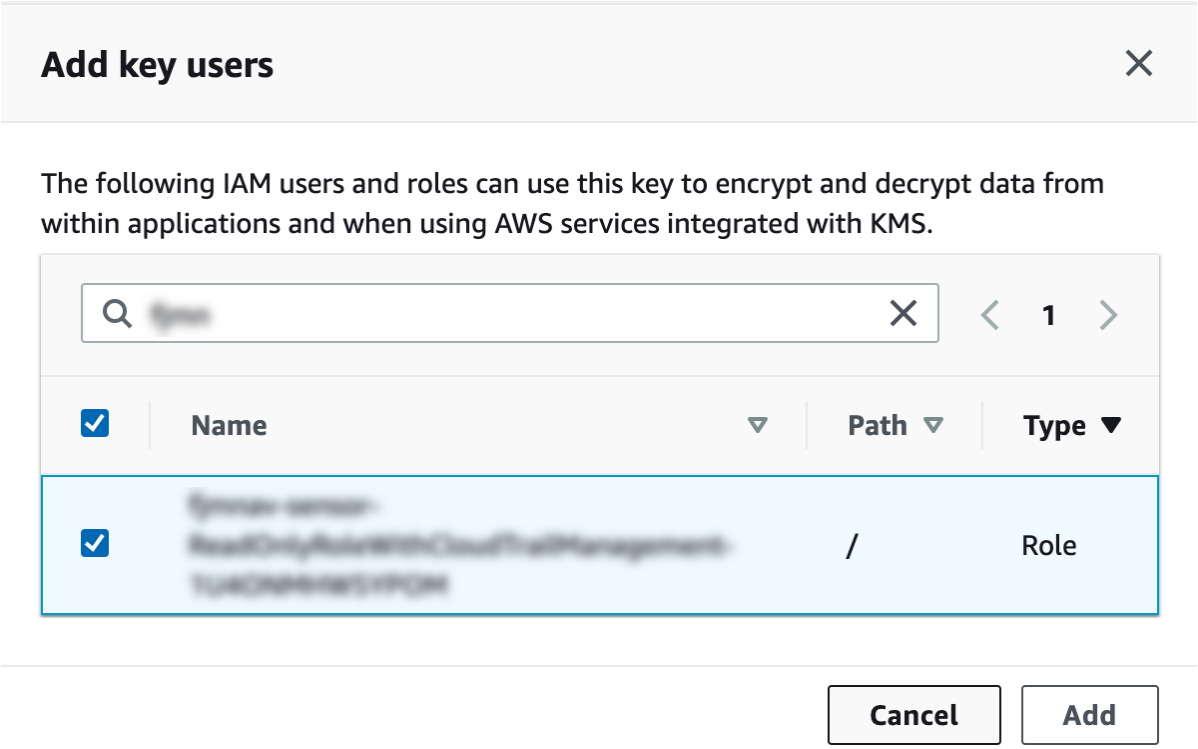

6. Click **Add**.

USM Anywhere Sensor Deployment on Microsoft Azure

The USM Anywhere Sensor provides operational visibility into the security of your Microsoft Azure environment. Based on the collected log information, USM Anywhere analyzes the data generated by your Azure environment and provides real-time alerting to identify malicious activity. The Azure Sensor is deployed into your environment to provide ultimate control over the installation and the data contained within it, and also to avoid any external access to the environment.

This section discusses the following topics:

About Azure Sensor Deployment 183

Requirements for Azure Sensor Deployment 184

Deploy the USM Anywhere Sensor from the Azure Marketplace 194

Create an Application and Obtain Azure Credentials 197

Connect the Azure Sensor to USM Anywhere 203

Complete the Azure Sensor Setup 207

Azure Log Discovery and Collection in USM Anywhere 225

About Azure Sensor Deployment

All USM Anywhere Sensors allow for authenticated scans of assets by leveraging stored credentials that you define in USM Anywhere. This enables USM Anywhere to detect potential vulnerabilities, installed software packages, and running processes and services.

In addition to these standard sensor functions, the Azure Sensor also provides capabilities that leverage the Azure environment:

- Automatic discovery of virtual machines (VMs) running in your Microsoft Azure environment
- Optional monitoring of Azure logs
- Integration with Collect Logs from Azure Event Hubs

Log Collection and Scans

USM Anywhere automatically discovers your use of the following logs without requiring enablement on the Azure subscription side, as long as the Azure resource subscription has contributor-level permissions:

- Azure Representational State Transfer (REST) Monitor (formerly Azure Insight) logs
- Azure security alerts
- Azure web apps logs
- Azure SQL Server logs

> **Note:** The Azure SQL Server job is deprecated. Use the Event Hub Integration to collect Azure SQL Server logs. See Collect Logs from Azure Event Hubs for more information.

- Azure Internet Information Services (IIS) logs
- Azure Windows logs
- Asset scans on your VMs to inventory installed software packages, running processes, and services

Log Analysis

USM Anywhere analyzes these logs in these stages:

Stage 1: Collects logs from systems and software running in your environment

Stage 2: Configures log line processing and generates events

- Includes IP addresses and timestamps culled from extracted log-line data
- Adds other data to the event, such as security context and environmental information

Stage 3: Analyzes events and stores them

Deployment Overview

AT&T Cybersecurity distributes the Azure Sensor through the Azure Marketplace as a D2 Standard or DS2 Premium VM template.

> **Note:** If your organization uses multiple subnets to allow communication between headquarters and remote offices, AT&T Cybersecurity recommends that you deploy a sensor to each. Alternatively, you can deploy a USM Anywhere Sensor in a single virtual network. When you deploy a sensor to a single virtual network in your Azure subscription, you'll see Azure logs for the entire subscription.

The deployment process for an initial USM Anywhere Sensor in your Azure environment consists of these primary tasks:

1. Review requirements for an Azure Sensor deployment
2. Deploy the USM Anywhere Sensor within your Azure environment
3. Register the deployed sensor with your sensor authentication code to provision the USM Anywhere instance and connect the deployed sensor
4. (Optional.) Manually create a new application and credentials in the Azure console
5. Complete your Azure Sensor configuration, including initial asset discovery

Requirements for Azure Sensor Deployment

To ensure that you can successfully deploy USM Anywhere in your Microsoft Azure subscription and monitor all of your Azure resources, make sure you have the following available in your Azure environment:

- An Azure account with privileges in the resource group or subscriptions that you want to install the USM Anywhere Sensor.

  > **Note:** You can deploy a single USM Anywhere Sensor to monitor all of your Azure resource groups. To do this, you must assign the application you create to the entire subscription.

- Administrative access to Active Directory (AD) within Azure.

This AD access enables you to create an application required to install resource groups or a subscription for monitoring.

- A virtual network inside the resource group.
- A subnet inside the virtual network.
- A storage account.

Important: USM Anywhere does not support Azure Classic accounts.

Important: Because the needs of a sensor differ based on the varying demands of different deployment environments and the complexity of events being processed, the number of events per second (EPS) a sensor can process varies.

Depending on your environment, you may need to deploy additional sensors to ensure that all events are processed.

Warning: Be sure not to install any application outside of those already provided within your image where you are deploying your Azure Sensor.

You may want to check your system for automatically installed applications, such as OMIAgent, which must be uninstalled. Left uninstalled, such applications may make your environment or your sensor unstable.

Sensor Ports and Connectivity

Note: To launch the USM Anywhere Sensor web UI during the initial setup, you need to allow inbound traffic to the sensor IP address through TCP port 80. You can remove access to this port after the sensor successfully connects to USM Anywhere. You do not need to allow inbound traffic to this port from the Internet.

The following tables list the inbound and outbound ports.

Sensor Ports and Connectivity (Outbound Ports)

Type	Ports	Endpoints	Purpose
TCP	443	update.alienvault.cloud	Communication with AT&T Cybersecurity for initial setup and future updates of the sensor.
TCP	443	reputation.alienvault.com	Ongoing communication with AT&T Alien Labs™ Open Threat Exchange® (OTX™).
TCP	443	otx.alienvault.com	Ongoing communication with OTX to retrieve vulnerability scores. Connecting to otx.alienvault.com is not required but highly recommended. OTX uses the AWS CloudFront services. Refer to the AWS IP address ranges page when you deploy a new sensor. This page contains the current IP address ranges for the service and instructions on how to filter the addresses.
TCP	443	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
SSL	443	storage-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send and retrieve backups.
SSL	443	metrics-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send metrics and messages.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
SSL/TCP	443	api-parameters-<REGION>-prod.alienvault.cloud[4] api-message-proxy-<REGION>-prod.alienvault.cloud api.message-proxy.<REGION>.prod.alienvault.cloud	Ongoing communication with USM Anywhere. It is only necessary to allowlist the address that corresponds to the region where your USM Anywhere instance is hosted.
SSL/TCP	7100	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
UDP	53	DNS Servers (Google Default)	Ongoing communication with USM Anywhere.
UDP	123	0.ubuntu.pool.ntp.org 1.ubuntu.pool.ntp.org 2.ubuntu.pool.ntp.org 3.ubuntu.pool.ntp.org	Sync with NTP services in the Azure Cloud.
TCP	22 and 443	prod-usm-saas-tractorbeam.alienvault.cloud prod-gov-usm-saas-tractorbeam.gov.alienvault.us (for AT&T TDR for Gov)	SSH communications with the USM Anywhere remote support server. See Troubleshooting and Remote Sensor Support for more information about remote technical support through the USM Anywhere Sensor console.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	<event-hub-namespace>.servicebus.windows.net	(Optional.) Connect to Microsoft Azure Event Hubs for log collection. Replace <event-hub-namespace> with the name of your Event Hubs namespace. If your environment includes additional services such as AMQP or Kafka, you may need to make additional ports available. See Microsoft's Troubleshooting Guide for detailed information about these potential additional port requirements.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	geoip-ap-northeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-south-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-2-prod.alienvault.cloud/geo-ip/sensor geoip-ca-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-me-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-sa-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-us-gov-west-1-prod-gov.gov.alienvault.us/geo-ip/sensor (for AT&T TDR for Gov)	Allows resolution of IP addresses for geolocation services. It is only necessary to allowlist the GeoIP address that corresponds to the region where your USMA instance is hosted.

1

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

2

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

3

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

4

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

5

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

Sensor Ports and Connectivity (Inbound Ports)

Type	Ports	Purpose
SSH	22	Inbound method for secure remote login from a computer to USM Anywhere.
HTTP	80	Inbound communication for HTTP traffic.
UDP (RFC 3164)	514	USM Anywhere collects data through syslog over UDP on port 514 by default.
TCP (RFC 3164)	601	Inbound communication for reliable syslog service. USM Anywhere collects data through syslog over TCP on port 601 by default.
TCP (RFC 5424)	602	USM Anywhere collects data through syslog over TCP on port 602 by default.
Traffic Mirroring	4789	Inbound communication for virtual extensible local area network (VXLAN).

Sensor Ports and Connectivity (Inbound Ports) (Continued)

Type	Ports	Purpose
WSMANS	5987	Inbound WBEM WS-Management HTTP over Secure Sockets Layer/Transport Layer Security (SSL/TLS) (NXLog).
TLS/TCP (RFC 3164)	6514	USM Anywhere collects TLS-encrypted data through syslog over TCP on port 6514 by default.
TLS (RFC 5424)	6515	USM Anywhere collects data through syslog over TLS on port 6515 by default.
TCP	9000	Inbound communication used internally for HTTP sensor traffic.
Graylog	12201	Inbound communication for Graylog Extended Log Format (GELF).

USM Anywhere IP Addresses for Allowlisting

Your sensor is connected to a USM Anywhere instance deployed in one of the Amazon Web Services (AWS) endpoint regions based on your location. If you need to configure your firewall to allow communication between the sensor and the USM Anywhere instance, refer to the following table with the reserved IP address ranges for each region.

Important: The Update Server and the AlienVault Agent always use the 3.235.189.112/28 range no matter where your USM Anywhere is deployed. The AT&T TDR for Gov Update Server uses the 3.32.190.224/28 range.

Note: The regional IP ranges listed in this table are limited to the control nodes (subdomain). You must also meet all requirements provided in the Sensor Ports and Connectivity (Outbound Ports) table.

AWS Regions Where USM Anywhere Instance Is Available

Code	Name	Reserved Static IP Address Ranges
ap-northeast-1	Asia Pacific (Tokyo)	18.177.156.144/28 3.235.189.112/28 44.210.246.48/28
ap-south-1	Asia Pacific (Mumbai)	3.7.161.32/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-1	Asia Pacific (Singapore)	18.143.203.80/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-2	Asia Pacific (Sydney)	3.25.47.48/28 3.235.189.112/28 44.210.246.48/28
ca-central-1	Canada (Central)	3.96.2.80/28 3.235.189.112/28 44.210.246.48/28
eu-central-1	Europe (Frankfurt)	18.156.18.32/28 3.235.189.112/28 44.210.246.48/28
eu-west-1	Europe (Ireland)	3.250.207.0/28 3.235.189.112/28 44.210.246.48/28

AWS Regions Where USM Anywhere Instance Is Available (Continued)

Code	Name	Reserved Static IP Address Ranges
eu-west-2	Europe (London)	18.130.91.160/28 3.235.189.112/28 44.210.246.48/28
me-central-1	Middle East (UAE)	3.29.147.0/28 3.235.189.112/28 44.210.246.48/28
sa-east-1	South America (São Paulo)	18.230.160.128/28 3.235.189.112/28 44.210.246.48/28
us-east-1	US East (N. Virginia)	3.235.189.112/28 44.210.246.48/28
us-west-2	US West (Oregon)	44.234.73.192/28 3.235.189.112/28 44.210.246.48/28
us-gov-west-1	AWS GovCloud (US-West)	3.32.190.224/28

Azure Portal URLs for Proxy Bypass

The URL endpoints to allowlist on your Azure portal are specific to the Azure cloud where your environment is deployed. To allow network traffic to reach these endpoints, select your cloud environment, and then add the following list of URLs to your proxy server or firewall.

```
*.aadcdn.microsoftonline-p.com
*.aka.ms
*.applicationinsights.io
*.azure.com
*.azure.net
*.azureafd.net
```

```
*.azure-api.net
*.azuredatalakestore.net
*.azureedge.net
*.loganalytics.io
*.microsoft.com
*.microsoftonline.com
*.microsoftonline-p.com
*.msauth.net
*.msftauth.net
*.trafficmanager.net
*.usgovcloudapi.net (AT&T TDR for Gov only)
*.visualstudio.com
*.windows.net
*.windows-int.net
```

Deploy the USM Anywhere Sensor from the Azure Marketplace

After you review the requirements and make sure that your Microsoft Azure environment is configured as needed, you can deploy the USM Anywhere Sensor for Azure. AT&T Cybersecurity provides the virtual machine (VM) template for the sensor and makes it available through the Microsoft Azure Marketplace for easy deployment.

Note: Azure limits the availability of the Azure Marketplace to customers according to country. On the Marketplace FAQs page, the "Azure Marketplace for Customers" section provides a current list of supported countries.

To deploy a USM Anywhere Sensor from the Azure Marketplace

1. Go to the USM Anywhere Sensor Downloads page and click the icon of your specific sensor.

   If you are not already logged in to the Azure console, this link launches the Microsoft Azure Login page. Provide your Azure account credentials (username and password) and click **Sign in**.

2. On the page, review the details of the license and click **Create**.

   This takes you to the *Create a virtual machine* page, which guides you through the steps for deploying the USM Anywhere Sensor VM.

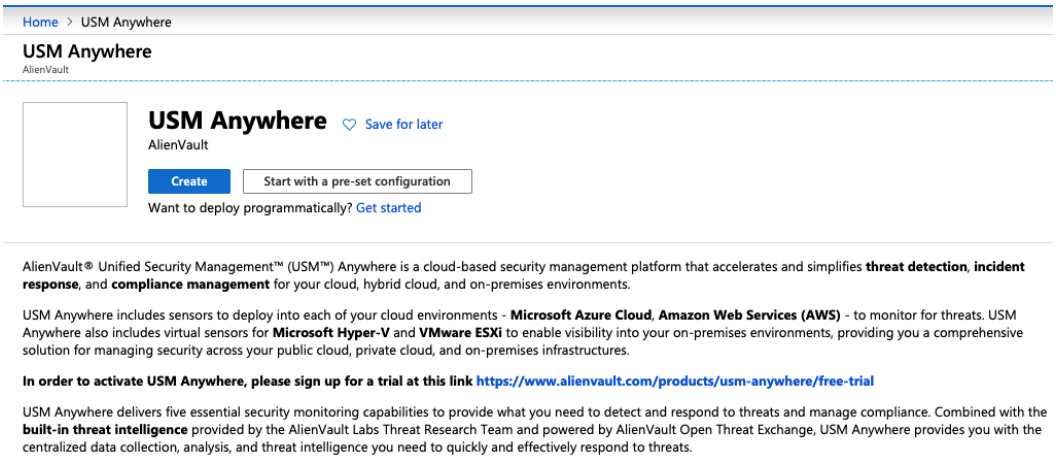

3. On the Basics tab, specify the required fields for the VM:

   - **Subscription**: Select the subscription into which the USM Anywhere Sensor should be installed.
   - **Resource Group**: Indicate whether you want to install the USM Anywhere Sensor into an existing resource group or into a new resource group. If new, enter a unique name.
   - **Virtual machine name**: Enter the name you want to use for the USM Anywhere Sensor VM.
   - **Region**: Select the region you want to deploy the USM Anywhere Sensor VM.
   - **Image**: This field is set to Unified Security Management (USM) Anywhere.
   - **Size**: This field is set to Azure Standard D2 v2.
   - **Authentication type**: Set this option to specify an SSH public key or a password for SSH access.
   - **Username**: Enter a username.

     Important: AT&T Cybersecurity recommends using *sysadmin* as the username. If you use a different name, you will need to "sudo up" to access the sensor console. See Checking Connectivity to the Remote Server for more information.

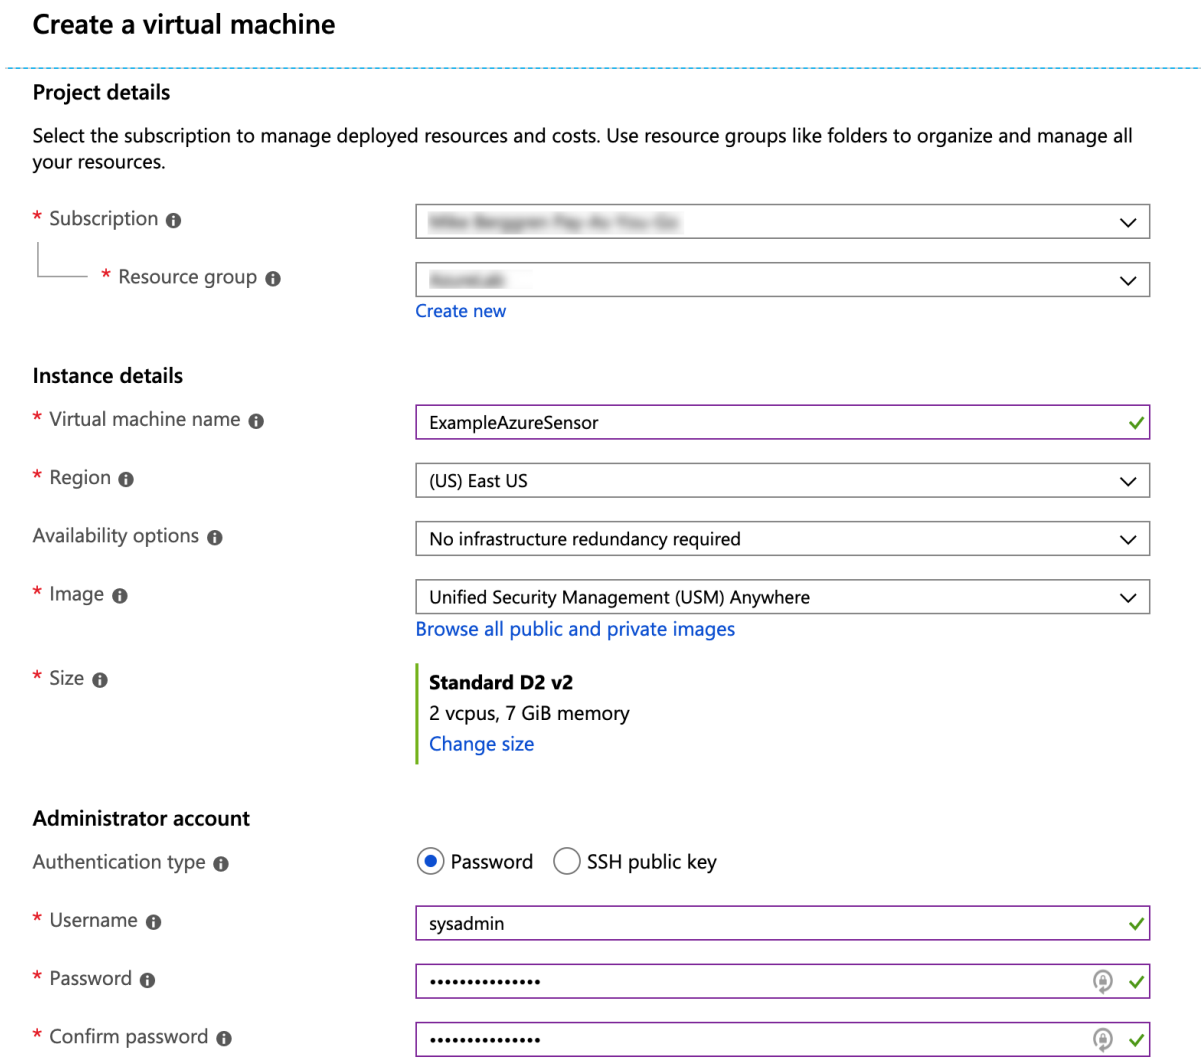

4. Click **Next : Disks**.
5. On the Disks tab, select **Standard SSD** as the disk type.
6. Click **Next : Networking**.
7. On the Networking tab, select the virtual network or subnet upon which the USM Anywhere Sensor VM should be installed. Keep the other defaults.

> **Important:** Make sure you install the USM Anywhere Sensor in the network that has sufficient connectivity to the assets that you want to monitor.

8. Click **Review + create** to keep the defaults on the remaining tabs.
9. On the Review + create tab, review your specifications and the cost summary.

10. Click **Create**.

    This starts the deployment of the USM Anywhere Sensor, which can take up to six minutes.

11. After deployment finishes, click **Go to resource** or go to the overview page of the VM and locate its public IP address.

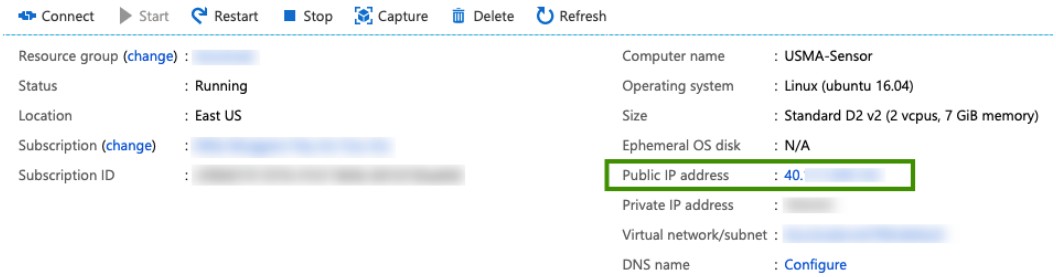

12. Paste the IP address into your browser to launch the USM Anywhere Sensor Setup page.

> **Important:** This link requires that inbound port 80 is open for the sensor VM, which is not a default network setting on Azure. See Sensor Ports and Connectivity for more information.

Create an Application and Obtain Azure Credentials

Role Availability ✗ Read-Only ✗ Investigator ✗ Analyst ✓ **Manager**

To enable USM Anywhere to monitor your Microsoft Azure subscription, you must create an application that grants permission to USM Anywhere to fetch data using the Azure software development kit (SDK) and Azure Representational State Transfer (REST) API. USM Anywhere requires the following credentials:

Required Azure Credentials

Azure Credential	USM Anywhere Field Name
azure_tenant_id	Azure Tenant ID
azure_subscription_id	Azure Subscription ID
azure_application_id	Azure Application ID
azure_application_key	Azure Application Key

The following instructions focus on the requirements for USM Anywhere. See Microsoft documentation for detailed steps and descriptions to register an application using the Azure portal, including a video demonstration.

Important: You must have global administrator privileges to create an application and obtain credentials.

Obtain the Azure Subscription ID

The subscription identifier (ID) is required when you complete the Azure Credentials step of the sensor setup in USM Anywhere.

To get the Azure subscription ID

1. Log in to the Azure portal (https://portal.azure.com).
2. From the Azure Dashboard, select your subscription.
3. From the Subscription page, copy your subscription ID and save it somewhere that you can access later.

Create the Application in Azure

To allow USM Anywhere to access Azure resources, you must first set up an Azure Active Directory (AD) application and complete the Azure standard procedure for adding a new application registration. Then you can create a client secret for Azure AD.

To create the application in Azure

1. Log in to the Azure portal (https://portal.azure.com).
2. Go to **Azure Active Directory > App registrations > New registration**.
3. Enter a name for the application.
4. In Supported account types, select **Accounts in any organizational directory (Any Azure AD directory - Multitenant)**.
5. Click **Register**.
6. After the application is created, you can locate the application(client) ID, directory (tenant) ID, and object ID needed to complete the Azure Credentials step of the sensor setup in USM Anywhere.

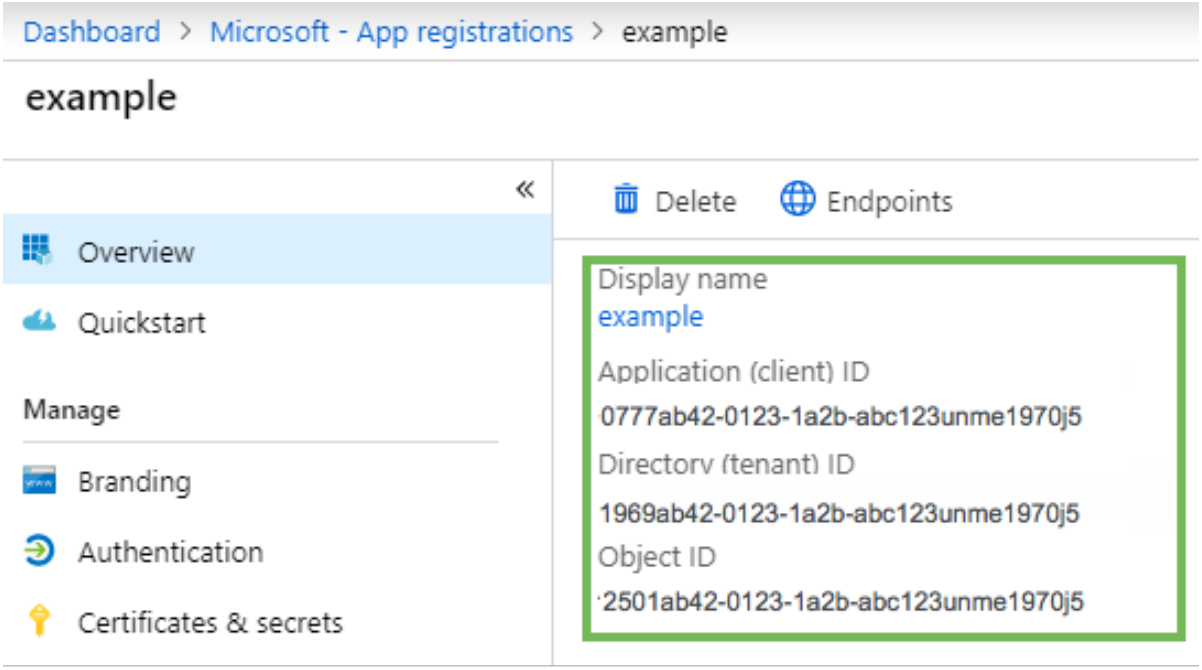

7. Go to **Certificates & secrets** and click **New client secret**.
8. Enter a description for the secret and select a duration.
9. Click **Add**.

   The value displayed in the Azure portal is the *Azure Application Key* used by USM Anywhere.

   Important: Copy this value and save it because you won't be able to copy the key later.

Grant API Permissions

To let your application collect user information in your Azure environment, you need to grant Microsoft Graph API permissions.

To grant API permissions

1. Log in to the Azure portal (https://portal.azure.com) and select your application.
2. Go to **API Permissions** and click **Add a permission**.

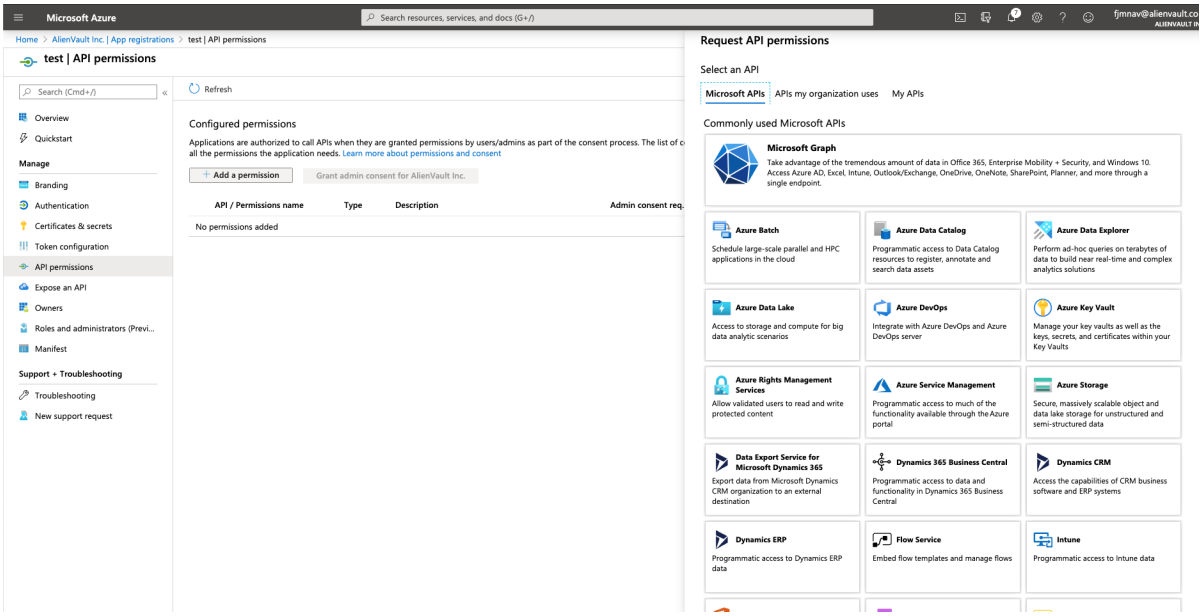

3. Select **Microsoft Graph**.
4. Select **Application permissions** and then **User.Read.All**. Use the search function to help locate the permissions.

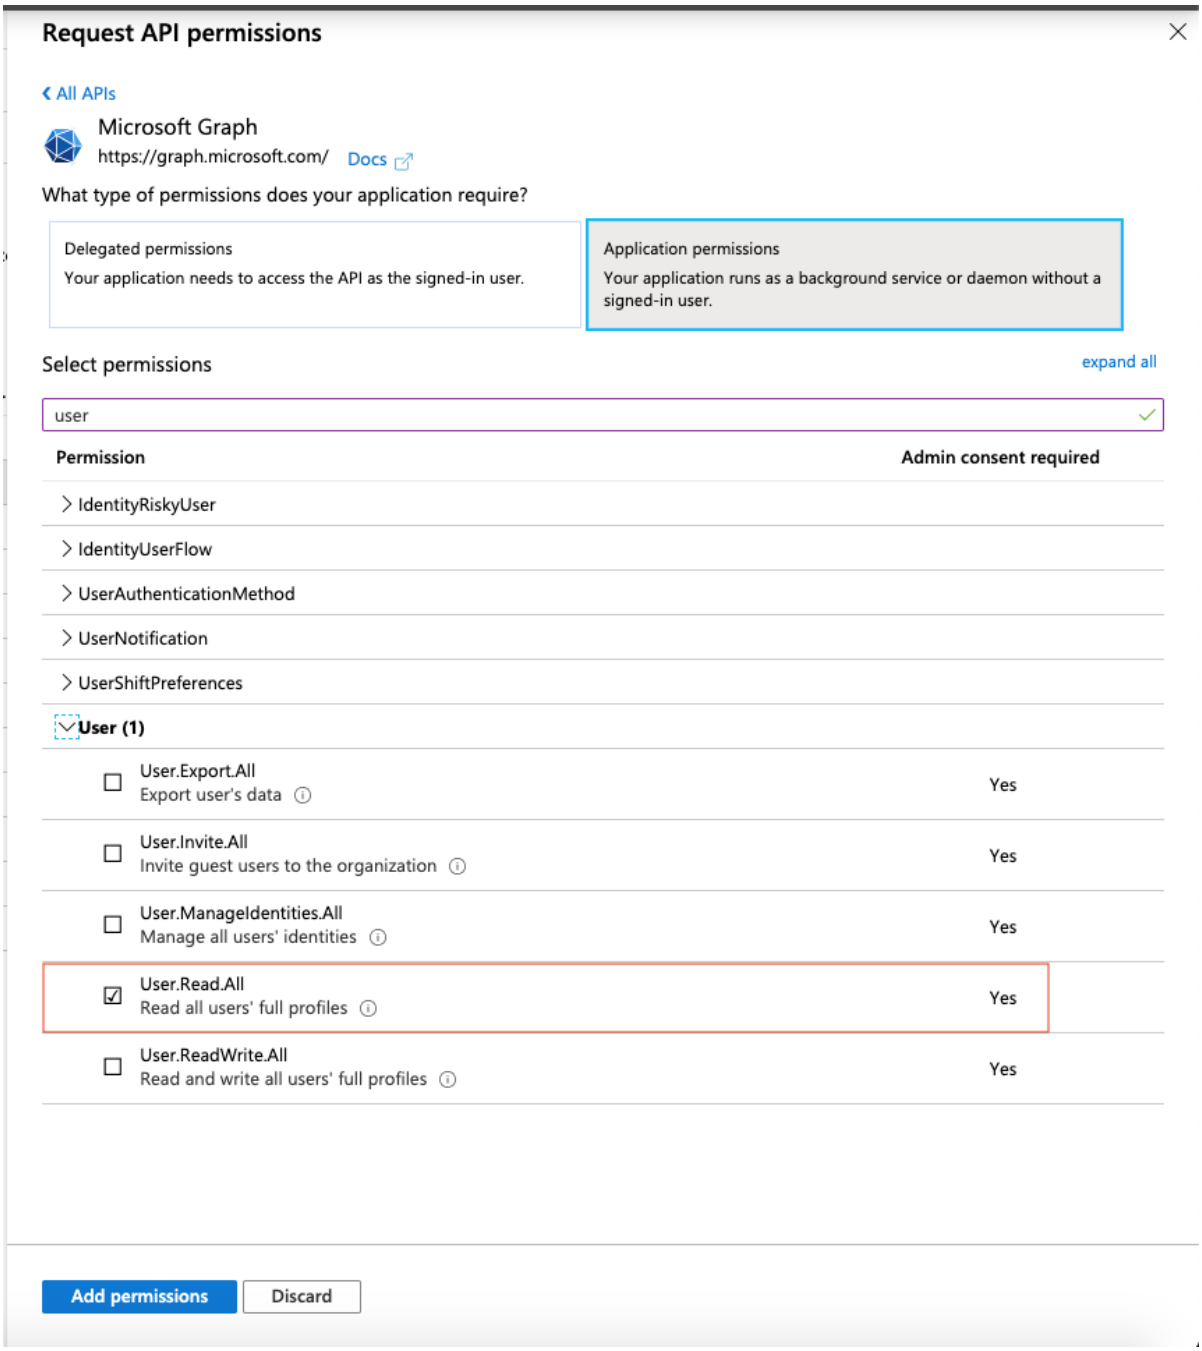

5. Click **Add Permissions**.
6. These permissions require admin approval, so make sure to click **Grant admin consent for**.

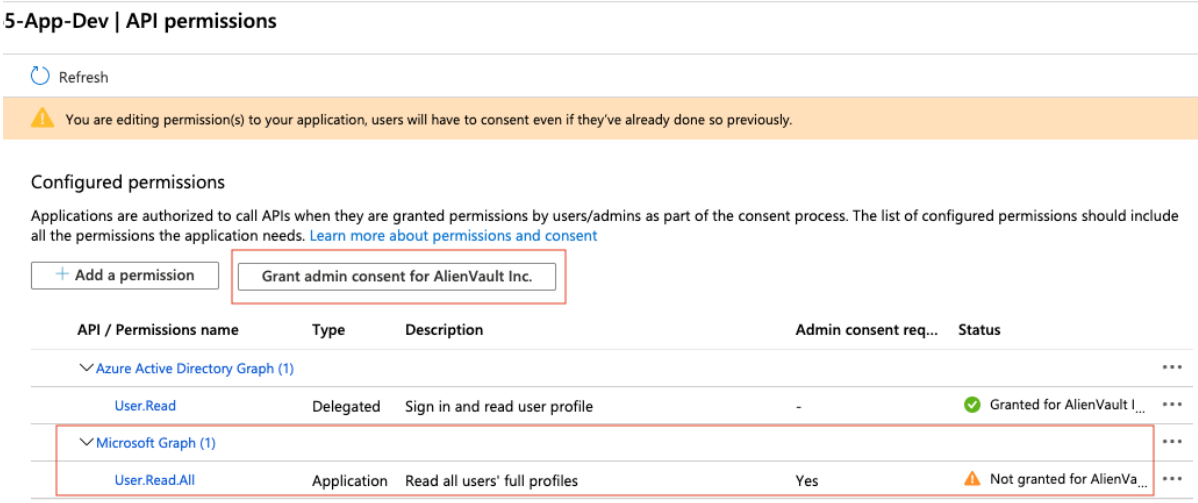

Associate the Application with the Entire Subscription

If you want to use USM Anywhere to monitor all of your Azure resources, you should associate it with your Azure subscription as a whole.

To associate the application with the entire subscription

1. Log in to the Azure portal (https://portal.azure.com).
2. Go to **More Services > Subscriptions**, locate the subscription, and select it.
3. Select **Access control (IAM)** in the navigation list.

   This displays the roles and permissions for the subscription.

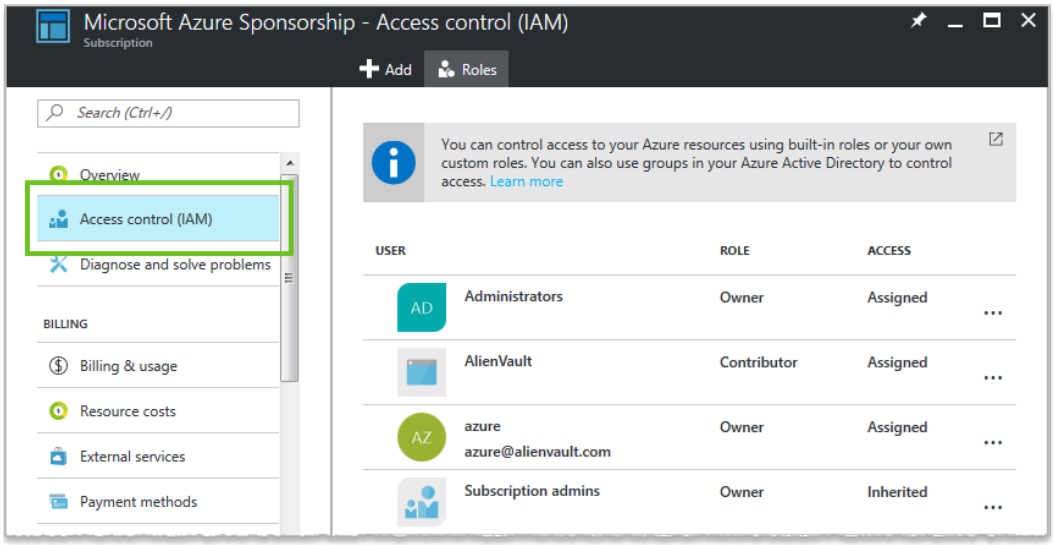

4. At the top of the page, click **Add**.
5. Select the **Reader** role (recommended).

   This role allows assigned users to fetch new Azure logs.

   > **Warning:** You must select the **Contributor** role if you want to collect Microsoft Internet Information Services (IIS), Azure SQL Server, or Windows logs.

6. Select the application you created previously to assign the role to the subscription.
7. Click **Save** and **OK**.

Connect the Azure Sensor to USM Anywhere

After deploying the Microsoft Azure Sensor, you must connect it to USM Anywhere through registration.

Obtain the Authentication Code

You must enter an authentication code when registering the USM Anywhere Sensor. How to obtain the authentication code depends on your USM Anywhere instance and whether this is the first sensor you're deploying.

Instructions for USM Anywhere customers:

If this is your first USM Anywhere Sensor, you must register the sensor using the initial authentication code (starts with a "C") received from AT&T Cybersecurity. With this code, the registration process provisions a new USM Anywhere instance and defines its attributes, such as how many sensors to allow for connection, how much storage to provide, and what email address to use for the initial user account. After registration, you will gain access to the sensor through the USM Anywhere web user interface (UI), where you can complete the sensor setup.

If you are deploying additional sensors, you must generate the authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Instructions for AT&T TDR for Gov customers:

AT&T Cybersecurity has already provisioned the AT&T Threat Detection and Response for Government (AT&T TDR for Gov) instance for you, therefore you won't receive an authentication code for your sensor. This is true regardless if it's the first sensor or additional

sensors you're deploying. However, for the first sensor, you'll receive a link to access your instance.

For every sensor you deploy, you must generate an authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Register Your Sensor

You perform this procedure after deploying the USM Anywhere Sensor within your Azure subscription. The IP address link is displayed after you create the virtual machine (VM) and the instance is running in your Azure environment.

To register your sensor

1. Click the public IP address displayed for the running sensor VM in the Azure console.

   Important: This link requires that inbound port 80 is open for the sensor VM, which is not a default network setting on Azure. See Sensor Ports and Connectivity for more information.

   This opens the *Welcome to USM Anywhere Sensor Setup* page, which prompts you to provide the information for registering the sensor with your new USM Anywhere instance.

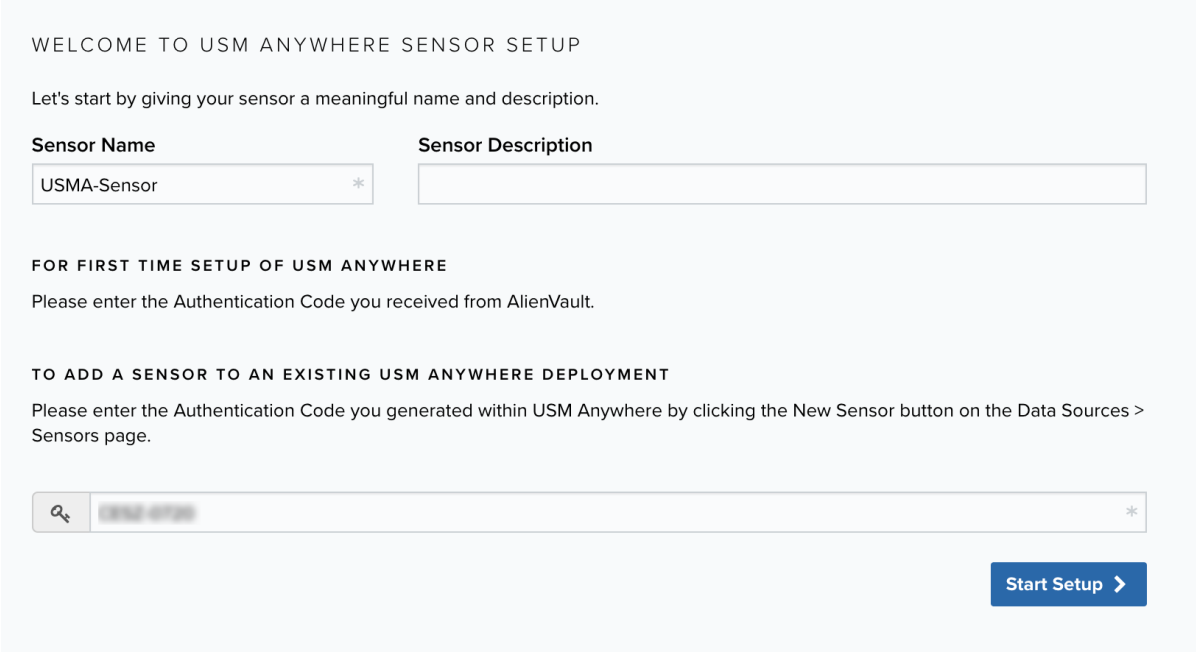

2. Enter a sensor name and sensor description.

3. Paste the authentication code into the field with the key icon ( ).

4. Click **Start Setup** to start the process of connecting the USM Anywhere Sensor.

   It takes about 20 minutes to provision your USM Anywhere instance upon registration of your initial sensor. When this instance is provisioned and running, you'll see a welcome message that provides an access link.

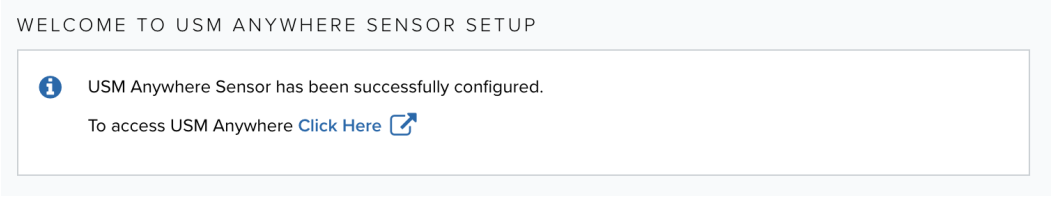

   Use this link to open the secured web console for your USM Anywhere instance. You and the other USM Anywhere users in your organization can access this console from a web browser on any system with internet connectivity.

   Note: If this is your first deployment, you'll also receive an email from AT&T Cybersecurity that provides the access link to USM Anywhere.

Configure the Initial Login Credentials

When you link to a newly provisioned USM Anywhere instance, you must configure the password for the initial user account. This is the default administrator as defined in your subscription.

To configure login credentials

1. In the welcome message, click the link.

   This displays a prompt to set the password to use for the default administrator of USM Anywhere.

2. Enter the password, and then enter it again to confirm.

   Keep in mind these points when you are logging in:

   - The login credentials that you set will apply to any USM Anywhere™ and USM Central™ you have access to.
   - USM Anywhere requires all passwords to have a minimum length of 8 characters and a maximum length of 128 characters.

- The password must contain numerical digits (0-9).
- The password must contain uppercase letters (A-Z).
- The password must contain lowercase letters (a-z).
- The password must contain special characters, such as hyphen (-) and underscore (_).

> **Note:** USM Anywhere passwords expire after 90 days. When your password expires, USM Anywhere enforces a password change when you next log in. A new password must be different from the previous four passwords.After 45 days of inactivity, your user account will be locked. Manager users can unlock inactive accounts.

3. Click **Save & Continue**.
4. When the login page opens, enter the password you just set and click **Login**.

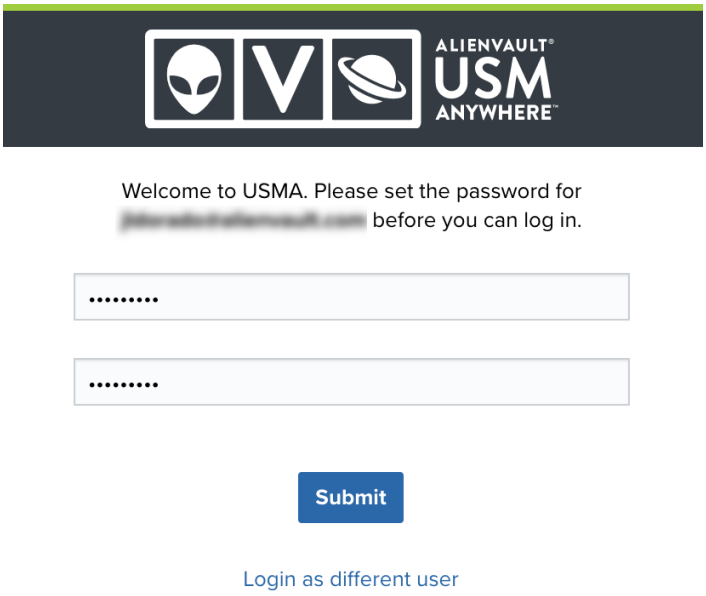

Verify That Your Sensor Is Running

It's a good idea to verify that the USM Anywhere Sensor is running. It also gives you the chance to watch the sensor actively working to find all of your assets and to record events from the start.

> **Note:** Verify that the sensor is running before performing the configuration. You can keep one web browser tab with the Welcome to USM Anywhere page in the background while you perform the verification on a different tab.

To verify that your new sensor is running

1. In USM Anywhere, go to **Data Sources > Sensors**.

   You should now see your sensor in the page. See in the *USM Anywhere User Guide* for more information.

   After a few minutes, USM Anywhere locates your assets and starts generating events.

2. You can review the activity in two locations:

   - From the primary task bar, select **Environment > Assets**.
   - From the primary task bar, select **Activity > Events**.

   Note: It could take up to six minutes before events appear. Make sure to refresh your browser from time to time to display the current data.

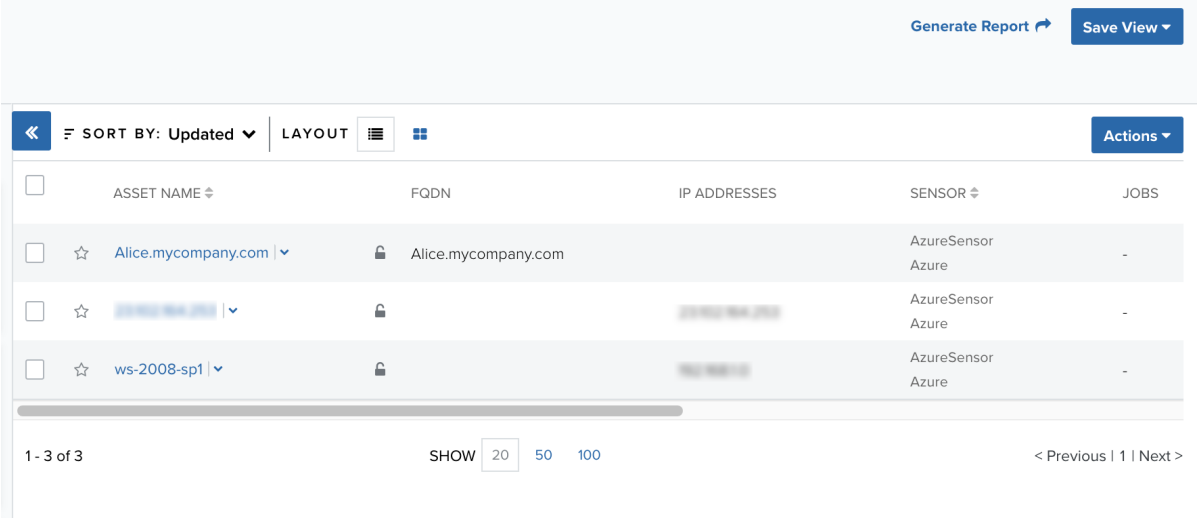

   See the *USM Anywhere User Guide* for more information about using the Assets and Events pages in USM Anywhere.

The Azure SQL Server job is deprecated. Use the Event Hub Integration to collect Azure SQL Server logs. See Collect Logs from Azure Event Hubs for more information.

Complete the Azure Sensor Setup

Role Availability ✗ **Read-Only** ✗ Investigator ✗ Analyst ✓ **Manager**

After you initialize a new USM Anywhere Sensor, you must configure it in the Setup Wizard. As you configure the sensor, you can enable USM Anywhere to perform specific actions through scheduled jobs, such as running an asset discovery scan or collecting security events from a predefined cloud storage location.

Accessing the Setup Wizard

The Setup Wizard is accessible under the following circumstances:

- After you first log in to the USM Anywhere web user interface (UI) and see the Welcome to USM Anywhere page, click **Get Started** to launch the Setup Wizard.
- If you have already registered one USM Anywhere Sensor but did not complete the setup before logging out, the USM Anywhere Sensor Configuration page launches automatically at your next login to remind you to finalize configuration of the sensor. From that page, you click **Configure** to launch the Setup Wizard and complete the sensor configuration.
- If you registered an additional USM Anywhere Sensor, but did not complete the setup, the Sensors page displays an error (⊗) in the Configured column. See in the *USM Anywhere User Guide* for more information.

  Go to **Data Sources > Sensors**, and then click the sensor name to complete the sensor configuration. See in the *USM Anywhere User Guide* for more information.

Configuring the Azure Sensor in the Setup Wizard

The first time you log in from the Welcome to USM Anywhere web page, the Setup Wizard prompts you to complete the configuration of the first deployed sensor. Thereafter, you can use the Sensors page to configure an additional sensor or to change the configuration options for a deployed sensor. See in the *USM Anywhere User Guide* for more information.

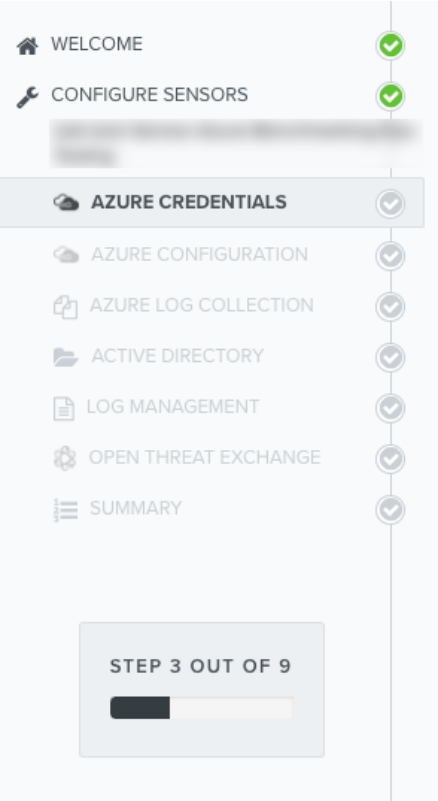

Azure Credentials

To complete the Microsoft Azure Sensor configuration, you must obtain Azure API credentials for the subscription that you want USM Anywhere to monitor. Select the option on the Azure Credentials page that matches your current Azure credential creation status:

- If you already generated your Azure credentials, click **Yes, I have my Azure credentials and am ready to enter them**.
- If you don't yet have your Azure credentials, click **No, I don't have my Azure credentials and need to create them**.
- If you're not sure, click **I am not sure. Show me how to create my Azure credentials**.

If you select **No** or **I am not sure**, the page provides options for two creation methods:

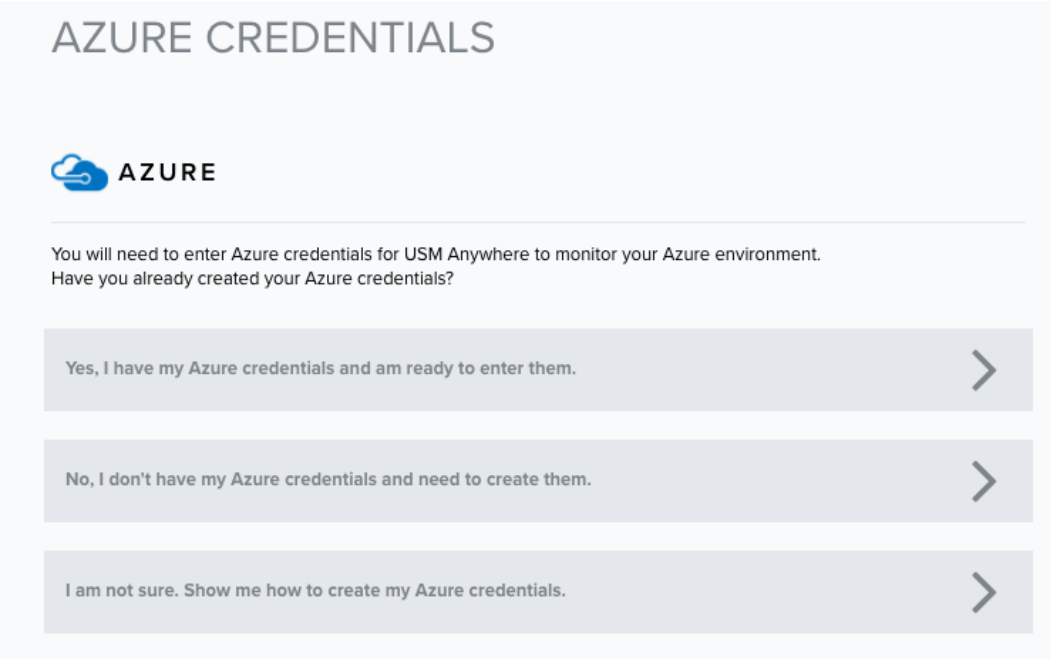

If you select **Yes**, follow the steps in Configuring the Azure Credentials After Manual Credential Generation.

Generating the Azure Credentials for Windows Users

To generate Azure Credentials for Microsoft Windows users

This procedure is for Windows users who want to use the provided Power Shell script to automatically generate their credentials for sensor configuration:

1. Select **Create credentials automatically using a Power Shell script (Recommended)**.

   The page automatically launches a download of the Power Shell script. You can use the browser tools to save the file to the appropriate location on your system.

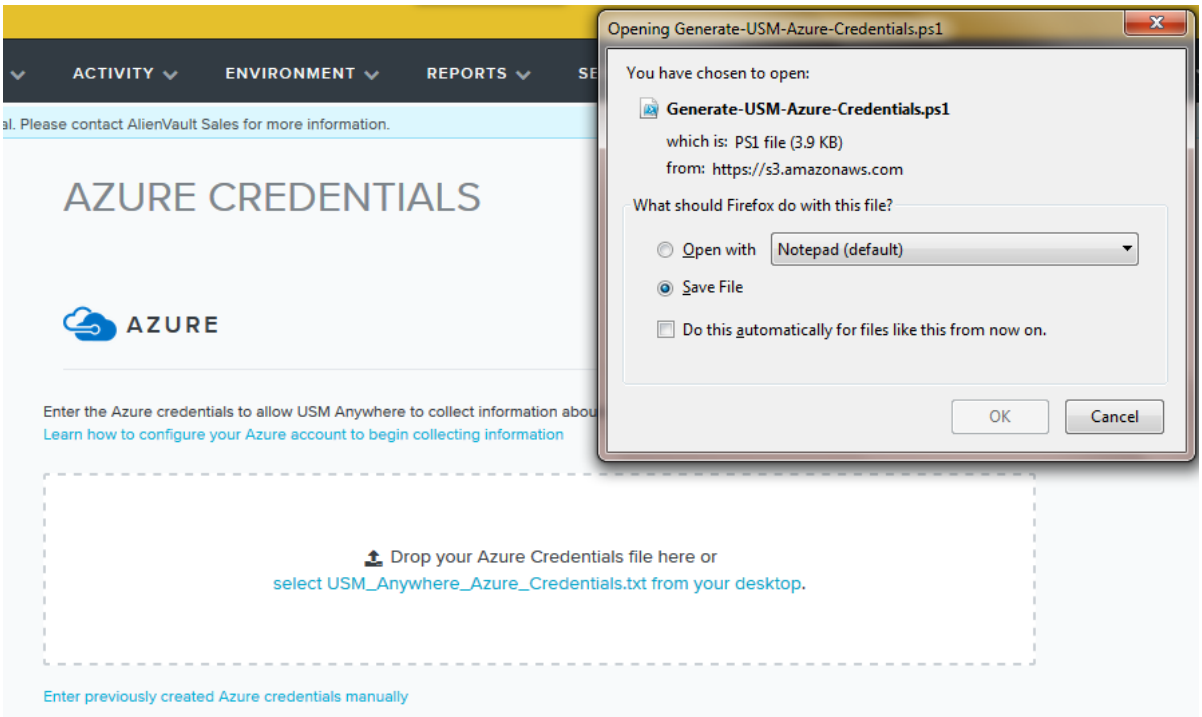

2. Run the Power Shell script as administrator on your Windows operating system (OS) from the command-line interface (CLI) shell prompt.

> **Important:** You won't be able to answer the prompts from the script if you use Windows PowerShell Integrated Scripting Environment (ISE) to run the script.

> **Note:** If you have multiple Azure subscriptions, the script prompts you to identify which one you want USM Anywhere to monitor.

When the script finishes, it creates a text file that saves to your desktop.

3. In USM Anywhere, drop the Azure credentials text file onto the displayed page or click the **select USM_Anywhere_Azure_Credentials.txt from your desktop** link to locate, select, and upload the file.

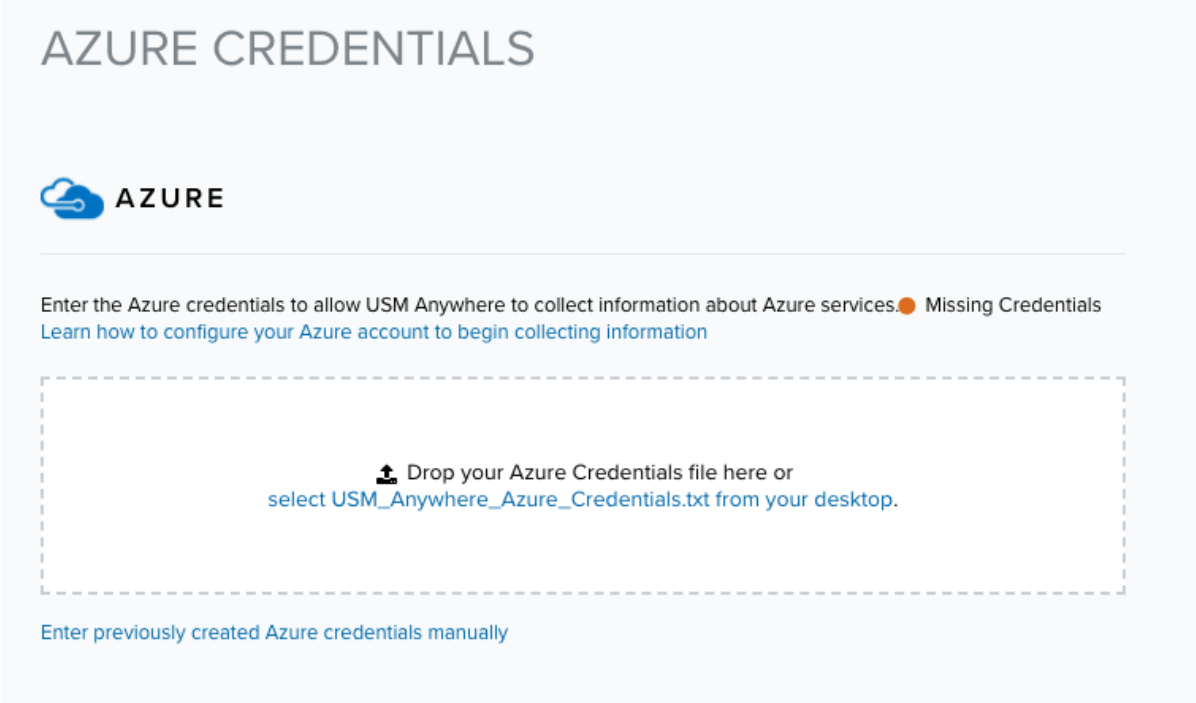

4. Verify that the status at the top of the page displays the following message:

   `Valid Credentials`

Creating the Azure Credentials Manually

To create the Azure credentials manually

1. Select **Learn how to create Azure credentials manually**.

   This opens the Create an Application and Obtain Azure Credentials page in a new browser tab or window.

2. Follow the instructions for creating the needed credentials.
3. Return to USM Anywhere, then click the **Back** button to display the first Azure Credentials page.

Configuring the Azure Credentials After Manual Credential Generation

To configure the Azure credentials after they were generated manually

> **Note:** This procedure is for non-Windows users who generated their Azure credentials manually and who are ready to configure the sensor.

1. Select the **Yes** option, and in the next page click the **Enter previously created Azure credentials manually** link at the bottom of the page.
2. Enter the Azure API credentials you generated in the Azure console into the appropriate fields.

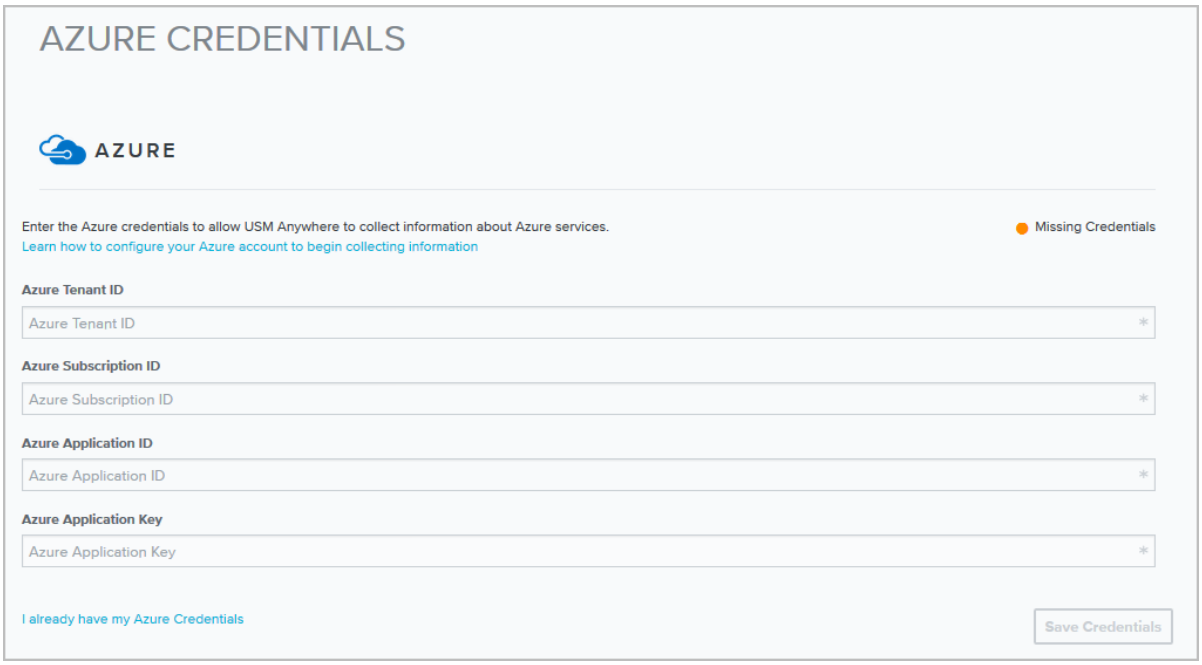

3. Click **Save Credentials**.
4. Verify that the status at the top of the page displays the following message:

```
Valid Credentials
```

When the credentials are configured, click **Next**. The wizard displays the next page in the setup process, Azure Configuration.

Azure Configuration

After you've successfully configured the Azure credentials, the Azure Configuration page opens. This page summarizes the number of Azure virtual machines (VMs), resource groups, and VM sizes in your environment.

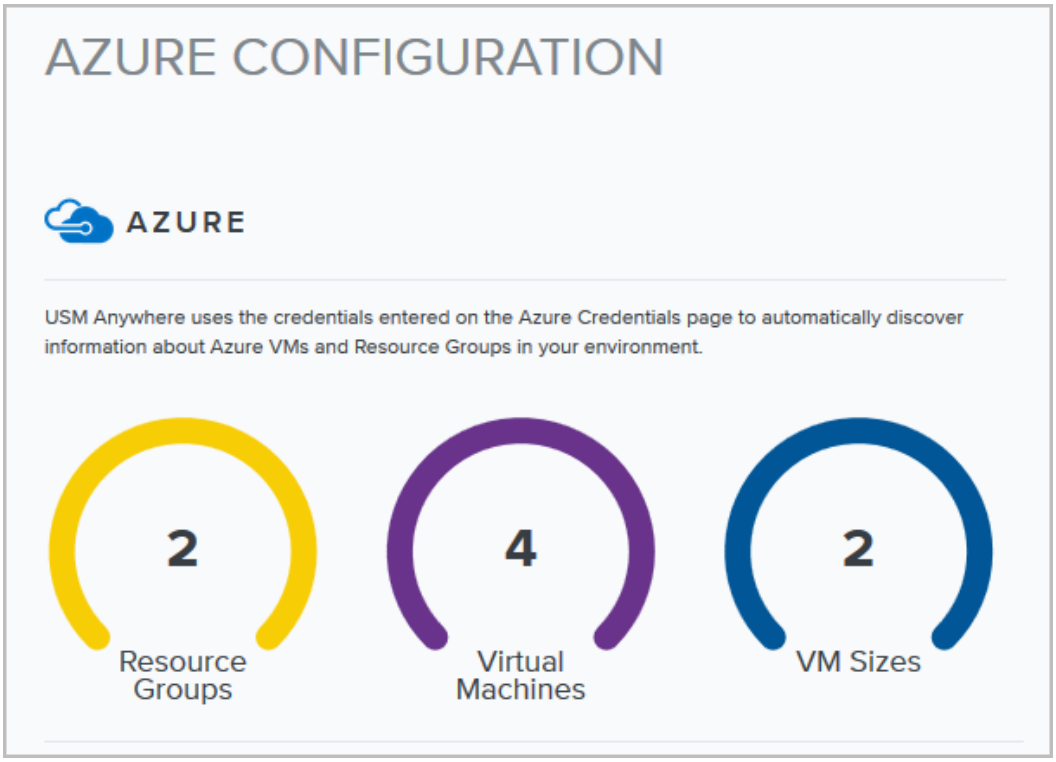

Important: If you are using VM scale sets to provide redundancy and load balancing in your Azure environment, the Azure Sensor does not automatically discover the scale set hosts through network scans. It does collect syslog from these hosts, but you must manually add the VMs to the USM Anywhere asset inventory.

See the Azure documentation for more information about virtual machine scale sets. See Adding Assets Manually in the UI in the *USM Anywhere User Guide* for detailed information about adding these VMs to the asset inventory.

Click **Next**.

The wizard displays the next page in the setup process, Azure Log Collection.

Azure Log Collection

The **Azure Log Collection** page displays the following Azure logs that are automatically discovered by USM Anywhere in your environment:

- Azure REST Monitor (formerly Azure Insight)
- Azure security alerts
- Azure SQL Server logs
- Azure Internet Information Services (IIS) logs
- Azure Windows logs

Important: The Azure SQL Server job is deprecated. Use the Event Hub Integration to collect Azure SQL Server logs. See Collect Logs from Azure Event Hubs for more information.

See Azure Log Discovery and Collection in USM Anywhere for more information about Azure log discovery and collection.

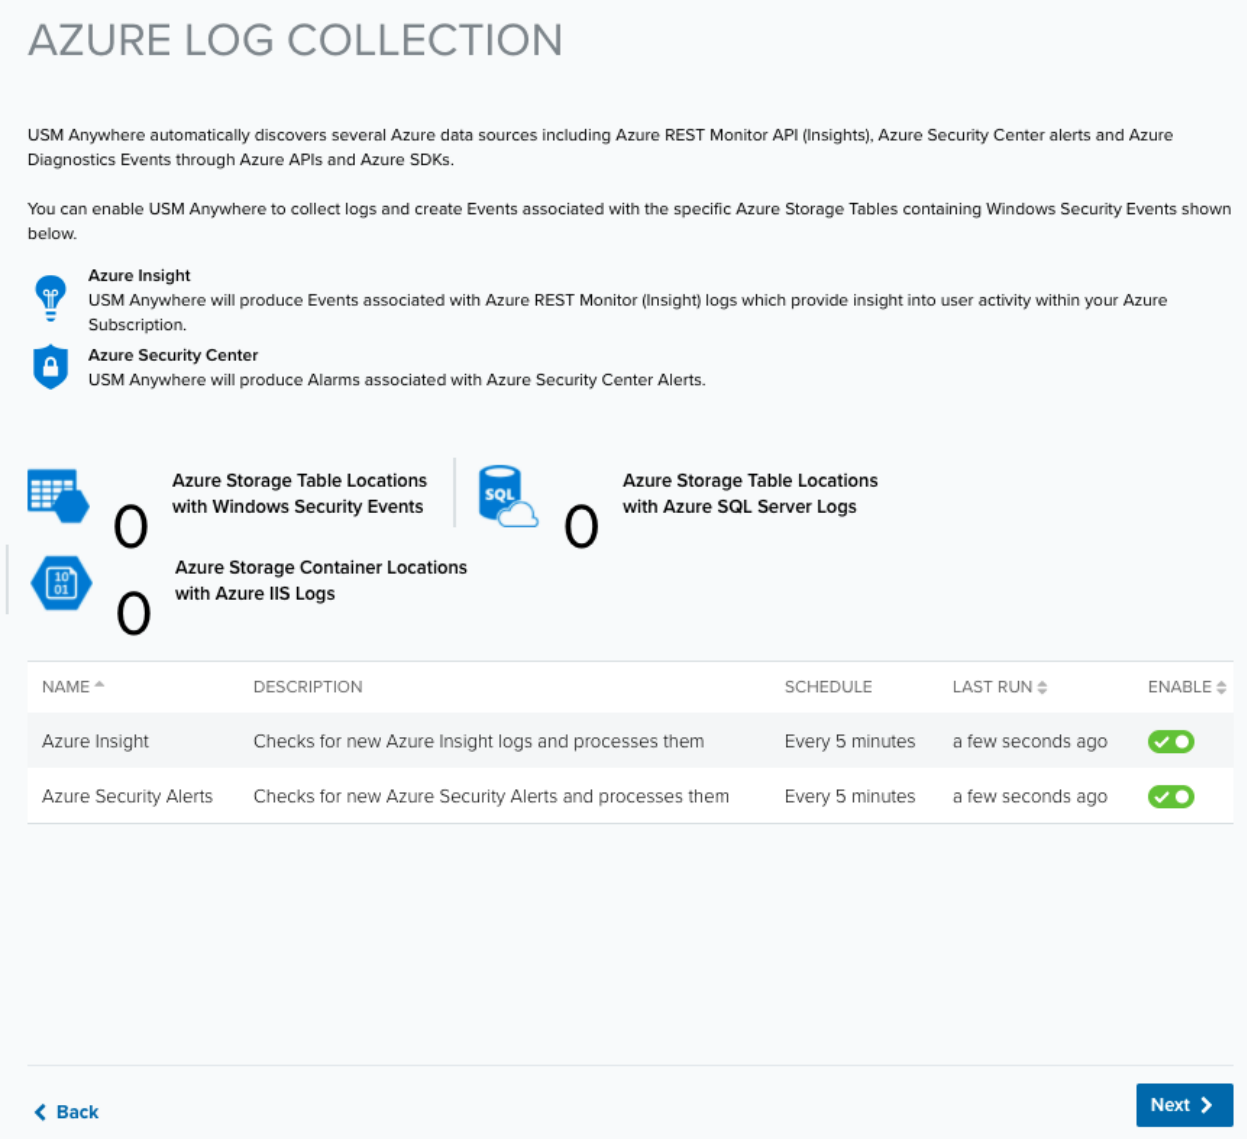

To enable these out-of-box Azure log collection jobs, toggle the gray **Enable** icon so that it turns green. When you enable any of these log collection jobs, USM Anywhere starts collecting the log data immediately according to the preconfigured frequency. See Create a New Azure Log Collection Job if you want to add other Azure log collection jobs after the sensor configuration, including jobs for Azure Web Apps.

Note: If you go to **Activity > Events** in USM Anywhere post-configuration, you can see all of the events associated with each log type, including its Event ID and many other useful details. You can also review related log collection jobs in the Job Scheduler page (**Settings > Scheduler**). See in the *USM Anywhere User Guide* for more information.

After you enable each job that you want, click **Next**.

The wizard displays the next page in the setup process, Active Directory.

Active Directory

The optional Active Directory (AD) setup page configures USM Anywhere to collect information from your AD account. To monitor Microsoft Windows systems effectively, USM Anywhere needs access to the AD server to collect inventory information.

> **Note:** This configuration is only for one AD server. If you want to scan different AD servers, you must create an AD scan job for each of them. See Scheduling Active Directory Scans from the Job Scheduler Page for more information.

AT&T Cybersecurity recommends that you create a dedicated AD account with membership in the Domain Admins group to be used by USM Anywhere to log in to the Windows systems. You also need to activate Microsoft Windows Remote Management (WinRM) in the domain controller and in all the hosts that you want to scan. You can do this by using a group policy for all the systems in your AD.

> **Important:** Before this feature is fully functional, you must configure access to the USM Anywhere Sensor on the AD server. See Granting Access to Active Directory for USM Anywhere for more information.

To complete the AD access configuration

1. Provide the AD credentials for USM Anywhere:
   - **Active Directory IP Address**: Enter the IP address for the AD server.
   - **Username**: Enter your username as admin of the account.
   - **Password**: Enter your admin's password.
   - **Domain**: Enter the domain for the AD instance.

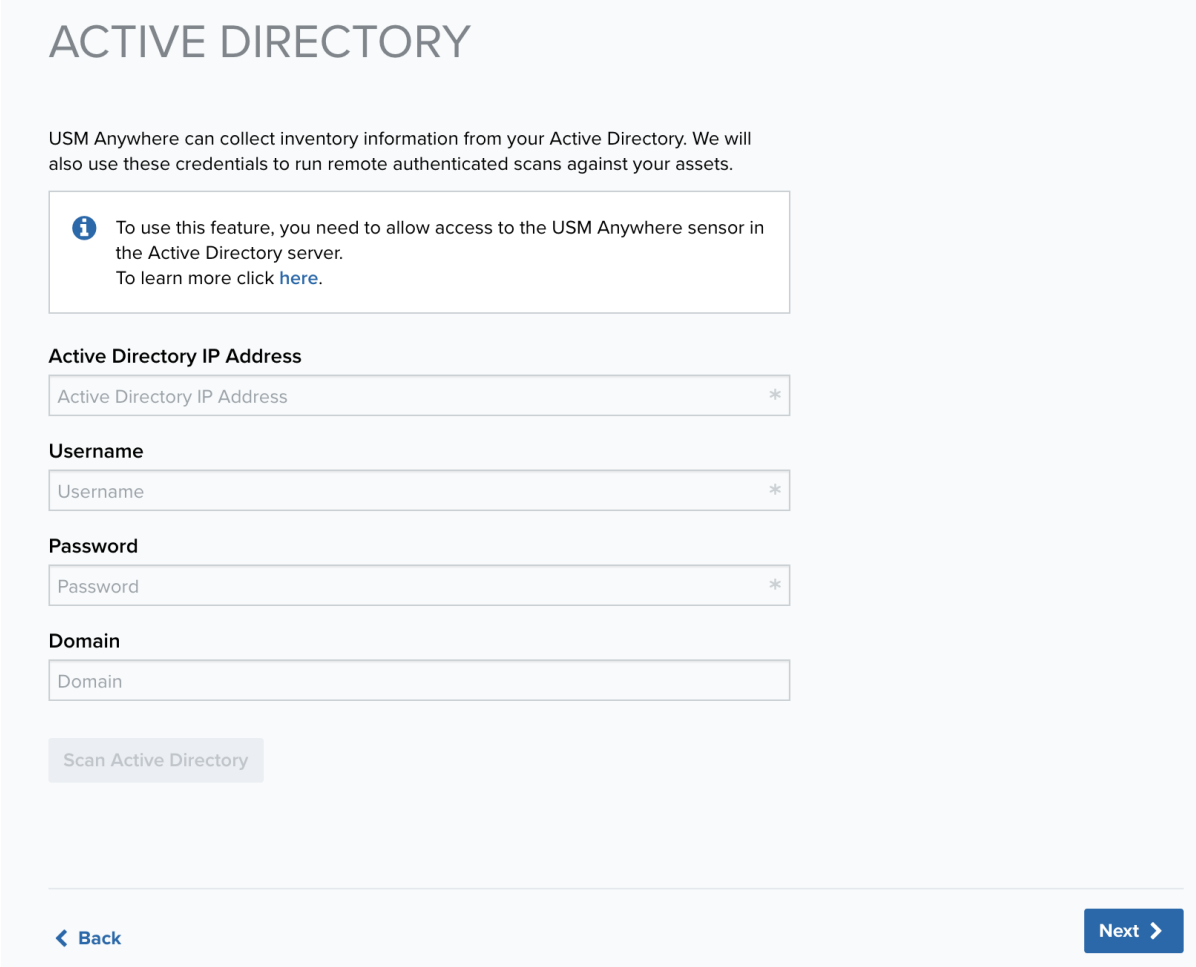

2. Click **Scan Active Directory**.

   After a successful launch of the scan, a confirmation dialog box opens.

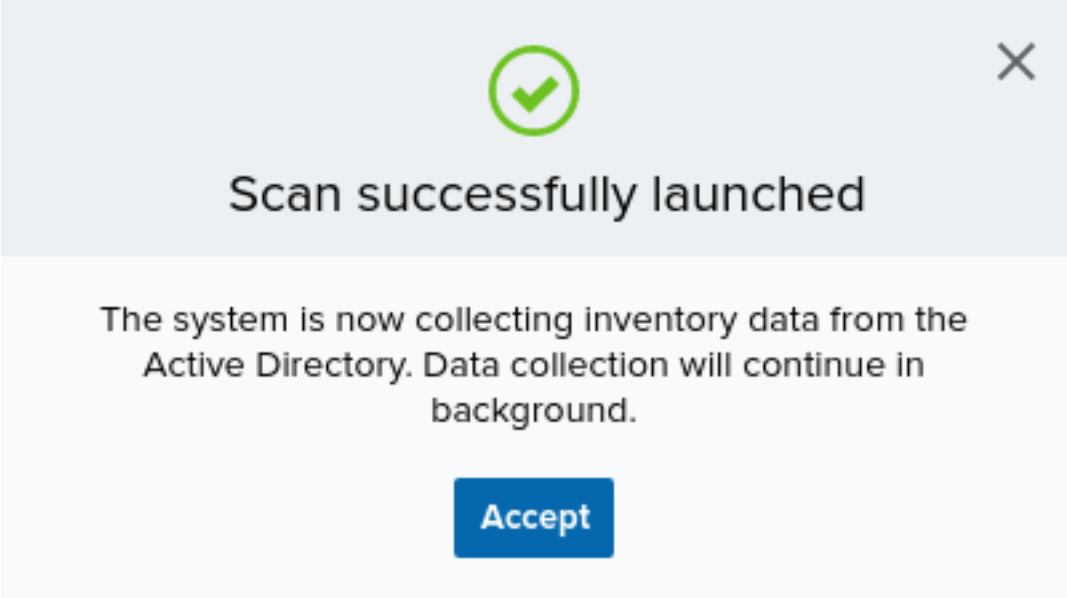

3. Click **Accept**.

   The scan continues in the background.

   Upon completion, another dialog box opens and provides information about the number of assets USM Anywhere discovered. It also prompts you to decide if you want to scan for hosts and services running in your environment.

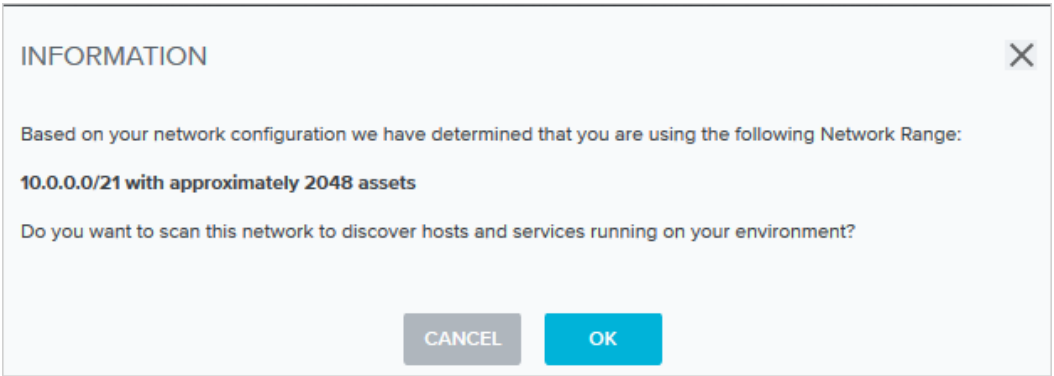

   Click **Cancel** to opt out of this scan.

4. (Optional.) If you want to scan for other hosts and services, click **OK**.

5. Click **Next** after the scan ends.

   The wizard opens the next page in the setup process, Log Management.

Log Management

On the Log Management page are syslog port numbers. (These ports are the same for all USM Anywhere Sensors.)

USM Anywhere collects third-party device, system, and application data through syslog over UDP on port 514 and over TCP on ports 601 or 602 by default. It collects Transport Layer Security (TLS)-encrypted data through TCP on ports 6514 or 6515 by default. These ports support the RFC 3164 and RFC 5424 formats. To configure any third-party devices to send data to USM Anywhere, you must provide the IP address and the port number of your USM Anywhere Sensor.

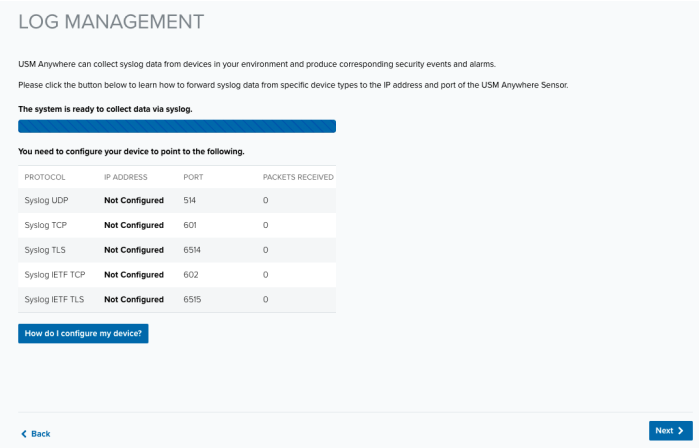

To enable log collection and configure your log management

1. Make sure that you have granted the necessary permissions for your OS to allow USM Anywhere to access its logs. You can also integrate a wide variety of data sources to send log data over syslog to the USM Anywhere Sensor.

   To learn how to configure your operating systems and supported third-party devices to forward syslog log data, see the following related topics:

   - **The Syslog Server Sensor App**: Log collection (UDP, TCP, and TLS-encrypted TCP) from rsyslog
   - **Collecting Linux System Logs**: Log collection from a Linux system
   - Collecting Windows System Logs: Log collection from a Windows system
   - Go to the specific AlienApp in USM Anywhere for instructions about syslog forwarding

> **Note:** Because the log scan can take some time, you might not see all of the automatically discovered log sources immediately after deploying the first USM Anywhere Sensor.

2. When you have finished the log collection setup and integrated any needed plugins, verify that the data transfer is occurring.
3. Click **Next** when this step is complete.

OTX

AT&T Alien Labs™ Open Threat Exchange® (OTX™) is an open information-sharing and analysis network providing users with the ability to collaborate, research, and receive alerts on emerging threats and indicators of compromise (IoCs) such as IP addresses, file hashes, and domains.

You must have an OTX account to receive alerts based on threats identified in OTX. This account is separate from your USM Anywhere account. Go to The World's First Truly Open Threat Intelligence Community to create an OTX account.

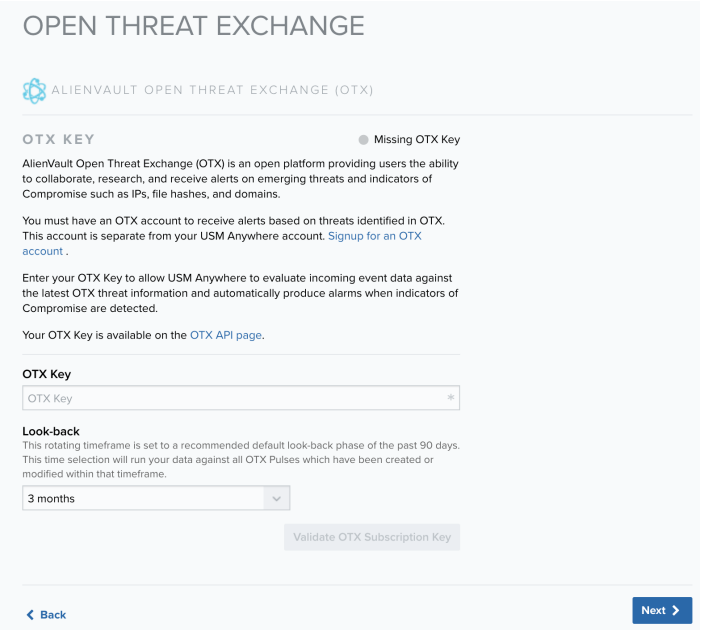

Note: If you do not already have an OTX account, click the **Sign up** link. This opens another browser tab or window that displays the OTX signup page. After you confirm your email address, you can log in to OTX and retrieve the unique API key for your account.

See Open Threat Exchange® and USM Anywhere in the *USM Anywhere User Guide* for more information about OTX integration in USM Anywhere.

To enable USM Anywhere to evaluate event data against the latest OTX intelligence

1. Log in to OTX and open the API page (https://otx.alienvault.com/api).
2. In the DirectConnect API Usage pane, click the icon to copy your unique OTX connection key.

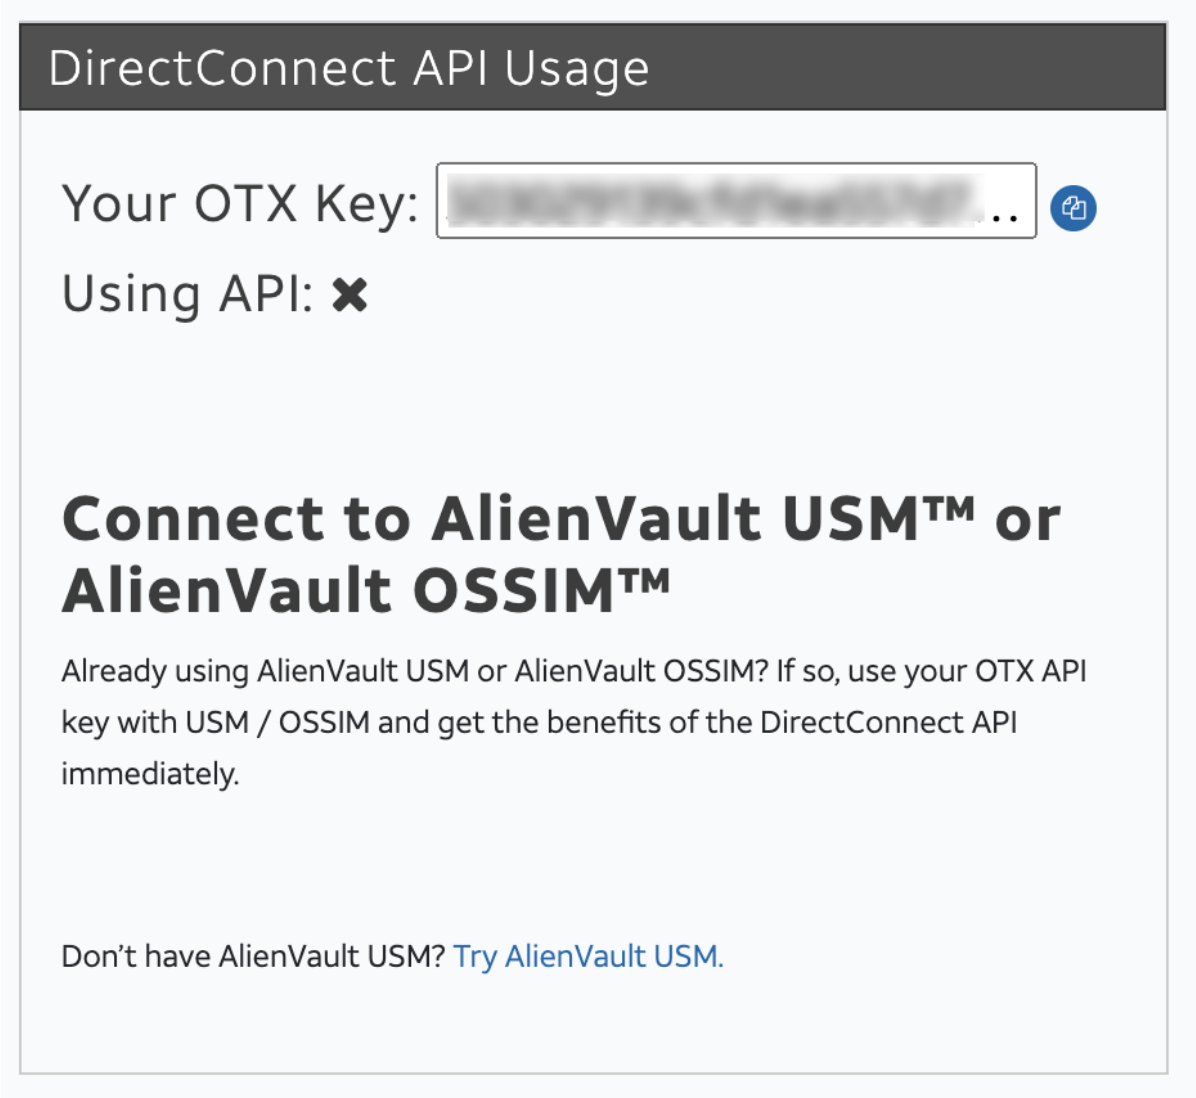

3. Return to the Open Threat Exchange (OTX) page of the USM Anywhere Sensor Setup Wizard and paste the value in the OTX Key text box.

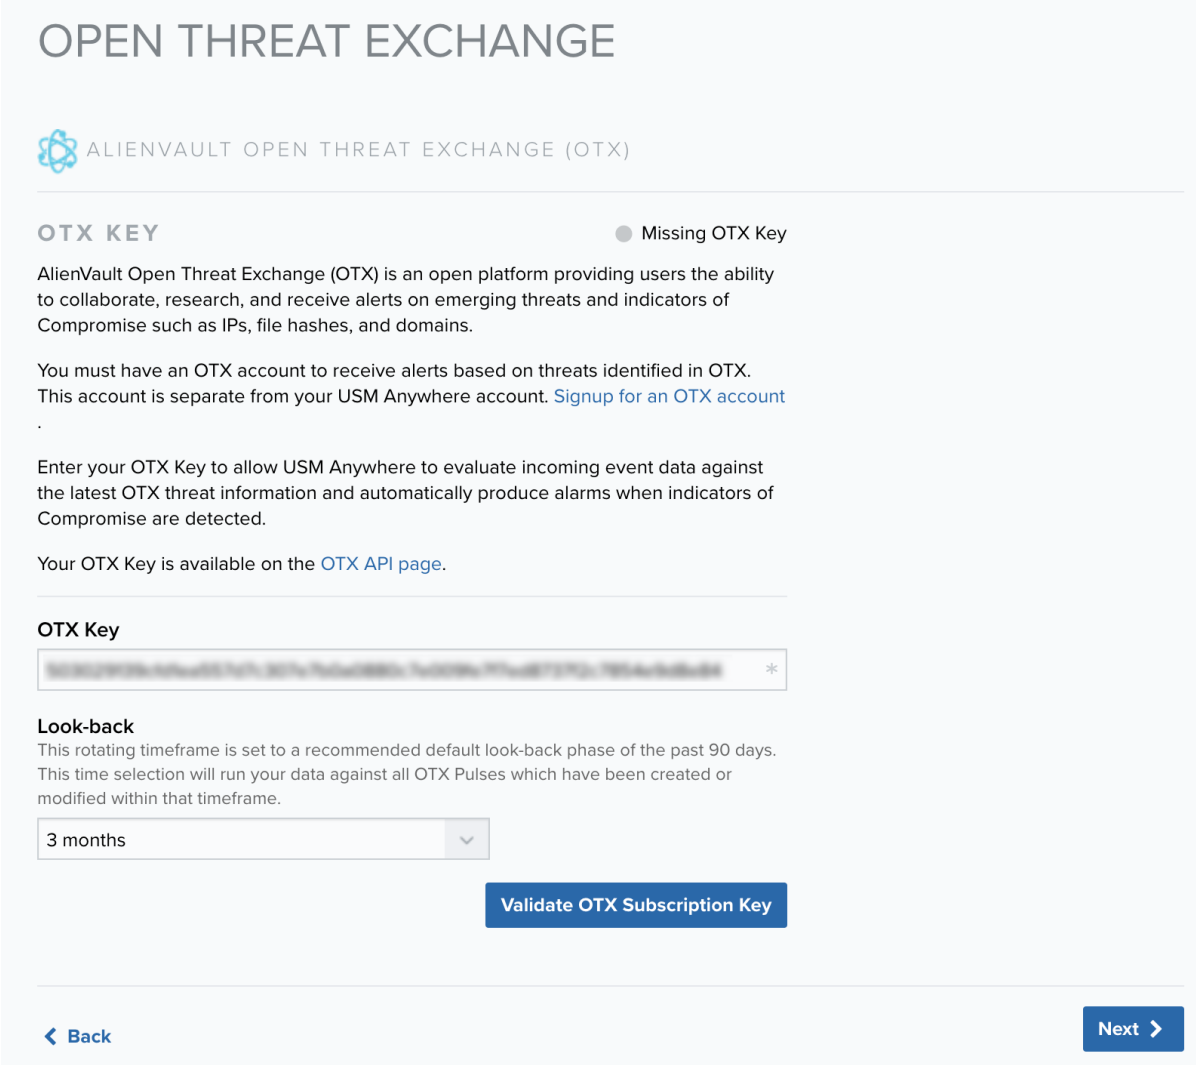

4. Click **Validate OTX Subscription Key**.

   With a successful validation of the key, the status at the top of the page changes to "Valid OTX key".

5. Click **Next** when this task is complete.

Setup Complete

The Congratulations page summarizes the status of your configuration.

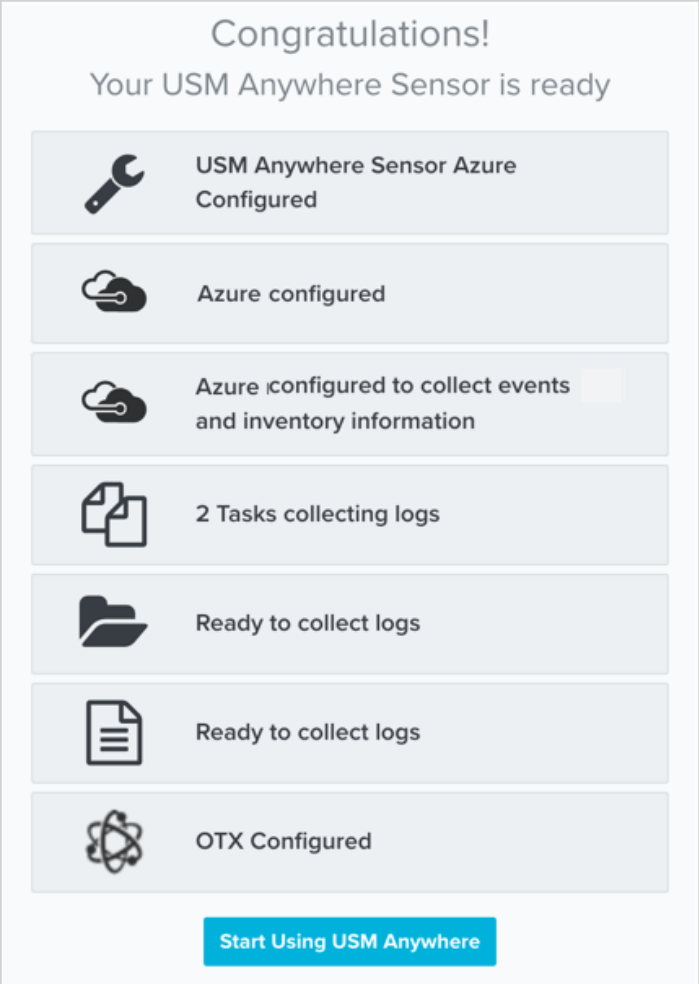

Click **Start Using USM Anywhere**, which takes you to the Overview dashboard.

Next...

Now is a great time to run a vulnerability scan. See Vulnerability Assessment in the *USM Anywhere User Guide* for detailed information about running a vulnerability scan.

Azure Log Discovery and Collection in USM Anywhere

With a USM Anywhere Sensor deployed in your Microsoft Azure environment, referred to as the Azure Sensor, USM Anywhere can discover and collect logs in two different ways.

An Azure Sensor is preconfigured to automatically discover and collect these types of Azure resource logs (previously referred to as diagnostic logs):

- Azure Monitor (Insight)
- Azure Security Alerts
- Azure Internet Information Services (IIS) logs
- Azure SQL Server logs
- Azure Web Apps logs
- Azure Windows logs

See Collect Azure Resource Logs for more information.

Furthermore, if you stream data to Azure Event Hubs, you can connect an Azure Sensor to your event hub and collect the following logs:

- Azure Active Directory (AD) logs, including audit logs and sign-in logs
- Azure Monitor logs
- Azure SQL Database logs
- Microsoft Defender Advanced Threat Protection (ATP) logs

See Collect Logs from Azure Event Hubs for more information.

Collect Azure Resource Logs

Microsoft Azure resource logs (previously referred to as diagnostic logs) provide insight into operations performed within an Azure resource, such as Microsoft Azure Internet Information Services (IIS) or Microsoft Azure SQL Server. USM Anywhere discovers and collects these logs through the Azure APIs. A USM Anywhere Sensor deployed in your Azure environment is preconfigured to automatically discover logs from your Azure storage account. You can enable or disable the predefined jobs from the Azure Sensor Setup Wizard (see Azure Log Collection) or within the USM Anywhere scheduler (see USM Anywhere Scheduler).

To supplement the default log location or to add log collection for Microsoft Azure Web Apps, you can create custom log collection jobs that operate through the Azure Sensor app.

Note: What an Azure log job collects depends on whether you granted contributor permissions to one of your resources or to your entire Azure subscription for the USM Anywhere application. Depending on the Azure credentials configured for the deployed Azure Sensor, the sensor could have access to individual resource groups or the whole subscription. See Create an Application and Obtain Azure Credentials for more information.

Microsoft Azure Monitor (Insight)

Microsoft Azure Monitor (formerly Azure Insights) provides base-level infrastructure metrics and logs for most services in Azure. It helps you to track user activities within an Azure subscription, including when users log on, deploy or shut down virtual machines (VMs), and more. Through the Microsoft Azure Monitor Representational State Transfer (REST) API, USM Anywhere captures those logs and creates events.

You need to perform a specific configuration of Azure Monitor in the Azure console for USM Anywhere to collect the Azure-related logs. You need to enable the archive to a storage account option on the Azure subscription, which then enables USM Anywhere to automatically detect and create a job for the Azure-related jobs. When you complete the Log Collection step for your Azure Sensor setup, you can enable this default job, which runs every 20 minutes.

Diagnostic setting name LogtoStoageAndEventHub

Category details

log

- [x] Administrative
- [x] Security
- [x] ServiceHealth
- [x] Alert
- [x] Recommendation
- [x] Policy
- [x] Autoscale
- [x] ResourceHealth

Destination details

- [ ] Send to Log Analytics workspace
- [x] Archive to a storage account

Cannot find resource 'defaultstorageaccountea'. Either you do not have permission or this resource no longer exists.

Location
All

Subscription
[redacted]

Storage account *
[redacted]

- [x] Stream to an event hub

For potential partner integrations, see documentation here

Subscription
[redacted]

Event hub namespace *
[redacted]

You can also enable or disable this default job in the USM Anywhere Scheduler page. When you select the job in this page, you can review the history for the scheduled job.

Azure-Sensor Azure	Azure	Azure Insight	Checks for new Azure I nsight logs and process es them	Every 5 minutes	3 minutes ago	

Process Azure Insight Logs

SCHEDULE PARAMETERS

SENSOR	Azure-Sensor Azure

SCHEDULE HISTORY

EVENT ACTION	EVENT NAME	FINISHED	FINISHED RELATIVE
Successful	Azure Insight Events: Processed 17 new events.	Tue, May 17 2022, 04:26 PM	3 minutes ago
Successful	Azure Insight Events: Processed 13 new events.	Tue, May 17 2022, 04:21 PM	8 minutes ago
Successful	Azure Insight Events: Processed 3 new events.	Tue, May 17 2022, 04:16 PM	13 minutes ago
Successful	Azure Insight Events: Processed 4 new events.	Tue, May 17 2022, 04:11 PM	18 minutes ago
Successful	Azure Insight Events: Processed 4 new events.	Tue, May 17 2022, 04:06 PM	23 minutes ago
Successful	Azure Insight Events: Processed 6 new events.	Tue, May 17 2022, 04:01 PM	28 minutes ago
Successful	Azure Insight Events: Processed 2 new events.	Tue, May 17 2022, 03:56 PM	33 minutes ago
Successful	Azure Insight Events: Processed 4 new events.	Tue, May 17 2022, 03:51 PM	38 minutes ago
Successful	Azure Insight Events: Processed 4 new events.	Tue, May 17 2022, 03:46 PM	43 minutes ago
Successful	Azure Insight Events: Processed 6 new events.	Tue, May 17 2022, 03:41 PM	an hour ago

1 - 10 of 199,507 < Previous | 1 2 3 4 5 | Next >

Azure Security Alerts

Microsoft Azure Security Center is an Azure service that continuously monitors your Azure environment and applies analytics to automatically detect a wide range of potentially malicious activity. It surfaces these detections as *security alerts*. Security Center performs this function by collecting data from your VMs, which is enabled for all VMs in your subscription by default. You can also customize this data collection in the Security Center policy.

You do not need to perform a specific configuration of the Azure Security Center alerts in the Azure console to be able to collect these logs. USM Anywhere automatically detects these logs and creates a job for Azure Security Center alerts logs. When you complete the Log Collection step for your Azure Sensor setup, you can enable this default job, which runs every 20 minutes.

You can also enable or disable this default job in in the USM Anywhere Scheduler page. When you select the job in this page, you can review the history for the scheduled job.

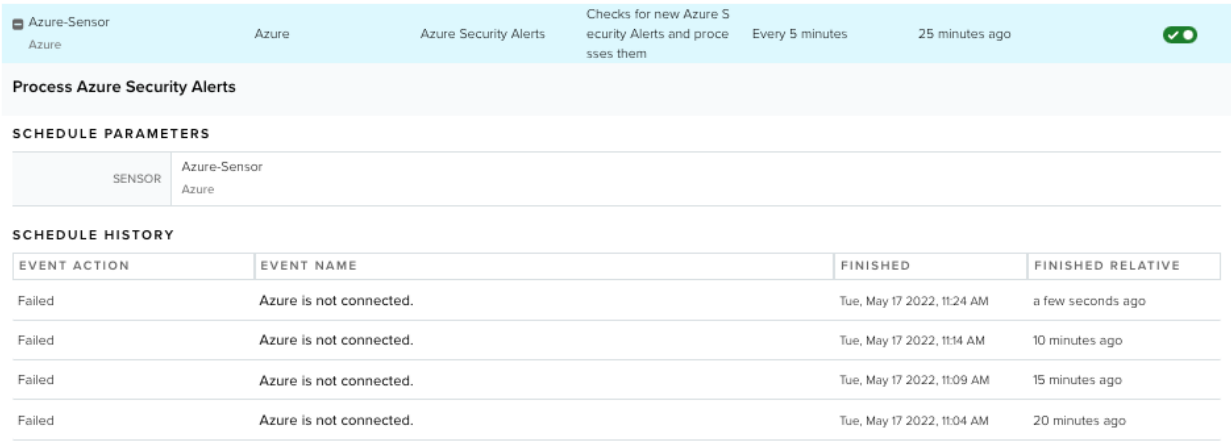

Azure IIS Logs

For individual VMs running IIS with Azure diagnostics enabled, you can designate storage for the IIS logs. USM Anywhere automatically detects these logs through the Azure APIs and Azure software development kits (SDKs). For each Azure Storage container locations with Azure IIS logs that it detects, USM Anywhere creates a default log collection job. When you complete the Log Collection step for your Azure Sensor setup, you can enable these default jobs, which run every five minutes.

Warning: If there are network restrictions in your environment restricting access to the storage account, those restrictions must allow access to the sensor.

Note: This type of IIS implementation is different than Azure Web Apps, which is a platform service and uses a different logging configuration. See Azure Web Apps Logs for information about collecting logs for web apps.

You can also enable or disable this default job in the Job Scheduler. When you select the job in this page, you can review the history for the scheduled job. You could choose to disable this default job based on the IIS log locations that USM Anywhere discovers, and create a custom Azure IIS log collection job for a location that you specify.

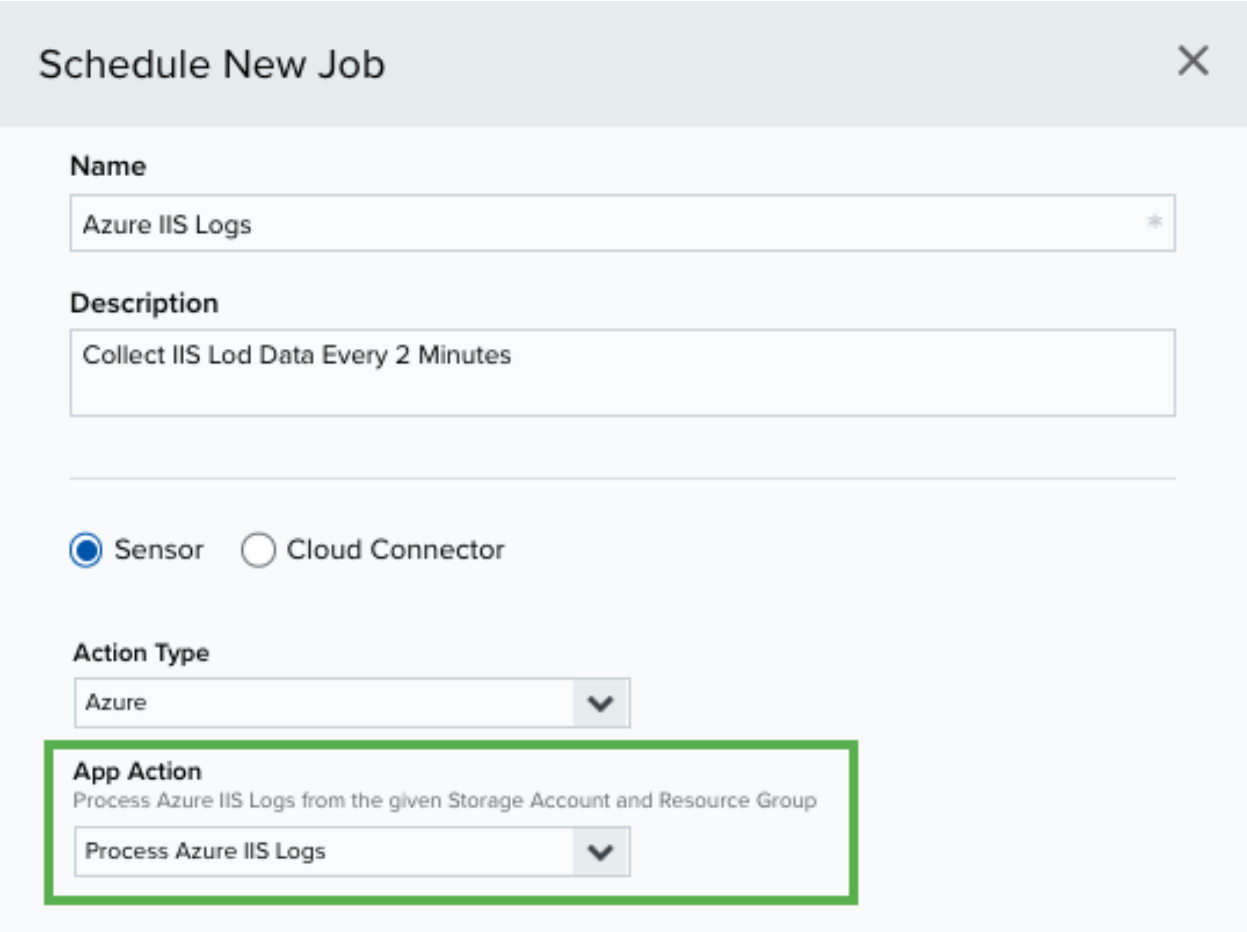

When you configure the new job, set the App Action option to **Process Azure IIS Logs**. You must also specify the **Resource Group**, **Storage Account**, and **Blob Container** for the custom log collection job. See Create a New Azure Log Collection Job for more information about scheduling an Azure log collection job.

Azure SQL Server Logs

For individual VMs running an Azure SQL Server with Azure diagnostics enabled, you can designate storage for the IIS logs. You must configure this to use Microsoft Azure Table storage. To simplify the tracking of related security issues, USM Anywhere treats the SQL service as an asset, and maps events and other security issues directly with the SQL service. When it detects Azure Table storage locations with Azure SQL Server logs, USM Anywhere creates a default log collection job for each. When you complete the Log Collection step for your Azure Sensor setup, you can enable these default jobs, which run every five minutes.

Important: The Azure SQL Server job is deprecated. Use the Event Hub Integration to collect Azure SQL Server logs. See Collect Logs from Azure Event Hubs for more information.

If you want to supplement this automatic Azure log collection in USM Anywhere, you can create an additional Azure SQL Server log collection job.

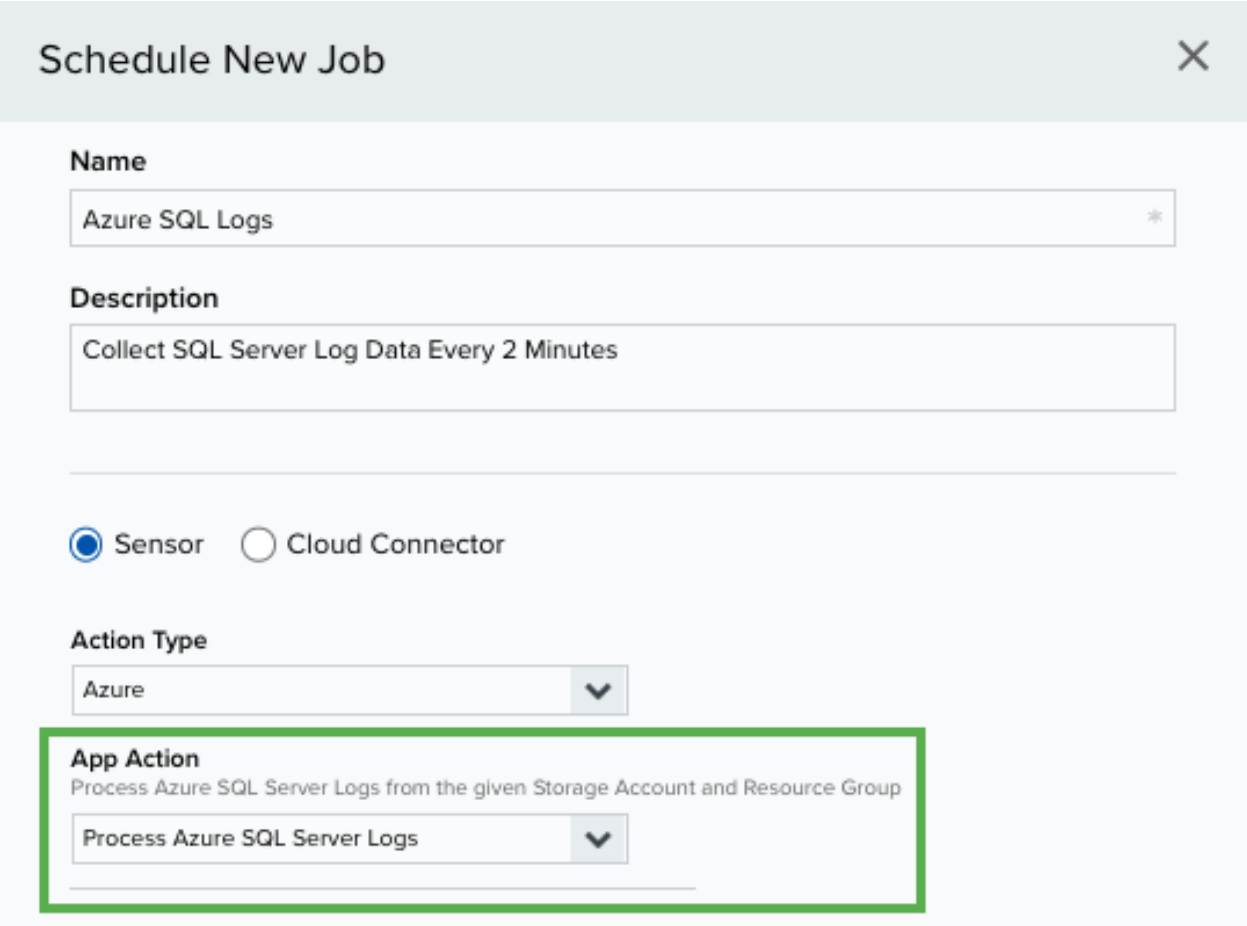

When you configure the new job, set the App Action option to **Process Azure SQL Server Logs**. You must also specify the **Resource Group**, **Storage Account**, and **Table Container** for the custom log collection job. See Create a New Azure Log Collection Job for more information about creating a new Azure log collection job.

Azure Web Apps Logs

Warning: If there are network restrictions in your environment restricting access to the storage account, those restrictions must allow access to the sensor.

Azure App Service Web Apps is a fully managed compute platform that is optimized for hosting websites and web applications. A web app represents the compute resources that Azure provides for hosting a website or web application. These compute resources may be on shared or dedicated VMs. For each deployed web application in your Azure environment, you can enable diagnostic logging to capture and store the web server and application information.

Important: When configuring Azure Web Apps logs, you must use the World Wide Web Consortium (W3C) format and select the following fields:

date, time, s-sitename, cs-method, cs-uri-stem, cs-uri-query, s-port, cs-username, c-ip, cs(User-Agent), cs(Cookie), cs(Referer), cs-host, sc-status, sc-substatus, sc-win32-status, sc-bytes, cs-bytes, time-taken

Unlike the other supported Azure logs, the USM Anywhere Sensor does not perform an automatic discovery job for Web Apps to look for the storage location. If you want USM Anywhere to collect the log data for your Web Apps, you must create a new log job and specify the storage location parameters.

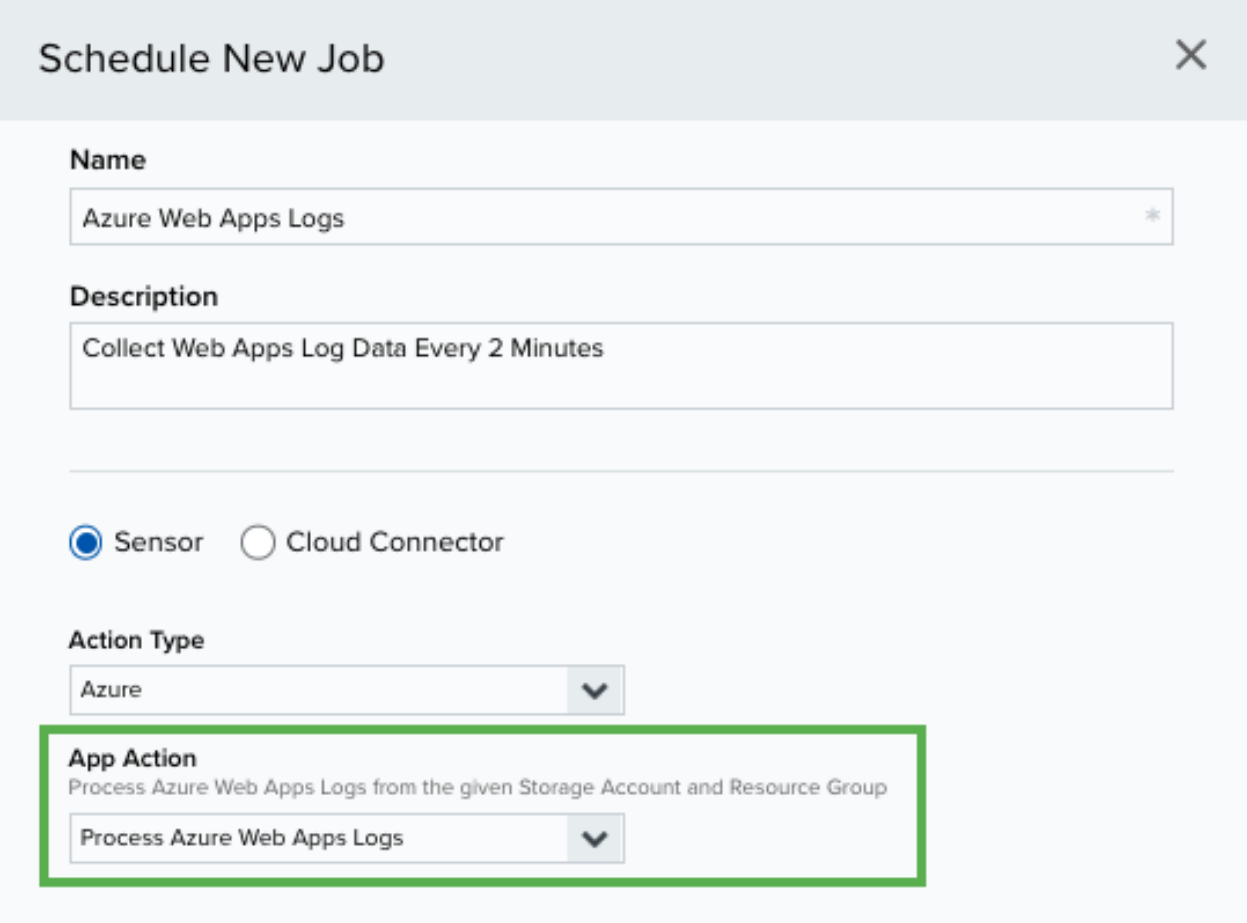

When you configure the new job, set the App Action option to **Process Azure Web Apps Logs**. You must also specify the **Resource Group**, **Storage Account**, and **Blob Container** for the custom log collection job. See Create a New Azure Log Collection Job for more information about creating a new Azure log collection job.

Azure Windows Logs

Warning: If there are network restrictions in your environment restricting access to the storage account, those restrictions must allow access to the sensor.

For individual VMs running Microsoft Windows with Azure diagnostics enabled, Azure stores the Windows Events logs by default. USM Anywhere automatically detects these logs through Azure APIs and Azure SDKs. When it detects Azure Storage container locations with Azure Windows logs, USM Anywhere creates a default log collection job for each. When you complete the Log Collection step for your Azure Sensor setup, you can enable these default jobs, which run every five minutes.

If you want to supplement this automatic Azure log collection in USM Anywhere, you can create an additional Azure Windows log collection job.

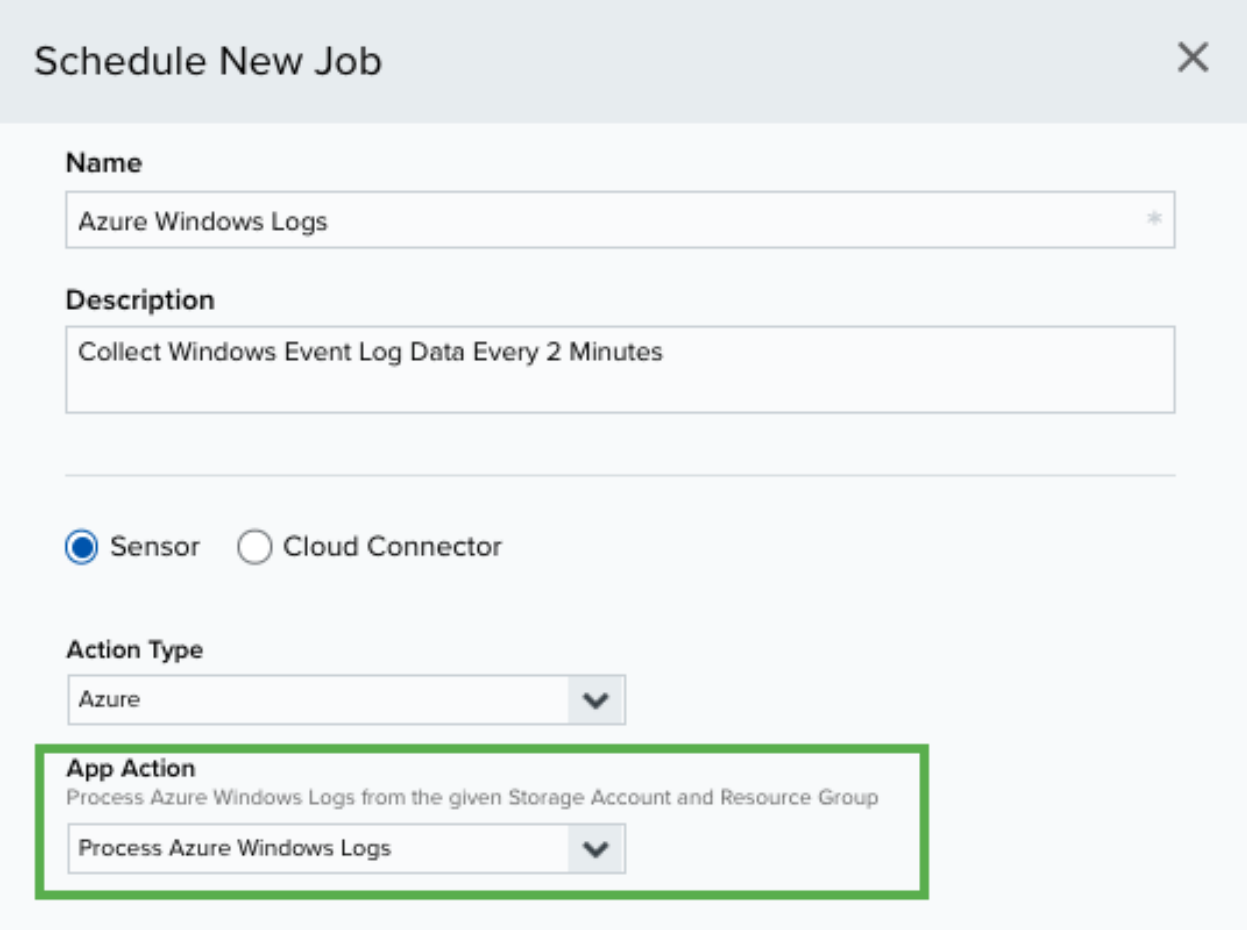

When you configure the new job, set the App Action option to **Process Azure Windows Logs**. You must also specify the **Resource Group**, **Storage Account**, and **Blob Container** for the custom log collection job. See Create a New Azure Log Collection Job for more information about creating a new Azure log collection job.

Enable Diagnostics for Azure Web Apps

Role Availability	Read-Only	Investigator	Analyst	Manager

If you have Azure Web Apps running in your Azure environment, you can enable diagnostics logging for these web apps in the Azure console and then create log collection jobs in USM Anywhere to retrieve and process the log data.

The Azure App Service web apps provide diagnostic functionality for logging information from both the web server and the web application. It logically separates this into web server diagnostics and application diagnostics. When you enable this feature in Azure, you specify a log data storage account and container for each of these. See the Microsoft Azure documentation at https://docs.microsoft.com/en-us/azure/app-service/web-sites-enable-diagnostic-log for more information.

To enable diagnostics for your Azure Web App

1. Log in to your account at https://portal.azure.com/.
2. Go to your Azure Web App and select **Settings > Diagnostics logs**.
3. For **Application Logging (Blob)**, click **On** and set the parameters:
   - Set the **Level** for the logging.
   - For **Storage Settings**, click **>** and select the **Storage Account** and **Container**.

     This is the storage account and container that Azure will use to store logs for the Web App. Make note of this information because you will need it to set up a log collection job in USM Anywhere. You can click **+ Storage Account** to create a new storage account or container, or select an existing one.
4. For Web server logging, select **Storage**.
5. Click **Storage Settings** and select the same storage account and container that you set

for the application logging.

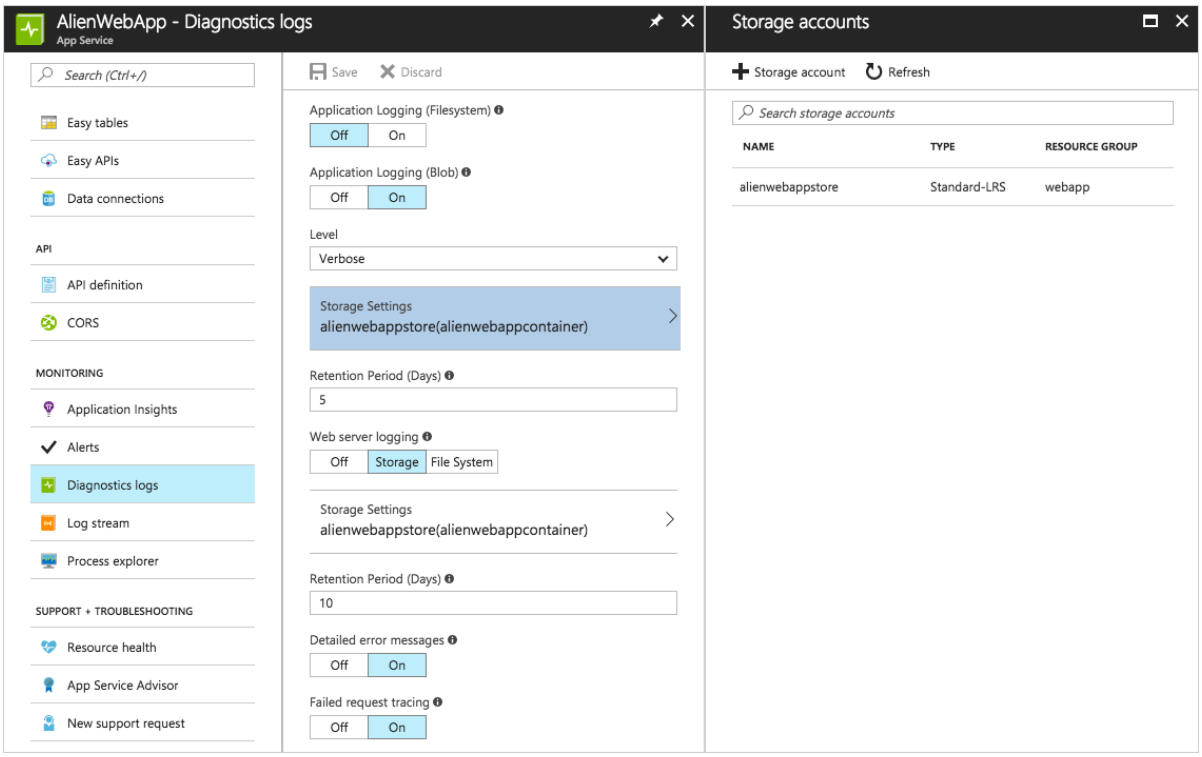

6. Click **Save**.

Create a New Azure Log Collection Job

Role Availability	Read-Only	Investigator	Analyst	Manager

USM Anywhere automatically creates log collection jobs for Azure Monitor and security logs. It also creates jobs for Internet Information Services (IIS), Microsoft Azure SQL Server, and Microsoft Windows if it detects storage locations for these log types. When you complete the Log Collection step for the Azure Sensor, you can enable these default jobs. You can review these jobs and their history in the Scheduler, but you cannot modify the parameters of these default jobs.

Note: What an Azure log job collects depends on whether you granted contributor permissions to one of your resources or to your entire Azure subscription for the USM Anywhere application. Depending on the Azure credentials configured for the deployed Azure Sensor, the sensor could have access to individual resource groups or the whole subscription. See Create an Application and Obtain Azure Credentials for more information.

To supplement the automatic Azure log collection in USM Anywhere and to set up log collection for Azure Web Apps, add new Azure log collection jobs.

Important: Before your scheduled jobs can collect logs, you may also have to perform specific configuration steps outside of USM Anywhere in your environment. See Collect Azure Resource Logs for detailed descriptions of the configuration steps your environment might require.

To schedule a new job to collect and process Azure logs

1. Go to **Settings > Scheduler**.
2. In the left navigation menu, click **Log Collection**.

   Note: You can use the Sensor filter at the top of the list to review the available log collection jobs on your Azure Sensor.

3. Click **Create Log Collection Job**.

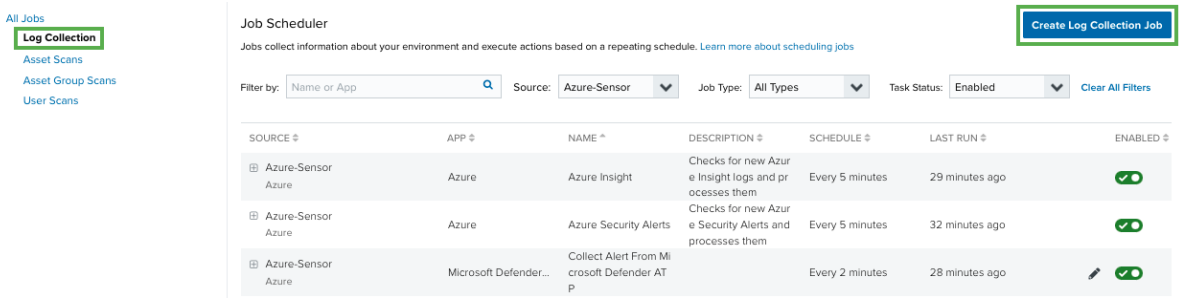

Note: If you have recently deployed a new USM Anywhere Sensor, it can take up to 20 minutes for USM Anywhere to discover the various log sources. After it discovers the logs, you must manually enable the Azure log collection jobs you want before the system collects the log data.

The Schedule New Job dialog box opens.

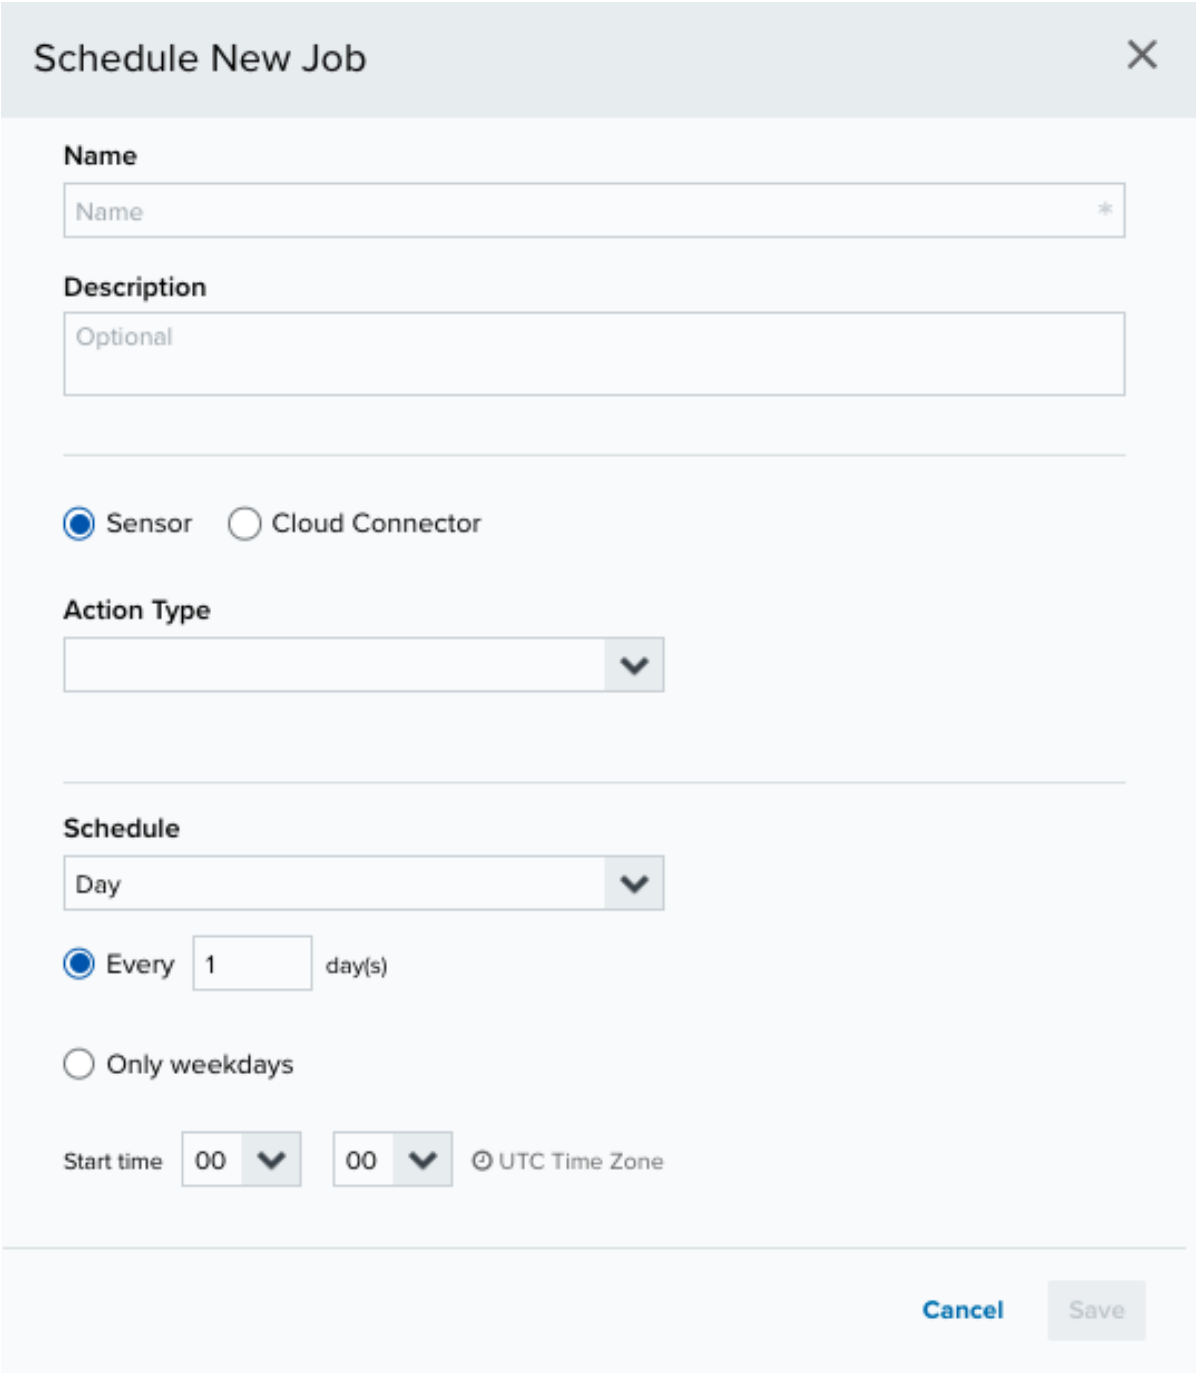

4. Enter the name and description for the job.

   The description is optional, but it is a best practice to provide this information so that others can easily understand what it does.

5. Select **Sensor** as the source for your new job.

6. In the Select App option, select **Azure**.
7. In the App Action option, select the action for Azure log type that you want to schedule for collection.

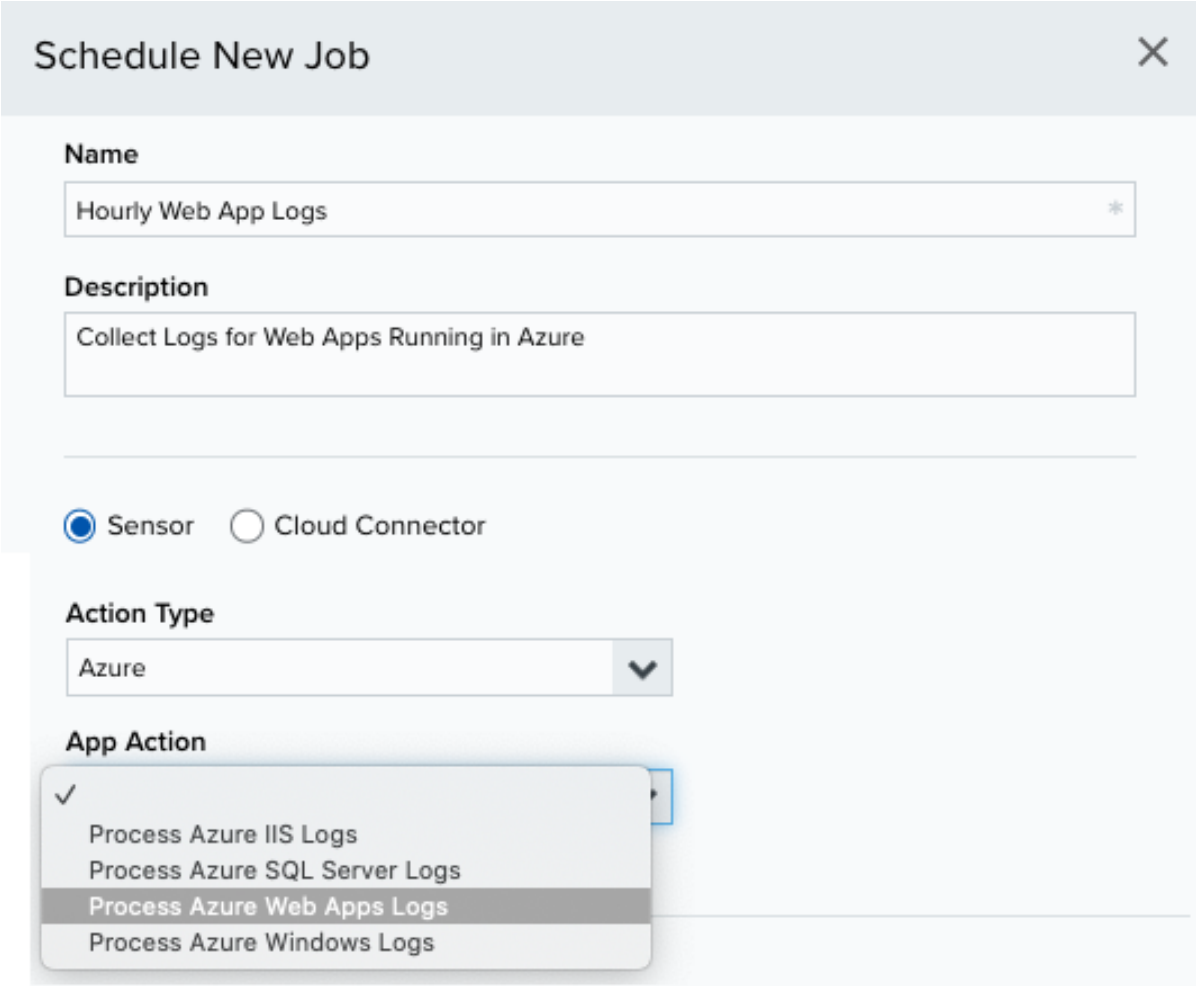

See Collect Azure Resource Logs to review details about the Azure log types that USM Anywhere can collect.

8. Depending on the selected app action (log type), specify the **Resource Group**, **Storage Account**, and **Container** for the logs.

   You can obtain this information by logging into the Azure console and reviewing the configuration for your diagnostic and storage resources.

> **Note:** For Azure IIS logs, Azure Web Apps logs, and Azure Windows logs, you must specify a binary large object (BLOB) container used for the log storage. For the Azure SQL Server log type, you must specify the table container used for the log storage.
>
> The Azure SQL Server job is deprecated. Use the Event Hub Integration to collect Azure SQL Server logs. See Collect Logs from Azure Event Hubs for more information.

9. In the Schedule section, specify when USM Anywhere runs the job:

   a. Select the increment as **Minute**, **Hour**, **Day**, **Week**, **Month**, or **Year**.

   > **Warning:** After a frequency change, monitor the system to check its performance. For example, you can check the system load and CPU. See USM Anywhere System Monitor for more information.

   b. Set the interval options for the increment.

   The selected increment determines the available options. For example, on a weekly increment, you can select the days of the week to run the job.

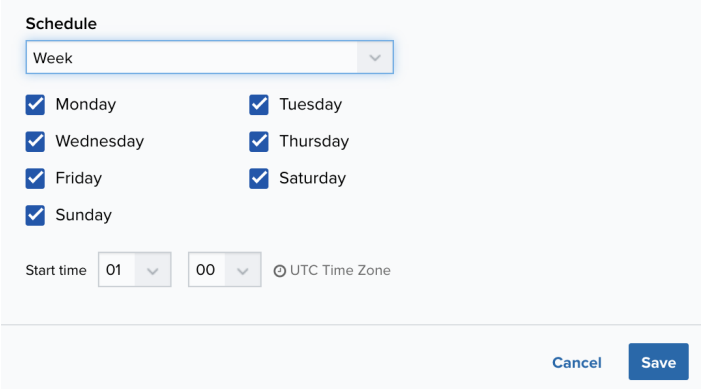

Or on a monthly increment, you can specify a date or a day of the week that occurs within the month.

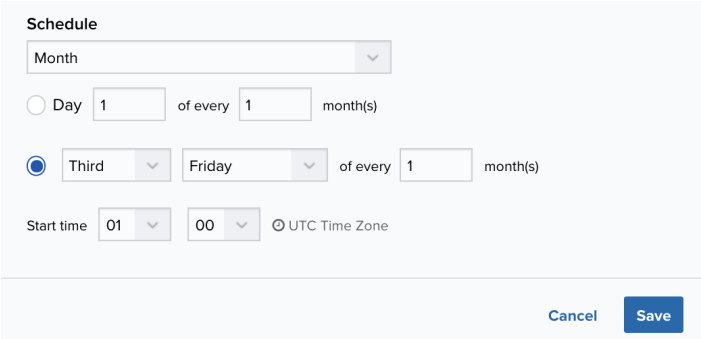

Important: USM Anywhere restarts the schedule on the first day of the month if the option "Every x days" is selected.

c. Set the start time.

   This is the time that the job starts at the specified interval. It uses the time zone configured for your USM Anywhere instance (the default is Coordinated Universal Time [UTC]).

10. Click **Save**.

Collect Logs from Azure Event Hubs

Role Availability ✗ Read-Only ✗ Investigator ✗ Analyst ✓ **Manager**

Microsoft Azure Event Hubs is a data and event processing service for Microsoft Azure. The integration between USM Anywhere and Azure Event Hubs enables the Azure Sensor to receive and process information from an event hub so that you can manage it in your USM Anywhere environment.

Warning: To process and display the custom events received from the Azure Event Hubs as generic events, USM Anywhere needs these custom events in a specific format. The correct format is an array as a value of a "records" key in JSON format. For example `{ "records": [ {<event-content>} ] }`.

Important: Be sure to review the Azure requirements page for any environmental requirements specific to Azure Event Hubs before implementing the streaming of your logs to Azure Event Hubs.

The Azure Sensor can process different types of logs sent through Azure Event Hubs, including but not limited to the following:

- Azure Active Directory (AD) logs, including audit logs and sign-in logs
- Azure Application Gateway logs
- Azure Monitor logs
- Azure SQL Database logs
- Microsoft Defender Advanced Threat Protection (ATP) logs
- Microsoft Intune logs

Important: The Azure Sensor will need to be connected to ports 5671 and 5672 in order to integrate with Azure Event Hubs.

Stream Logs to Azure Event Hubs

Before configuring the Azure Event Hubs integration in USM Anywhere, you must stream the logs you want to be analyzed to Azure Event Hubs. Make sure to stream your logs to the same event hub, because each Azure Sensor can only collect from a single event hub.

To stream logs to Azure Event Hubs

1. Log in to the Azure portal.
2. Create an event hub. See Microsoft Azure Quickstart: Create an event hub using Azure portal for instructions.
3. Go to the event hub you just created and click **Shared Access Policies** in the sidebar.
4. Create or edit a policy, and then select **Manage**, **Send**, and **Listen**. Streaming to Event Hubs requires these permissions.
5. Copy the connection string listed in the policy under *Connection String–Primary Key*.

   Note: You will need to enter this string when configuring the Event Hubs connection in USM Anywhere.

6. Configure streaming for the logs you want to collect. For example:

> **Note:** Make sure to enable *Stream to an event hub* and select the Event Hub you just created as the destination.

- **Azure AD logs:** See Stream Azure Active Directory Logs to an Azure Event Hub for instructions from Microsoft.
- **Azure Application Gateway logs:** See Enable Logging for Application Gateway for instructions from Microsoft.
- **Azure Monitor logs:** See Create Diagnostic Settings to Send Logs for instructions from Microsoft.
- **Azure SQL Database logs:** See Set up auditing for your database for instructions from Microsoft. Make sure to select Event Hub as the destination.
- **Microsoft Defender ATP logs:** See Configure Microsoft Defender ATP to stream Advanced Hunting events to your Azure Event Hubs for instructions from Microsoft.
- **Microsoft Intune logs:** See Send log data to storage, event hubs, or log analytics in Intune for instructions from Microsoft.

Set Up Azure Event Hubs Connection in USM Anywhere

After completing the initial setup of your Azure Event Hubs, return to your USM Anywhere Sensors page to enable the Azure Event Hubs connection in USM Anywhere.

To enable Azure Event Hubs in USM Anywhere

1. Go to **Data Sources > Sensors**, and then open the Azure Sensor.
2. Click the **Configuration** tab.

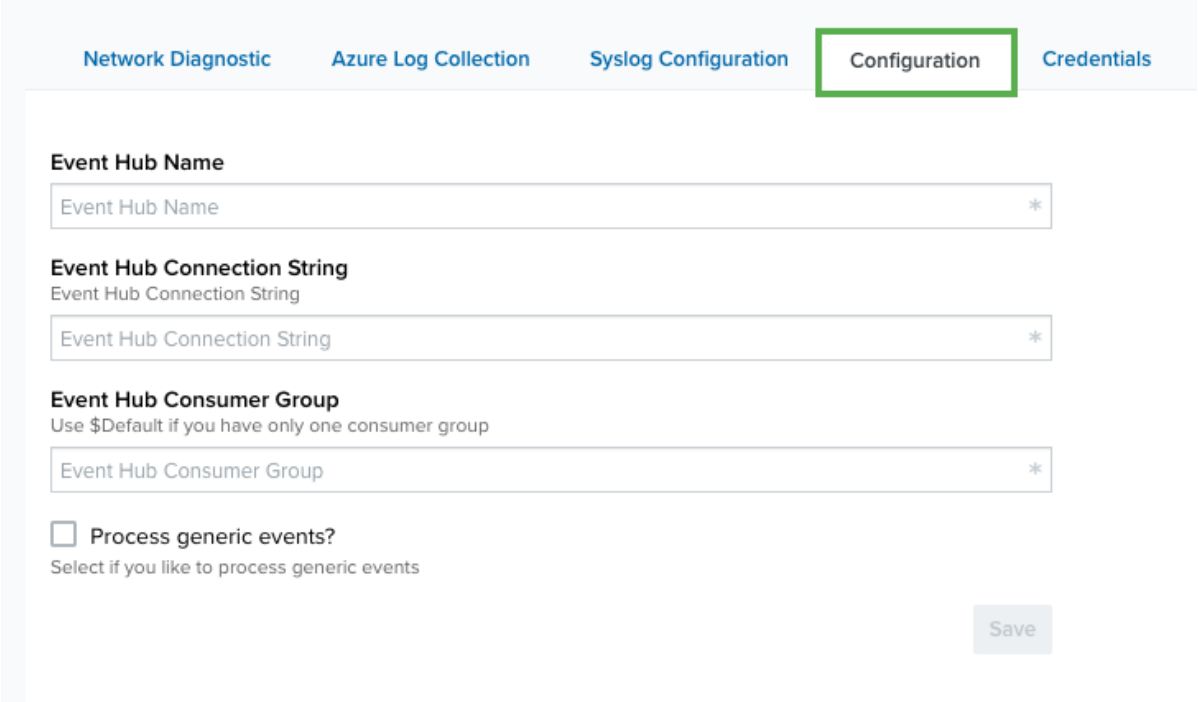

3. Complete the three fields:

- **Event Hub Name:** The name of the event hub created during initial setup.

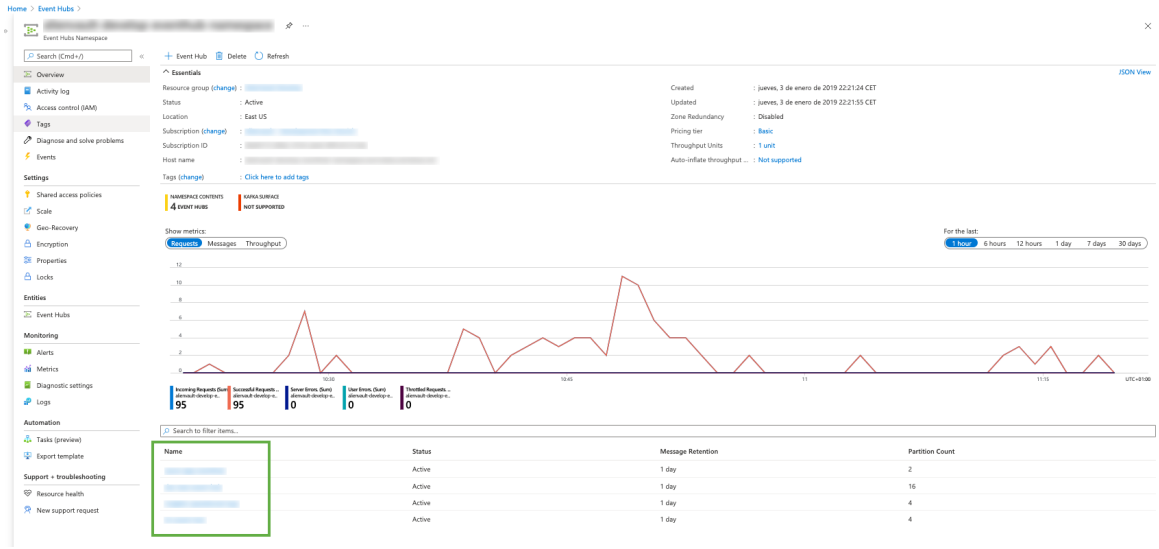

- **Event Hub Connection String:** A string containing unique configuration data about your Azure Event Hubs implementation. This is the connection string that was copied under *Connection String–Primary Key* in the Stream Logs to Azure Event Hubs procedure.
- **Event Hub Consumer Group:** The name of your Event Hubs consumer group. You can locate this name by opening your Event Hubs overview in the Azure portal and scrolling to the bottom of the page.

4. (Optional.) Select **Process Generic Events** to collect events for which USM Anywhere currently does not have a parser. These events will display as "GENERIC event" under Activity > Events.
5. Click **Save**.
6. Click the **Event Hub** tab to check the connection status and the number of events processed by each data source.

Viewing Azure Event Hubs Connectivity in USM Anywhere

The Event Hub tab on the Azure Sensor page provides a glimpse into the health of your sensor's connection to Azure Event Hubs. This page contains the name of your event hub, its connectivity status, and the number of events being processed by USM Anywhere.

To view your Azure Event Hubs connection

1. Go to the Sensors page, and then open your Azure Sensor.
2. Click the **Event Hub** tab.

These are the connectivity statuses you may see:

- **Connecting:** Azure Event Hubs is currently connecting to the sensor.
- **Processing:** Azure Event Hubs is successfully connected.
- **Shutting Down:** Azure Event Hubs has begun the shutdown process to allow a different event hub to connect to the sensor.
- **Shutdown:** The sensor is not currently connected to an event hub.
- **Error:** The connection has experienced an error.

USM Anywhere Sensor Deployment on GCP

The USM Anywhere Sensor provides operational visibility into the security of your Google Cloud Platform (GCP) environment. Based on the collected log information, USM Anywhere analyzes the data generated by your GCP environment and provides real-time alerting to identify malicious activity. The sensor is deployed into your GCP environment to provide ultimate control over the installation and the data contained within it, while avoiding any external access to your environment.

All USM Anywhere Sensors allow for authenticated scans of assets by leveraging stored credentials that you define in USM Anywhere. This enables USM Anywhere to detect potential vulnerabilities, installed software packages, and running processes and services.

This section includes the following topics:

About GCP Sensor Deployment 250

Requirements for GCP Sensor Deployment 251

Preparing Your GCP Environment for Sensor Deployment 261

Deploy the GCP Sensor 267

Connect the GCP Sensor to USM Anywhere 269

Complete the GCP Sensor Setup 274

GCP Log Discovery and Collection in USM Anywhere 287

About GCP Sensor Deployment

The USM Anywhere Sensor provides operational visibility into the security of your Google Cloud Platform (GCP) environment. Based on the collected log information, USM Anywhere analyzes the data generated by your GCP environment and provides real-time alerting to identify malicious activity. The sensor is deployed into your GCP environment to provide ultimate control over the installation and the data contained within it, while avoiding any external access to your environment.

All USM Anywhere Sensors allow for authenticated scans of assets by leveraging stored credentials that you define in USM Anywhere. This enables USM Anywhere to detect potential vulnerabilities, installed software packages, and running processes and services.

The GCP Sensor does not require you to install a sensor for every GCP project you wish to monitor. If you have multiple projects under a single GCP organization, the sensor can be configured to handle multiple projects within that organization.

Log Collection and Scans

The GCP Sensor collects GCP and system log, and generates asset scans and vulnerability assessments, consisting of the following:

- Google Cloud Audit Logs
- Amazon Virtual Private Cloud (VPC) Flow Logs
- Firewall logs
- Syslogs
- Apache Logs
- NGINX logs
- Operational logs for critical software packages deployed, such as HTTP servers and database servers
- Asset scans on your virtual machines (VMs) to inventory installed software packages, running processes, and services
- Periodic vulnerability assessments

Log Analysis

USM Anywhere analyzes these logs in these stages:

Stage 1: Collects logs from systems and software running in your environment

Stage 2: Configures log line processing and generates events

- Includes IP addresses and timestamps culled from extracted log-line data
- Adds other data to the event, such as security context and environmental information

Stage 3: Analyzes events and stores them

Deployment Overview

AT&T Cybersecurity distributes the GCP Sensor as a Google Cloud Deployment Manager template specifically for the Google Virtual Private Cloud (VPC).

The deployment process for an initial USM Anywhere Sensor in your GCP environment consists of these primary tasks:

1. Review requirements for a GCP Sensor deployment.
2. Prepare your GCP environment for sensor deployment.
3. Deploy the USM Anywhere Sensor within your GCP environment.
4. Register the sensor with your sensor authentication code to provision the USM Anywhere instance and connect the deployed sensor.
5. Complete your GCP Sensor configuration, including initial asset discovery.
6. Configure log collection with Google Cloud Pub/Sub.

Requirements for GCP Sensor Deployment

USM Anywhere deploys the Google Cloud Platform (GCP) Sensor through the Google Virtual Private Cloud (VPC) and requires at a minimum the following specifications.

VPC Minimum Specifications

Requirement	Description
n1-standard-2 instance	A standard instance with 2 vCPUs and 7.5 GB of memory.
Zonal SSD persistent disks	Persistent disk storage offers reliable network storage that your instances can access like physical disks. For optimal performance, 50GB and 128GB volumes are designated as the default size for the root and data partitions.
Internet connection to the AT&T Cybersecurity Secure Cloud	See Sensor Ports and Connectivity for more information.

Note: USM Anywhere does not support deploying a GCP Sensor on a shared VPC network.

Important: Because the needs of a sensor differ based on the varying demands of different deployment environments and the complexity of events being processed, the number of events per second (EPS) a sensor can process varies.

Depending on your environment, you may need to deploy additional sensors to ensure that all events are processed.

Application Service Dependencies

With the Google Cloud Deployment Manager Template provided by AT&T Cybersecurity, you can automatically deploy a USM Anywhere Sensor as a service into your environment. Review the following lists for information about the inbound and outbound IP addresses, ports, and services used by USM Anywhere.

Sensor Ports and Connectivity

For USM Anywhere Sensor deployment in the Google Cloud Platform (GCP), the Google Cloud Deployment Manager template automatically creates the firewall rules needed for network connectivity between the instances within the virtual private cloud (VPC).

Note: The required firewall rules are outlined below.

The following tables list the inbound and outbound ports.

Sensor Ports and Connectivity (Outbound Ports)

Type	Ports	Endpoints	Purpose
TCP	443	update.alienvault.cloud	Communication with AT&T Cybersecurity for initial setup and future updates of the sensor.
TCP	443	reputation.alienvault.com	Ongoing communication with AT&T Alien Labs™ Open Threat Exchange® (OTX™).
TCP	443	otx.alienvault.com	Ongoing communication with OTX to retrieve vulnerability scores. Connecting to otx.alienvault.com is not required but highly recommended. OTX uses the AWS CloudFront services. Refer to the AWS IP address ranges page when you deploy a new sensor. This page contains the current IP address ranges for the service and instructions on how to filter the addresses.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
SSL	443	storage-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send and retrieve backups.
SSL	443	metrics-proxy.services.prod.alienvault.cloud	Ongoing communication with USM Anywhere in order to send metrics and messages.
SSL/TCP	443	api-parameters-<REGION>-prod.alienvault.cloud[5] api-message-proxy-<REGION>-prod.alienvault.cloud api.message-proxy.<REGION>.prod.alienvault.cloud	Ongoing communication with USM Anywhere. It is only necessary to allowlist the address that corresponds to the region where your USM Anywhere instance is hosted.
SSL/TCP	7100	Your USM Anywhere subdomain .alienvault.cloud Your USM Anywhere subdomain .gov.alienvault.us (for AT&T TDR for Gov)	Ongoing communication with USM Anywhere.
UDP	53	DNS Servers (Google Default)	Ongoing communication with USM Anywhere.
UDP	123	metadata.google.internal	Sync with NTP services.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	22 and 443	prod-usm-saas-tractorbeam.alienvault.cloud prod-gov-usm-saas-tractorbeam.gov.alienvault.us (for AT&T TDR for Gov)	SSH communications with the USM Anywhere remote support server. See Troubleshooting and Remote Sensor Support for more information about remote technical support through the USM Anywhere Sensor console.

Sensor Ports and Connectivity (Outbound Ports) (Continued)

Type	Ports	Endpoints	Purpose
TCP	443	geoip-ap-northeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-south-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-1-prod.alienvault.cloud/geo-ip/sensor geoip-ap-southeast-2-prod.alienvault.cloud/geo-ip/sensor geoip-ca-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-1-prod.alienvault.cloud/geo-ip/sensor geoip-eu-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-me-central-1-prod.alienvault.cloud/geo-ip/sensor geoip-sa-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-east-1-prod.alienvault.cloud/geo-ip/sensor geoip-us-west-2-prod.alienvault.cloud/geo-ip/sensor geoip-us-gov-west-1-prod-gov.gov.alienvault.us/geo-ip/sensor (for AT&T TDR for Gov)	Allows resolution of IP addresses for geolocation services. It is only necessary to allowlist the GeoIP address that corresponds to the region where your USMA instance is hosted.

1

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

2

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

3

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

4

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

5

<REGION> corresponds to the AWS endpoint region based on your location (for example, api-parameters-ap-south-1-prod.alienvault.cloud).

Sensor Ports and Connectivity (Inbound Ports)

Type	Ports	Purpose
SSH	22	Inbound method for secure remote login from a computer to USM Anywhere.
HTTP	80	Inbound communication for HTTP traffic.
UDP (RFC 3164)	514	USM Anywhere collects data through syslog over UDP on port 514 by default.
TCP (RFC 3164)	601	Inbound communication for reliable syslog service. USM Anywhere collects data through syslog over TCP on port 601 by default.
TCP (RFC 5424)	602	USM Anywhere collects data through syslog over TCP on port 602 by default.
Traffic Mirroring	4789	Inbound communication for virtual extensible local area network (VXLAN).

Sensor Ports and Connectivity (Inbound Ports) (Continued)

Type	Ports	Purpose
WSMANS	5987	Inbound WBEM WS-Management HTTP over Secure Sockets Layer/Transport Layer Security (SSL/TLS) (NXLog).
TLS/TCP (RFC 3164)	6514	USM Anywhere collects TLS-encrypted data through syslog over TCP on port 6514 by default.
TLS (RFC 5424)	6515	USM Anywhere collects data through syslog over TLS on port 6515 by default.
TCP	9000	Inbound communication used internally for HTTP sensor traffic.
Graylog	12201	Inbound communication for Graylog Extended Log Format (GELF).

Google Cloud Services

The USM Anywhere Sensor uses the following GCP services:

- Google Stackdriver
- Google Compute Engine
- Google Cloud Pub/Sub

USM Anywhere IP Addresses for Allowlisting

Your sensor is connected to a USM Anywhere instance deployed in one of the Amazon Web Services (AWS) endpoint regions based on your location. If you need to configure your firewall to allow communication between the sensor and the USM Anywhere instance, refer to the following table with the reserved IP address ranges for each region.

Important: The Update Server and the AlienVault Agent always use the 3.235.189.112/28 range no matter where your USM Anywhere is deployed. The AT&T TDR for Gov Update Server uses the 3.32.190.224/28 range.

Note: The regional IP ranges listed in this table are limited to the control nodes (subdomain). You must also meet all requirements provided in the Sensor Ports and Connectivity (Outbound Ports) table.

AWS Regions Where USM Anywhere Instance Is Available

Code	Name	Reserved Static IP Address Ranges
ap-northeast-1	Asia Pacific (Tokyo)	18.177.156.144/28 3.235.189.112/28 44.210.246.48/28
ap-south-1	Asia Pacific (Mumbai)	3.7.161.32/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-1	Asia Pacific (Singapore)	18.143.203.80/28 3.235.189.112/28 44.210.246.48/28
ap-southeast-2	Asia Pacific (Sydney)	3.25.47.48/28 3.235.189.112/28 44.210.246.48/28
ca-central-1	Canada (Central)	3.96.2.80/28 3.235.189.112/28 44.210.246.48/28
eu-central-1	Europe (Frankfurt)	18.156.18.32/28 3.235.189.112/28 44.210.246.48/28
eu-west-1	Europe (Ireland)	3.250.207.0/28 3.235.189.112/28 44.210.246.48/28

AWS Regions Where USM Anywhere Instance Is Available (Continued)

Code	Name	Reserved Static IP Address Ranges
eu-west-2	Europe (London)	18.130.91.160/28 3.235.189.112/28 44.210.246.48/28
me-central-1	Middle East (UAE)	3.29.147.0/28 3.235.189.112/28 44.210.246.48/28
sa-east-1	South America (São Paulo)	18.230.160.128/28 3.235.189.112/28 44.210.246.48/28
us-east-1	US East (N. Virginia)	3.235.189.112/28 44.210.246.48/28
us-west-2	US West (Oregon)	44.234.73.192/28 3.235.189.112/28 44.210.246.48/28
us-gov-west-1	AWS GovCloud (US-West)	3.32.190.224/28

Installation Prerequisites

Before you install the GCP Sensor, make sure you have the prerequisites available in the following table.

Installation Prerequisites

Prerequisites	Description
Cloud Deployment Manager template provided by AT&T Cybersecurity	The Cloud Deployment Manager template automatically creates all required GCP resources for deployment, including an instance, volume, and firewall rules for use by the USM Anywhere Sensor instance.
Privileged user account on GCP	To deploy the Cloud Deployment Manager template, you must have a privileged user account in GCP with Compute Engine permissions, permission to create and edit service accounts, as well as write permission to the Cloud Deployment Manager.

Preparing Your GCP Environment for Sensor Deployment

Role Availability ✗ Read-Only ✗ Investigator ✗ Analyst ✓ **Manager**

After you have ensured that your Google Cloud Platform (GCP) environment meets the sensor requirements, you must complete both of the following tasks before deploying a GCP Sensor in your environment:

- Enable required APIs
- Create a new service account

Enable Required APIs

Certain APIs must be enabled in your GCP environment to enable the features dependent on them to operate as designed.

Important: APIs are enabled at the project level, so you must enable all five of these APIs for each project the GCP Sensor will monitor.

The following APIs are needed in your GCP environment:

- **Google Cloud Resource Manager API:** https://console.cloud.google.com/apis/library/cloudresourcemanager.googleapis.com
- **Google Cloud Pub/Sub Logging API:** https://console.cloud.google.com/apis/api/pubsub.googleapis.com
- **Google Stackdriver Logging API:** https://console.developers.google.com/apis/api/logging.googleapis.com
- **Google Compute Engine API:** https://console.cloud.google.com/apis/library/compute.googleapis.com
- **Google Cloud Identity and Access Management (IAM) API:** https://-console.developers.google.com/apis/api/iam.googleapis.com/overview

To enable an API in your GCP environment

1. Log in to your GCP environment.
2. Navigate to that API in the GCP API library (or follow the corresponding link in the list above).

   Note: If the API is already enabled, you may see a green check mark and the text "API enabled" instead of the Enable button. In some views, you will see a "Disable API" button to indicate that the API has already been enabled.

3. Click **Enable**.

   If the Enable button is grayed out, ensure that you have the appropriate permissions required to manage APIs.

Create a New Service Account

The service account you have selected for your GCP Sensor must have adequate permissions for every GCP project it will monitor. Without these permissions, the sensor will not be able to accomplish the task that requires that access.

To create a new service account

1. In the Cloud Console, go to your project.
2. Go to the *IAM & admin* tab in the navigation pane and click **Service Accounts**.
3. Click **Create Service Account** and enter the required information for your new service account.
   a. **Service Account Name:** A display name for this service account
   b. **Service Account ID:** A name for your service account, which will be followed by "@<name-of-project>.iam.gserviceaccount.com"
   c. **Service Account Description:** A description for this service account
4. Click **Create and Continue** to save your new service account.
   From here, if you are facing a screen that allows you to grant the service account access to the project, or users access to the service account, you can click **Done** without making any changes on that screen to skip that step and move forward.

Generally, you will use the pre-defined roles *Project: Viewer* and *Pub/Sub: Pub/Sub Subscriber* for your service account. The *Project: Viewer* role allows your sensor to discover all your services, and the *Pub/Sub: Pub/Sub Subscriber* role allows your sensor to collect logs from Cloud Pub/Sub.

To assign the pre-defined roles to your service account

Important: This process must be followed for every project the GCP Sensor will be monitoring.

1. In the Cloud Console, go to your project.
2. Go to the *IAM & admin* tab in the navigation pane and click **IAM**.
3. Click **Grant Access**.
4. Enter the name of the service account you just created.
5. In the Role field, select **Project** and then **Viewer**.
6. Open a second Role field, this time selecting **Pub/Sub** and then **Pub/Sub Subscriber**.
7. Open a third Role field, this time selecting **Deployment Manager** and then **Deployment Manager Editor**.
8. (Optional.) Open a fourth Role field, this time selecting **Service Accounts** and then **List**.
9. Click **Save** once you are finished assigning roles.

Note: This role is only required if you intend to enable User Behavior Analytics (UBA).

If these roles are too expansive for your use, you can create a new role and limit its access according to your needs, so long as it has the minimum requirements necessary for the sensor to operate. See Creating a Custom Role for instructions detailing how to create a custom role for your sensor. Also be sure to review the Required IAM Policies table to see which functions depend on which IAM policies.

To create and download a new service account key

1. On the **Service Accounts** page, click the email address of the service account you just created and navigate to the **Keys** tab.
2. Using the **Add Key** drop-down, select **Create New Key**.
3. Select **JSON** for the key type and click **Create**.
   Clicking Create downloads a service account key file.
4. Save this key file in a safe location.
   You will need to reference this file when you Deploy the GCP Sensor.

Create and Add an SSH Key

You will need to create an SSH key and add it to your GCP project. This SSH key will be used to connect to your sensor once it is deployed.

To create and add an SSH key

1. Follow the steps outlined in the Google Cloud documentation (appropriate to your OS) to create an SSH key.
   You will save a copy of a newly generated SSH key and use it later in this process.
2. Within the Google Cloud console, navigate to your project.
3. Search for and select **SSH Keys**.
4. Click **Edit**, then **Add Item**.
5. Enter the key you copied earlier.
   This is the .pub file that was generated in step 1.
6. Click **Save**.

Creating a Custom Role

Role Availability	Read-Only	Investigator	Analyst	Manager

If the pre-defined roles Project: Viewer and Pub/Sub: Pub/Sub Subscriber are too broad for your use, or are otherwise unsuitable for you, you can define a new role whose access is limited according to your needs.

Warning: At minimum, your service account role must be assigned each of the IAM policies required for your sensor operations. Review the Required IAM Policies table to see which functions depend on which IAM policies.

These permissions can be granted at the organization level; however, if your organization is very large you may experience performance issues. In this case (as long as you don't need the sensor to monitor all projects), you can use either of the following approaches to avoid possible throttling:

Project-Level Permissions

This allows you to select which specific projects should be monitored by the sensor. This approach is not valid for any logging at the organization level, or any functionality dependent on organization level permissions will not be enabled.

To grant the service account permission to monitor a project

Important: This process must be followed for every project the GCP Sensor will be monitoring.

1. In the Google Cloud Console, go to your project.
2. Go to the IAM & admin tab in the navigation pane and click **IAM**.
3. Click **Add**.
4. Enter the name of the service account whose permissions you are editing.

Note: The name of the service account takes the form of an email address and will look like <name-of-sensor-service-account>@<name-of-project>.iam.gserviceaccount.com.

5. In the Role field, select the appropriate role for this service account.
6. Click **Save**.

Note: To grant the service account permission to monitor the entire organization, use these same steps but begin by opening the organization instead of the project.

Required IAM Policies

At the organization level, the GCP Sensor needs the specific IAM policies in the following table.

Required IAM Policies at the Organization Level

IAM Policy	Description	Dependency
logging.logEntries.list	Allows the sensor to fetch log entries from Stackdriver	Google Cloud Audit Logs for Organizations
resourcemanager.organizations.get	Allows the sensor to get the details for a specific organization	Application Status Cloud Audit Logs for organizations

At the project level, the GCP Sensor needs the specific IAM policies in the following table.

Required IAM Policies at the Project Level

IAM Policy	Description	Dependency
logging.logEntries.list	Allows the sensor to fetch log entries from Stackdriver	Cloud Audit Logs for Projects Firewall Logs for Projects VPC Flow Logs for Projects Stackdriver Agent Logs
resourcemanager.projects.list	Allows the sensor to access a list of the available projects	Application Status Asset Inventory Configuration Issues Cloud Audit Logs for Projects Firewall Logs for Projects VPC Flow Logs for Projects Stackdriver Agent Logs
resourcemanager.projects.get	Allows the sensor to fetch the details for a specific project	

Required IAM Policies at the Project Level (Continued)

IAM Policy	Description	Dependency
deploymentmanager.deployments.create	Allows the sensor to be created and deployed	Deployment of a sensor
compute.firewalls.list	Allows the sensor to list the existing firewall rules	Configuration Issues
compute.firewalls.get	Allows the sensor to get the details for a specific firewall rule	Configuration Issues
compute.instances.list	Allows the sensor to list the existing virtual machines	Asset Inventory Configuration Issues
compute.instances.get	Allows the sensor to get the details for a specific virtual machine	Asset Inventory Configuration Issues
compute.zones.list	Allows the sensor to list the available zones	Asset Inventory Configuration Issues

Deploy the GCP Sensor

After you review the requirements and make sure that your Google Cloud environment is configured as needed, you can deploy the Google Cloud Platform (GCP) Sensor. This sensor must be deployed using the gcloud command-line interface (CLI).

Important: You must download and install the Google Cloud Software Development Kit (SDK) on your system and initialize it before you can use the gcloud CLI. See the Google Cloud SDK documentation for instructions on how to install and initialize the SDK.

The following procedure describes how to launch the GCP Sensor when provisioning the USM Anywhere service for the first time. In this process, you launch the USM Anywhere product using Google Cloud commands from your preferred command line interface.

To create a new sensor using gcloud CLI commands

1. Go to the USM Anywhere Sensor Downloads page and click the icon of your specific sensor. After clicking, your browser starts to download the USM Anywhere Sensor package.

2. Use the following command to log into GCP using the service account you created in Preparing Your GCP Environment for Sensor Deployment, replacing the variables below with the information relating to the service account key you downloaded:

   - **path_to_sa_file:** The path to your Google service account key
   - **service_account_key:** The name of your Google service account key

   ```
   gcloud auth activate-service-account --key-file <path_to_sa_key>/<service_account_key>
   ```

3. Navigate to the location where you saved the zip file and unzip it.

4. Define the required properties, replacing the variables below with your information:

   - **service_account_id:** Google service account ID
     This ID is in the form of an email address.
   - **public_key:** The full contents of the public SSH key downloaded from Google in Create and Add an SSH Key
   - **network_id:** The name of your network
     You can find this network name in your GCP Console by going to **VPC network > VPC networks** and copying the name of the network.

   Linux/Mac command

   ```
   PROPS="service_account:<service_account_id>,ssh_key:<public_key>,network:<network_id>,public_ip:True"
   ```

   Windows command

   ```
   set PROPS="service_account:<service_account_id>,ssh_key:<public_key>,network:<network_id>,public_ip:True"
   ```

   You can also include the following optional parameters in this command:

   - **public_ip:** "True"
     By default, your sensor is deployed to a private IP address. Setting this value to "True" will deploy to a public IP address.
   - **ip_ranges:** Specify to which range of IP addresses your firewall rules apply
     By default, the sensor will allow traffic from all IP addresses (0.0.0.0/0).

5. Use the following command to deploy the sensor, replacing the variables below with your information:

   - **VM_name:** The name of your virtual machine (VM)

   > **Warning:** This name must not be used by another VM in your environment, or your deployment will fail.

   - **project_id:** The project ID of your GCP project
     You can find this project ID from anywhere in your GCP Console by clicking the drop-down in the upper left of the screen and copying the project ID displayed in the window that opens.

   Linux/Mac command

   ```
   gcloud deployment-manager deployments create "<VM_name>" --template
   "./usm-anywhere-sensor-gcp.template.py" --properties "${PROPS}" --project
   "<project_id>"
   ```

   Windows command

   ```
   gcloud deployment-manager deployments create "<VM_name>" --template
   ".\usm-anywhere-sensor-gcp.template.py" --properties %PROPS% --project
   "<project_id>"
   ```

   Your sensor is now deployed.

6. After the deployment has finished, locate the sensor's IP address by reviewing the output of the previous command. You will find the URL under `OUTPUTS VALUE`.

   ```
   The fingerprint of the deployment is b'CWA2KOQCDI7zYAWMRTAriQ=='
   Waiting for create [operation-1624951011359-5c5e263cf3c43-333918e7-
   9c21733e]...done.
   Create operation operation-1624951011359-5c5e263cf3c43-333918e7-9c21733e
   completed successfully.
   OUTPUTS VALUE
   URL http://<sensor_ip_address>/
   CLIUser sysadmin
   ```

   > **Note:** Make note of this IP address so that you have it for configuring your data sources to send data to the GCP Sensor.

7. Paste the IP address in your browser to launch the USM Anywhere Sensor Setup page.

Connect the GCP Sensor to USM Anywhere

After deploying the GCP Sensor, you must connect it to USM Anywhere through registration.

Obtain the Authentication Code

You must enter an authentication code when registering the USM Anywhere Sensor. How to obtain the authentication code depends on your USM Anywhere instance and whether this is the first sensor you're deploying.

Instructions for USM Anywhere customers:

If this is your first USM Anywhere Sensor, you must register the sensor using the initial authentication code (starts with a "C") received from AT&T Cybersecurity. With this code, the registration process provisions a new USM Anywhere instance and defines its attributes, such as how many sensors to allow for connection, how much storage to provide, and what email address to use for the initial user account. After registration, you will gain access to the sensor through the USM Anywhere web user interface (UI), where you can complete the sensor setup.

If you are deploying additional sensors, you must generate the authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Instructions for AT&T TDR for Gov customers:

AT&T Cybersecurity has already provisioned the AT&T Threat Detection and Response for Government (AT&T TDR for Gov) instance for you, therefore you won't receive an authentication code for your sensor. This is true regardless if it's the first sensor or additional sensors you're deploying. However, for the first sensor, you'll receive a link to access your instance.

For every sensor you deploy, you must generate an authentication code (starts with an "S") for the registration. See Adding a New Sensor in the *USM Anywhere User Guide* for more information.

Register Your Sensor

You perform this procedure after deploying the USM Anywhere Sensor within your Google Cloud Platform (GCP) account.

To register your sensor

1. Click the public IP address displayed for the running sensor VM in the GCP environment.

   This opens the *Welcome to USM Anywhere Sensor Setup* page, which prompts you to provide the information for registering the sensor with your new USM Anywhere instance.

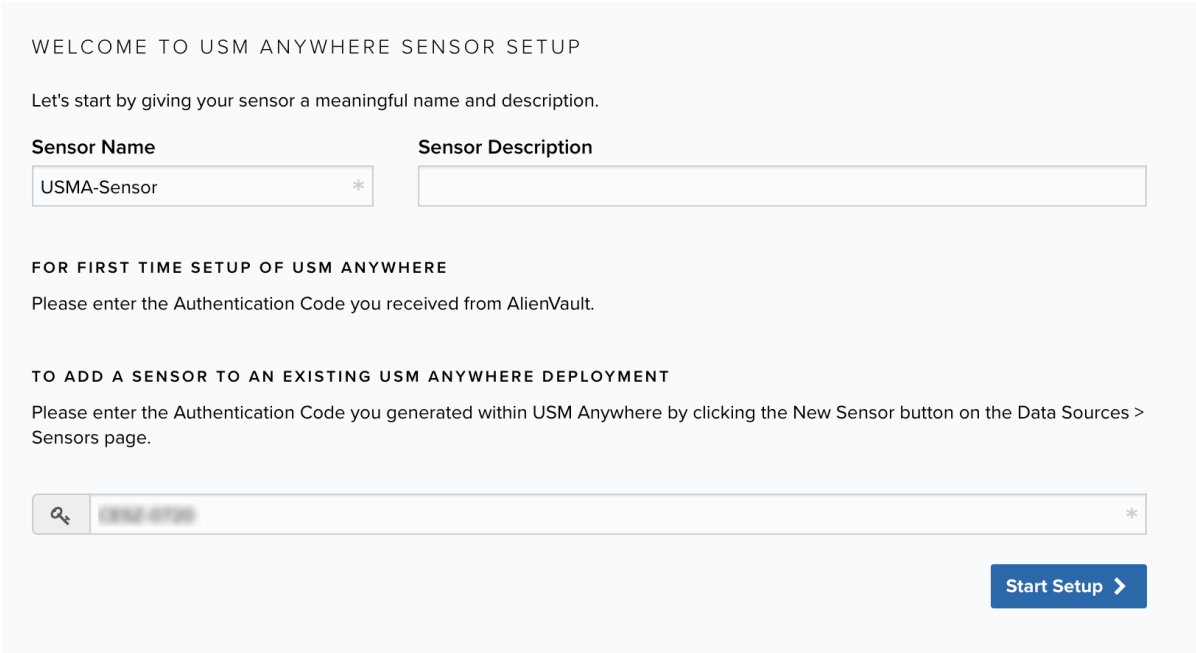

2. Enter a sensor name and sensor description.
3. Paste the authentication code into the field with the key icon ( ).
4. Click **Start Setup** to start the process of connecting the USM Anywhere Sensor.

   It takes about 20 minutes to provision your USM Anywhere instance upon registration of your initial sensor. When this instance is provisioned and running, you'll see a welcome message that provides an access link.

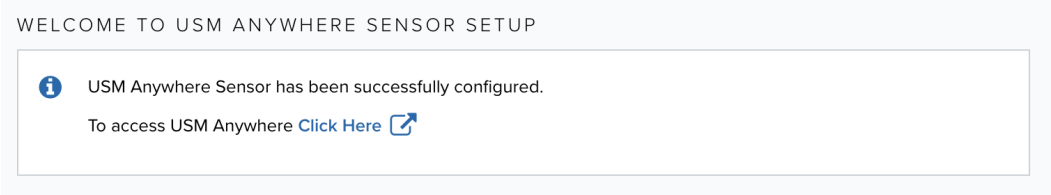

   Use this link to open the secured web console for your USM Anywhere instance. You and the other USM Anywhere users in your organization can access this console from a web

browser on any system with internet connectivity.

> **Note:** If this is your first deployment, you'll also receive an email from AT&T Cybersecurity that provides the access link to USM Anywhere.

Configure the Initial Login Credentials

When you link to a newly provisioned USM Anywhere instance, you must configure the password for the initial user account. This is the default administrator as defined in your subscription.

To configure login credentials

1. In the welcome message, click the link.

   This displays a prompt to set the password to use for the default administrator of USM Anywhere.

2. Enter the password, and then enter it again to confirm.

   Keep in mind these points when you are logging in:

   - The login credentials that you set will apply to any USM Anywhere™ and USM Central™ you have access to.
   - USM Anywhere requires all passwords to have a minimum length of 8 characters and a maximum length of 128 characters.
   - The password must contain numerical digits (0-9).
   - The password must contain uppercase letters (A-Z).
   - The password must contain lowercase letters (a-z).
   - The password must contain special characters, such as hyphen (-) and underscore (_).

> **Note:** USM Anywhere passwords expire after 90 days. When your password expires, USM Anywhere enforces a password change when you next log in. A new password must be different from the previous four passwords.After 45 days of inactivity, your user account will be locked. Manager users can unlock inactive accounts.

3. Click **Save & Continue**.
4. When the login page opens, enter the password you just set and click **Login**.

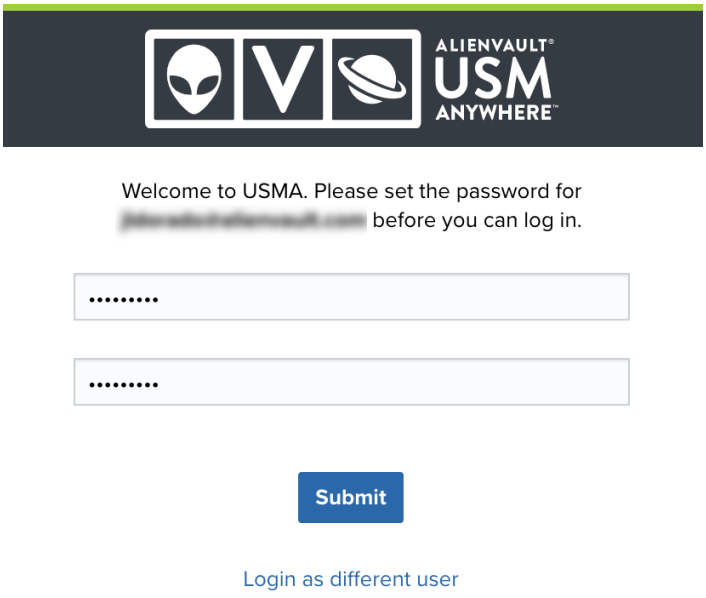

Verify That Your Sensor Is Running

It's a good idea to verify that the USM Anywhere Sensor is running. It also gives you the chance to watch the sensor actively working to find all of your assets and to record events from the start.

> **Note:** Verify that the sensor is running before performing the configuration. You can keep one web browser tab with the Welcome to USM Anywhere page in the background while you perform the verification on a different tab.

To verify that your new sensor is running

1. In USM Anywhere, go to **Data Sources > Sensors**.

   You should now see your sensor in the page. See in the *USM Anywhere User Guide* for more information.

   After a few minutes, USM Anywhere locates your assets and starts generating events.

2. You can review the activity in two locations:
   - From the primary task bar, select **Environment > Assets**.
   - From the primary task bar, select **Activity > Events**.

   > **Note:** It could take up to six minutes before events appear. Make sure to refresh your browser from time to time to display the current data.

   See the *USM Anywhere User Guide* for more information about using the Assets and Events pages in USM Anywhere.

Complete the GCP Sensor Setup

Role Availability	Read-Only	Investigator	Analyst	Manager
	✗	✗	✗	✓

After you initialize a new USM Anywhere Sensor, you must configure it in the Setup Wizard. As you configure the sensor, you can enable USM Anywhere to perform specific actions through scheduled jobs, such as running an asset discovery scan or collecting security events from a predefined cloud storage location.

About Accessing the Setup Wizard

The Setup Wizard is accessible under the following circumstances:

- After you first log in to the USM Anywhere web user interface (UI) and see the Welcome to USM Anywhere page, click **Get Started** to launch the Setup Wizard.
- If you have already registered one USM Anywhere Sensor but did not complete the setup before logging out, the USM Anywhere Sensor Configuration page launches automatically at your next login to remind you to finalize configuration of the sensor. From that page, you click **Configure** to launch the Setup Wizard and complete the sensor configuration.
- If you registered an additional USM Anywhere Sensor, but did not complete the setup, the Sensors page displays an error (✗) in the Configured column. See in the *USM Anywhere User Guide* for more information.

  Go to **Data Sources > Sensors**, and then click the sensor name to complete the sensor configuration. See in the *USM Anywhere User Guide* for more information.

Configuring the Sensor in the Setup Wizard

The first time you log in from the Welcome to USM Anywhere web page, the Setup Wizard prompts you to complete the configuration of the first deployed sensor. Thereafter, you can use the Sensors page to configure an additional sensor or to change the configuration options for a deployed sensor. See in the *USM Anywhere User Guide* for more information.

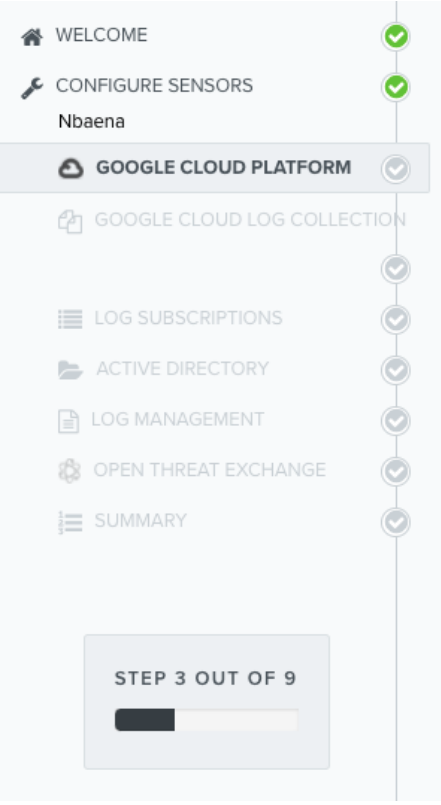

The Google Cloud Platform Configuration page provides information about the asset discovery that occurs upon the initial deployment of the USM Anywhere Sensor, summarizing the number of instances, instance types, and regions in your environment.

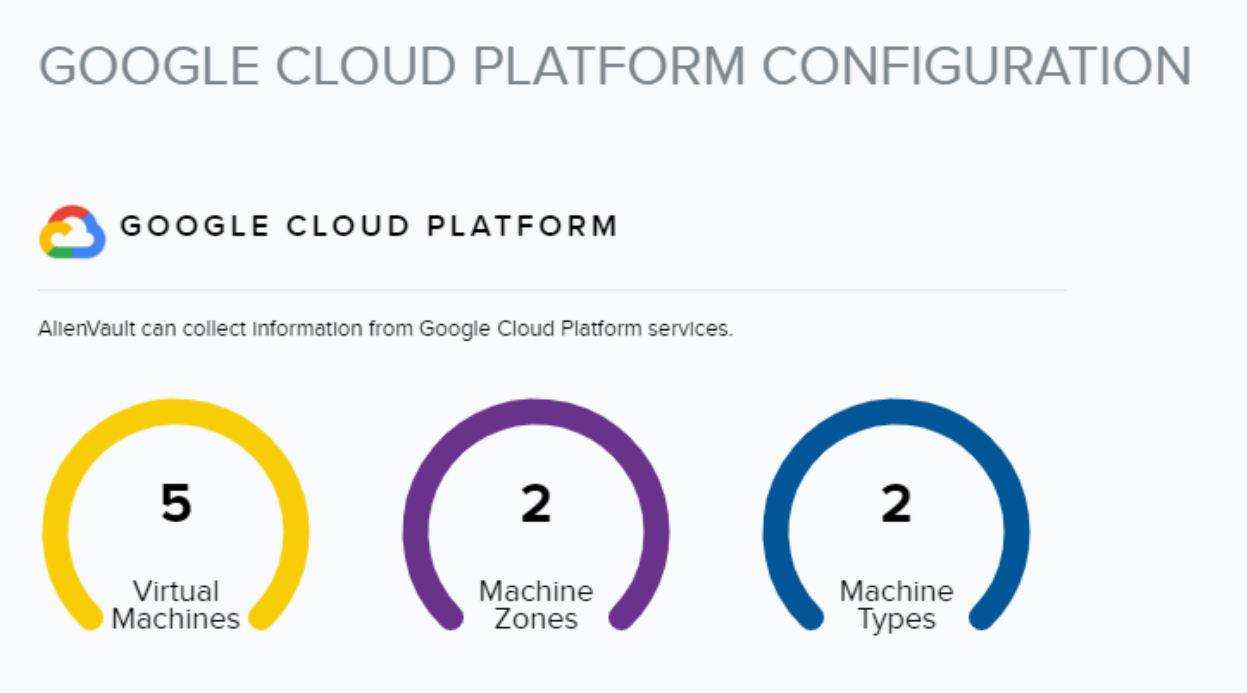

Click **Next** to proceed with the Setup Wizard and complete additional configuration on each page.

GCP Log Collection

USM Anywhere automatically discovers a number of out-of-box logs as long as you have enabled them within your GCP subscription. See GCP Log Discovery and Collection in USM Anywhere for more information about these logs and how they function within the GCP environment.

To enable the out-of-box log collection jobs

1. Locate the job you want to enable and click .

   This turns the icon green ( ).

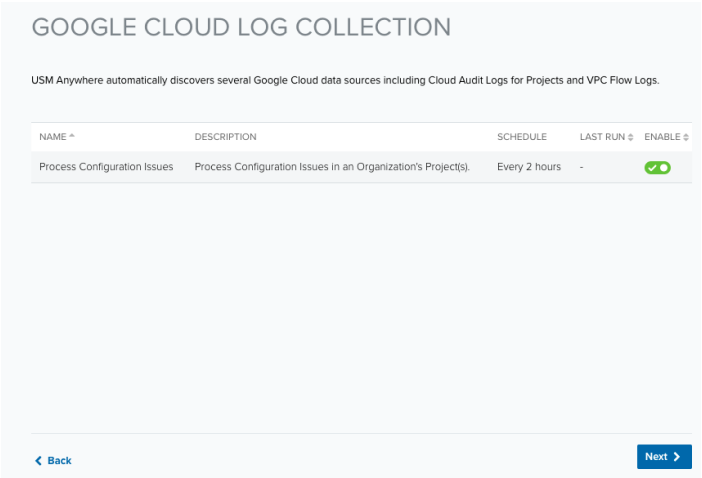

2. Click **Next**.

GCP Log Subscriptions

USM Anywhere uses subscriptions to collect log data from your environment and produce the corresponding events and alarms. See GCP Log Discovery and Collection in USM Anywhere for more information about enabling these subscriptions and how they function within the GCP environment.

To enable the out-of-box log subscriptions

1. Locate the job you want to enable and click **enable**.

   This displays the text as **disable**.

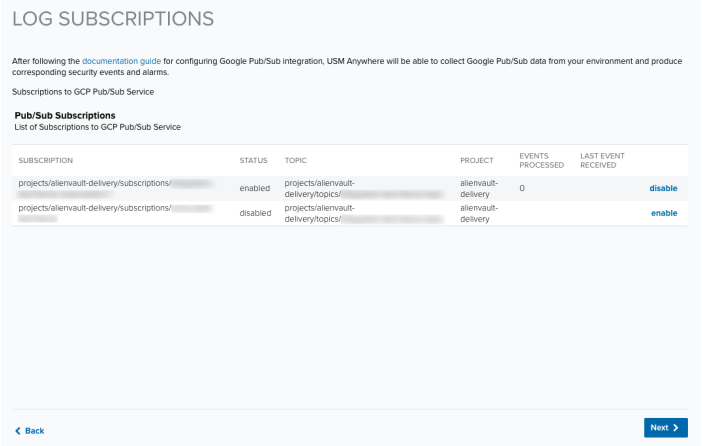

If you don't yet have any subscriptions in your GCP environment, this table will be empty at this step. The steps in GCP Log Discovery and Collection in USM Anywhere will create the subscriptions your sensor needs for Google Cloud Pub/Sub.

> **Note:** Unless you have previously configured your GCP environment for Cloud Pub/Sub integration (as described in GCP Log Discovery and Collection in USM Anywhere), this step will discover any subscriptions that currently exist in your GCP environment. While you may have valid subscriptions already created in your environment for other services, you can *only* use a subscription created with the configuration described here for Cloud Pub/Sub.

Active Directory

The optional Active Directory (AD) setup page configures USM Anywhere to collect information from your AD account. To monitor Microsoft Windows systems effectively, USM Anywhere needs access to the AD server to collect inventory information.

> **Note:** This configuration is only for one AD server. If you want to scan different AD servers, you must create an AD scan job for each of them. See Scheduling Active Directory Scans from the Job Scheduler Page for more information.

AT&T Cybersecurity recommends that you create a dedicated AD account with membership in the Domain Admins group to be used by USM Anywhere to log in to the Windows systems.

You also need to activate Microsoft Windows Remote Management (WinRM) in the domain controller and in all the hosts that you want to scan. You can do this by using a group policy for all the systems in your AD.

Important: Before this feature is fully functional, you must configure access to the USM Anywhere Sensor on the AD server. See Granting Access to Active Directory for USM Anywhere for more information.

To complete the AD access configuration

1. Provide the AD credentials for USM Anywhere:
   - **Active Directory IP Address**: Enter the IP address for the AD server.
   - **Username**: Enter your username as admin of the account.
   - **Password**: Enter your admin's password.
   - **Domain**: Enter the domain for the AD instance.

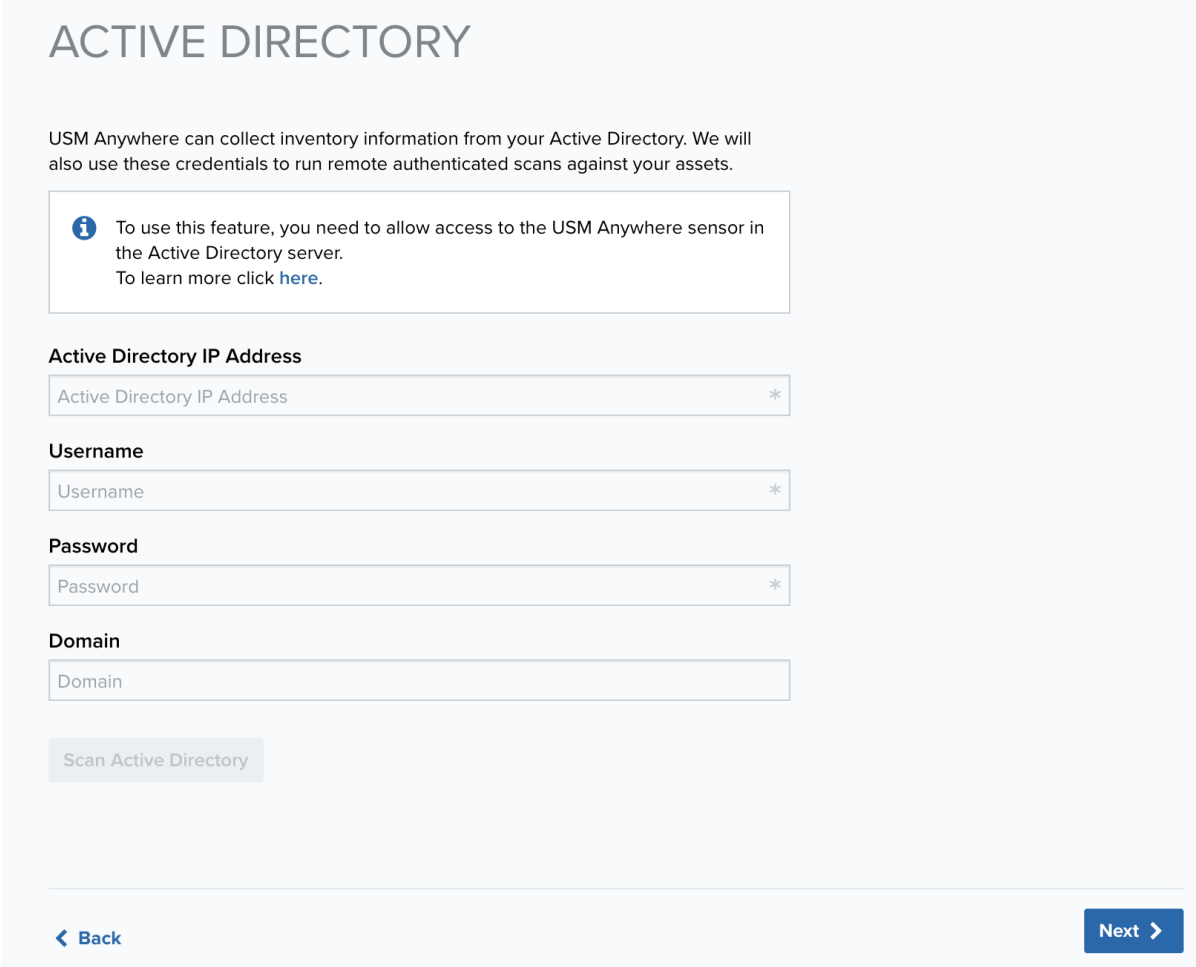

2. Click **Scan Active Directory**.

   After a successful launch of the scan, a confirmation dialog box opens.

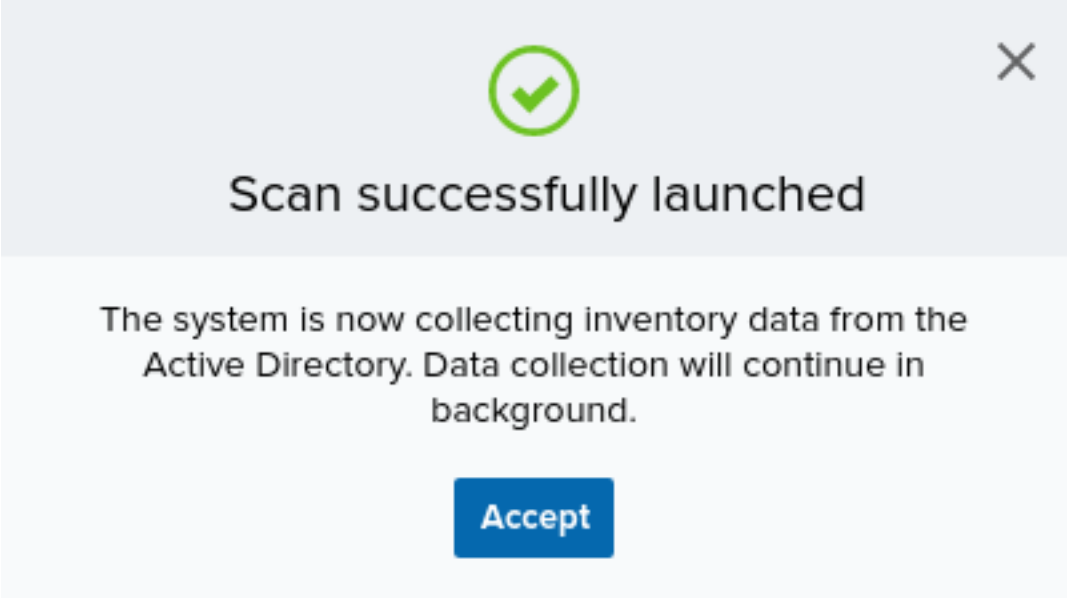

3. Click **Accept**.

   The scan continues in the background.

   Upon completion, another dialog box opens and provides information about the number of assets USM Anywhere discovered. It also prompts you to decide if you want to scan for hosts and services running in your environment.

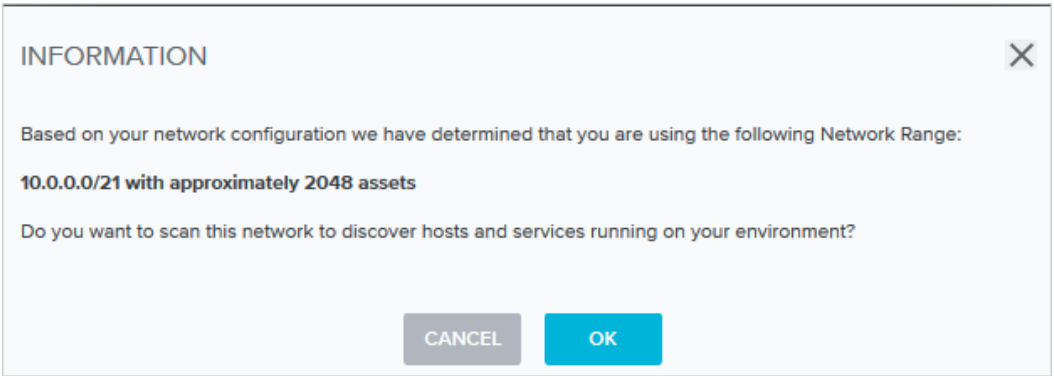

   Click **Cancel** to opt out of this scan.

4. (Optional.) If you want to scan for other hosts and services, click **OK**.

5. Click **Next** after the scan ends.

Log Management

On the Log Management page are syslog port numbers. (These ports are the same for all USM Anywhere Sensors.)

USM Anywhere collects third-party device, system, and application data through syslog over UDP on port 514 and over TCP on ports 601 or 602 by default. It collects Transport Layer Security (TLS)-encrypted data through TCP on ports 6514 or 6515 by default. These ports support the RFC 3164 and RFC 5424 formats. To configure any third-party devices to send data to USM Anywhere, you must provide the IP address and the port number of your USM Anywhere Sensor.

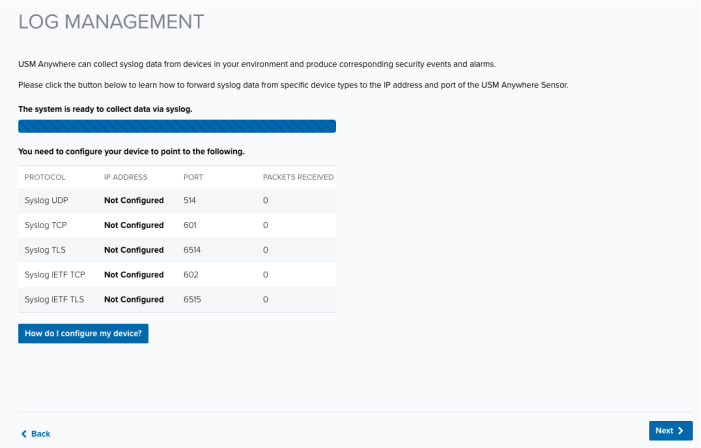

To enable log collection and configure your log management

1. Make sure that you have granted the necessary permissions for your OS to allow USM Anywhere to access its logs. You can also integrate a wide variety of data sources to send log data over syslog to the USM Anywhere Sensor.

   To learn how to configure your operating systems and supported third-party devices to forward syslog log data, see the following related topics:

   - **The Syslog Server Sensor App**: Log collection (UDP, TCP, and TLS-encrypted TCP) from rsyslog
   - **Collecting Linux System Logs**: Log collection from a Linux system
   - Collecting Windows System Logs: Log collection from a Windows system
   - Go to the specific AlienApp in USM Anywhere for instructions about syslog forwarding

> **Note:** Because the log scan can take some time, you might not see all of the automatically discovered log sources immediately after deploying the first USM Anywhere Sensor.

2. When you have finished the log collection setup and integrated any needed plugins, verify that the data transfer is occurring.
3. Click **Next** when this step is complete.

OTX

AT&T Alien Labs™ Open Threat Exchange® (OTX™) is an open information-sharing and analysis network providing users with the ability to collaborate, research, and receive alerts on emerging threats and indicators of compromise (IoCs) such as IP addresses, file hashes, and domains.

You must have an OTX account to receive alerts based on threats identified in OTX. This account is separate from your USM Anywhere account. Go to The World's First Truly Open Threat Intelligence Community to create an OTX account.

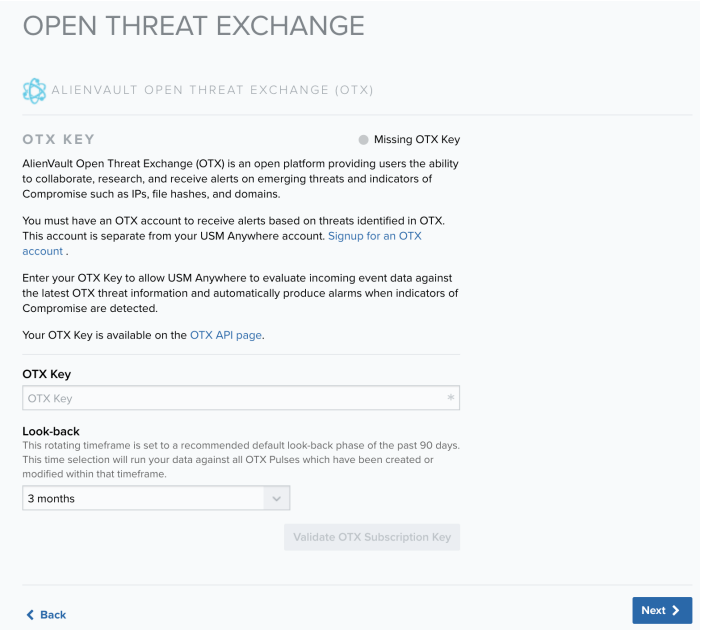

Note: If you do not already have an OTX account, click the **Sign up** link. This opens another browser tab or window that displays the OTX signup page. After you confirm your email address, you can log in to OTX and retrieve the unique API key for your account.

See Open Threat Exchange® and USM Anywhere in the *USM Anywhere User Guide* for more information about OTX integration in USM Anywhere.

To enable USM Anywhere to evaluate event data against the latest OTX intelligence

1. Log in to OTX and open the API page (https://otx.alienvault.com/api).
2. In the DirectConnect API Usage pane, click the icon to copy your unique OTX connection key.

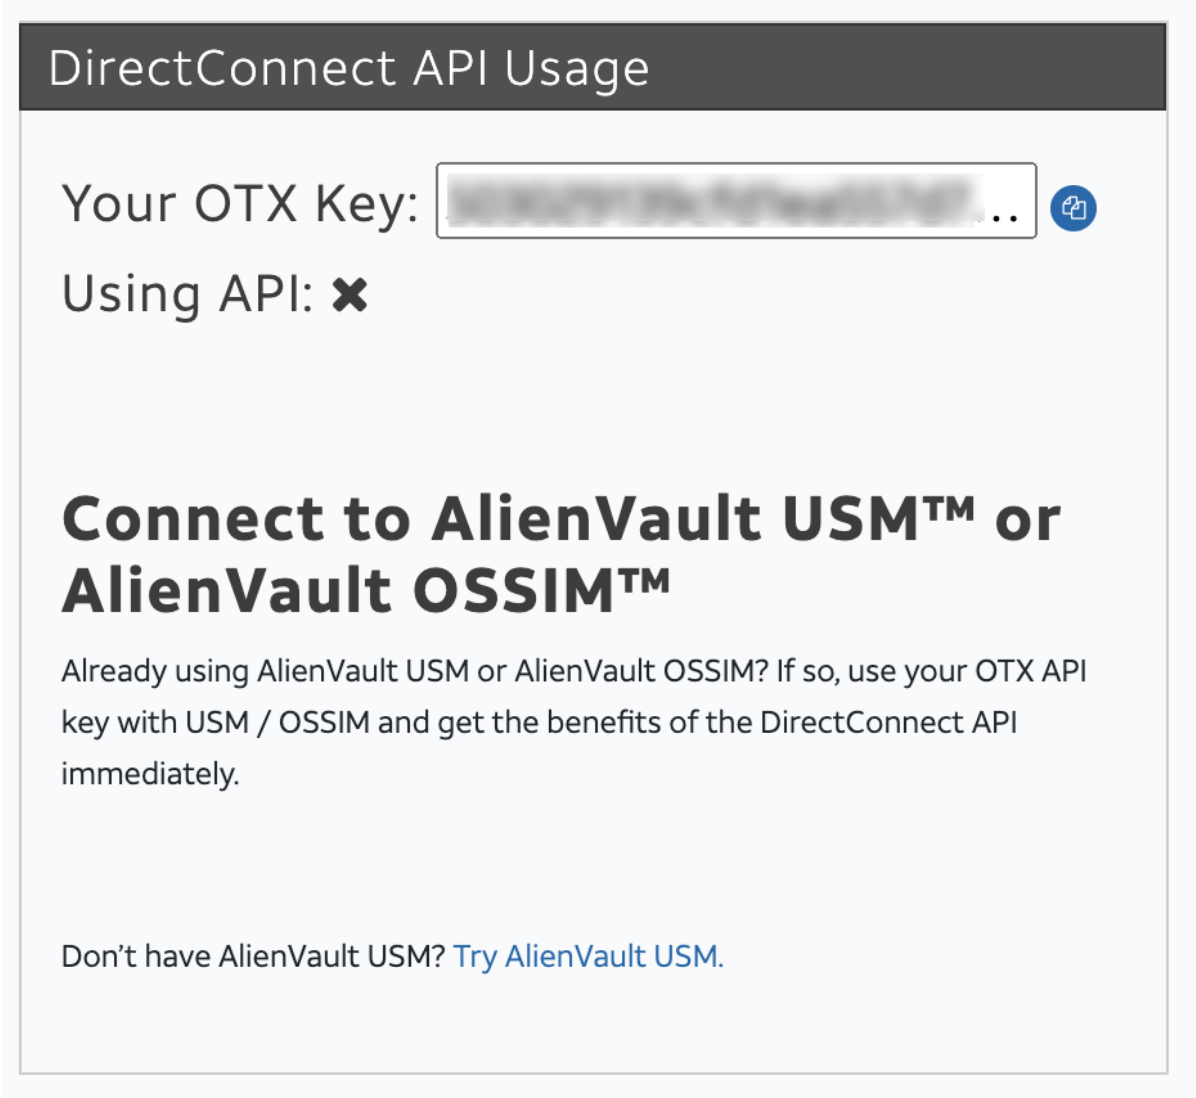

3. Return to the Open Threat Exchange (OTX) page of the USM Anywhere Sensor Setup Wizard and paste the value in the OTX Key text box.

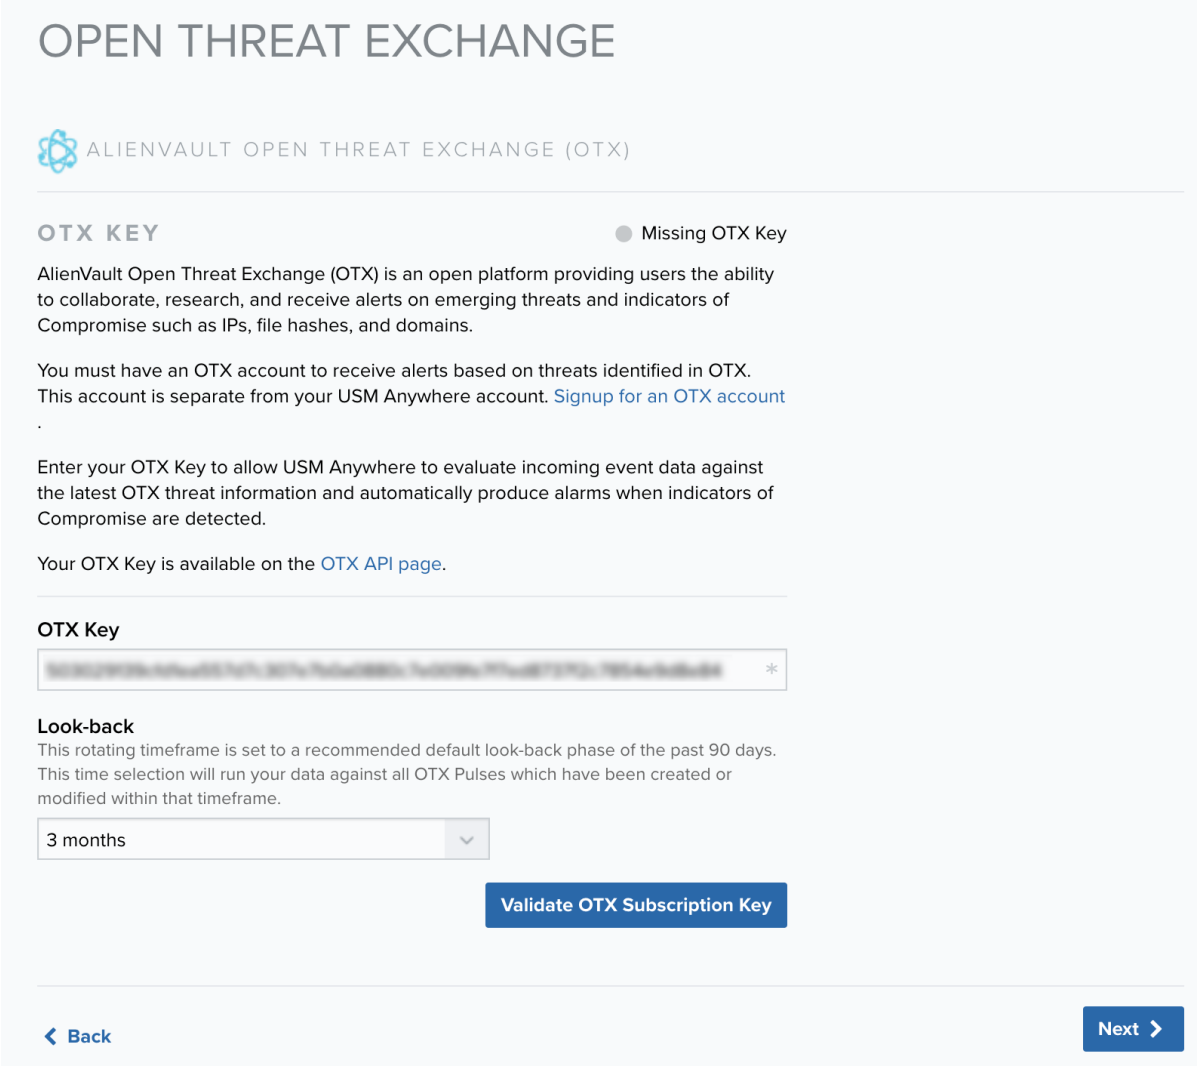

4. Click **Validate OTX Subscription Key**.

   With a successful validation of the key, the status at the top of the page changes to "Valid OTX key".

5. Click **Next** when this task is complete.

Setup Complete

The Congratulations page summarizes the status of your configuration.

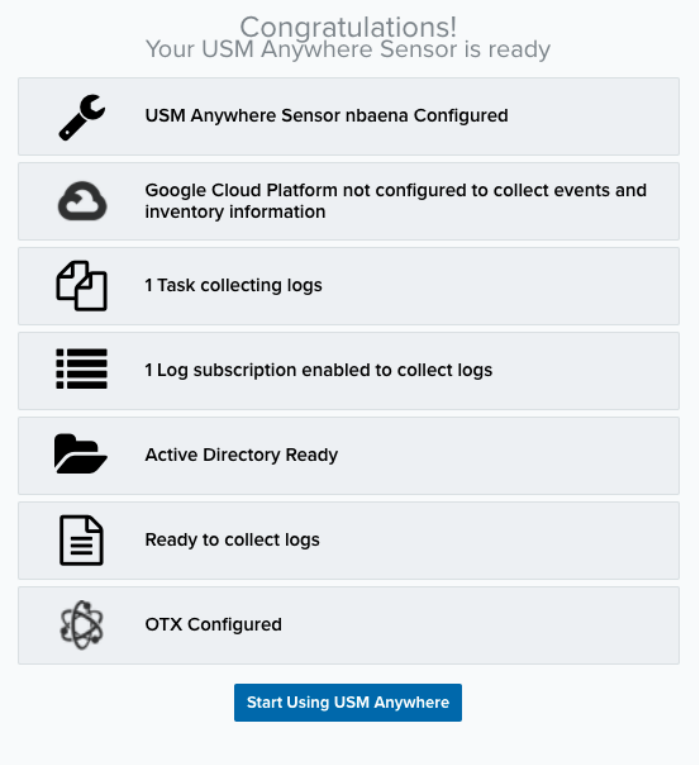

Click **Start Using USM Anywhere**, which takes you to the Overview dashboard.

Next...

Now is a great time to run a vulnerability scan. See Vulnerability Assessment in the *USM Anywhere User Guide* for detailed information about running a vulnerability scan.

GCP Log Discovery and Collection in USM Anywhere

The Google Cloud Platform (GCP) Sensor uses the Google Cloud Pub/Sub to power asynchronous log collection. Cloud Pub/Sub operates based on an export/subscribe model wherein your environment's logs are exported to one or more topics via export sinks. Your GCP Sensor pulls and processes these logs from the enabled Cloud Pub/Sub subscriptions. See the Google Cloud Pub/Sub documentation for more information.

To enable log collection, you can do the following:

- Configure Log Collection for Your GCP Sensor Using Templates
- Manually Create a Cloud Pub/Sub Topic for Your GCP Sensor
- Manually Create and Configure an Export Sink for Your GCP Sensor

Warning: You must have the Cloud Pub/Sub API enabled before beginning these steps

Configure Log Collection Using Templates

Role Availability ✗ Read-Only ✗ Investigator ✗ Analyst ✓ **Manager**

For your USM Anywhere Sensor to receive logs from your Google Cloud Platform (GCP) environment, you must have an export sink to define which logs are exported, a topic to receive those logs, and a subscription to deliver those exported logs to the sensor. The easiest way to create and configure all of these disparate pieces is by using the templates AT&T Cybersecurity provides.

See Manually Create a Cloud Pub/Sub Topic or Manually Create and Configure an Export Sink if you would like to perform these steps manually rather than using these templates.

Important: Because these templates are deployed using the Google Cloud Deployment Manager, you must ensure that both the user executing the deployment and the service account associated with the Cloud Deployment Manager have the required permissions:

- The user executing the deployment must be assigned the role "Deployment Manager Editor" for the project in which they will perform the deployment.
- The service account for the Cloud Deployment Manager must have the "Logging Admin" and "Pub/Sub Admin" roles for the project or organization from which you will be exporting logs.

To configure log collection using templates

1. Download the template files from AT&T Cybersecurity:
   - **Template**: https://storage.googleapis.com/usm-saas-gcp-util/log-export-templates/logExport.py
   - **Project Schema**: https://storage.googleapis.com/usm-saas-gcp-util/log-export-templates/exportProjectLogs.py.schema
   - **Organization Schema**: https://storage.googleapis.com/usm-saas-gcp-util/log-export-templates/exportOrganizationsLog.py.schema
2. Create a Type Registry to deploy the templates by going to the Type Registry page under your Cloud Deployment Manager.
3. Click **Add Composite Type**.

4. Import the templates you previously downloaded.
5. Provide the following information:
   - **Deployment name**: A name for this deployment
   - **source_id**: The identification (ID) of the project exporting these logs.

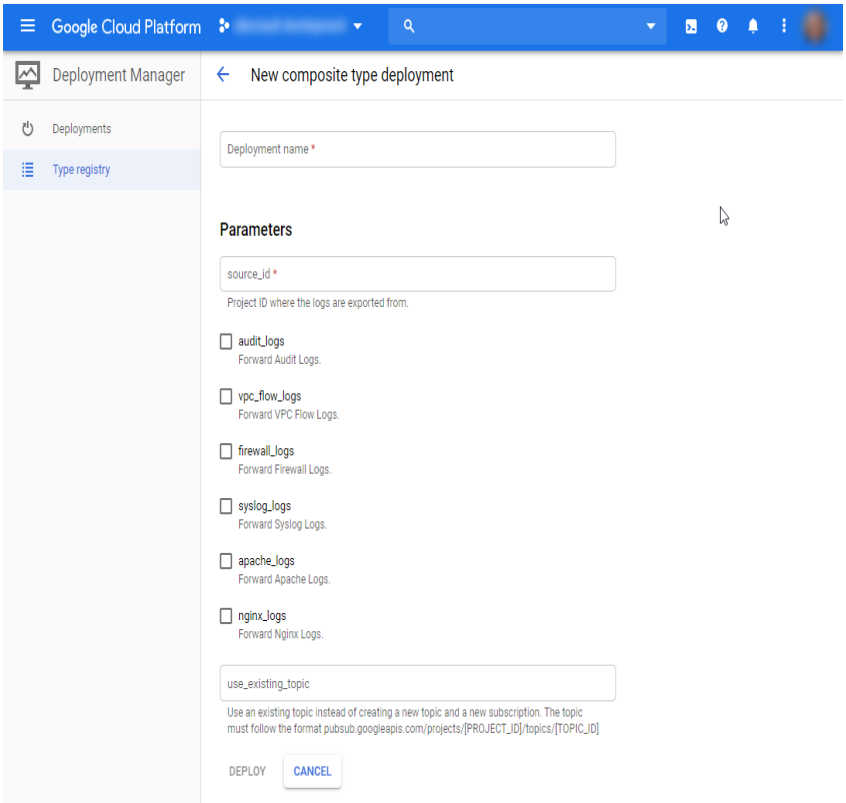

6. If you are executing this deployment at the project level, use the list to select the log types to export.

   Note: See the Log Export Filters table to see how these log queries are formatted.

7. (Optional.) Specify the name of an existing topic to use instead of creating a new one.
8. If you choose to use an existing topic, you must ensure that you grant the Writer Identity service account "Pub/Sub Publisher" permissions.
9. Click **Deploy**.

   You can verify that your topic and subscription have been created by checking the *Topics* page under *Pub/Sub*.

10. In USM Anywhere, go to your GCP Sensor under **Data Sources > Sensors** or the Google Cloud Platform Log Collection app under **Data Sources > AlienApps > Available Apps**.
11. On the Log Subscriptions tab, click **Enable** to enable the subscription you just created.

Manually Create a Cloud Pub/Sub Topic for Your GCP Sensor

Role Availability	Read-Only	Investigator	Analyst	Manager
	✗	✗	✗	✓

In Google Cloud Pub/Sub, a topic receives the logs your Google Cloud Platform (GCP) environment exports. Your GCP Sensor then retrieves those logs via subscriptions. Depending on the needs of your particular implementation, you may only need to create a single topic to receive all of your exported logs from all of the export sinks you configure. However, you may find that it would be advantageous for your implementation to include multiple topics, in which case any number of topics are supported.

To create a Cloud Pub/Sub topic

1. Log in to your GCP environment and go to the Topics page under Pub/Sub.
2. Click **Create Topic**.

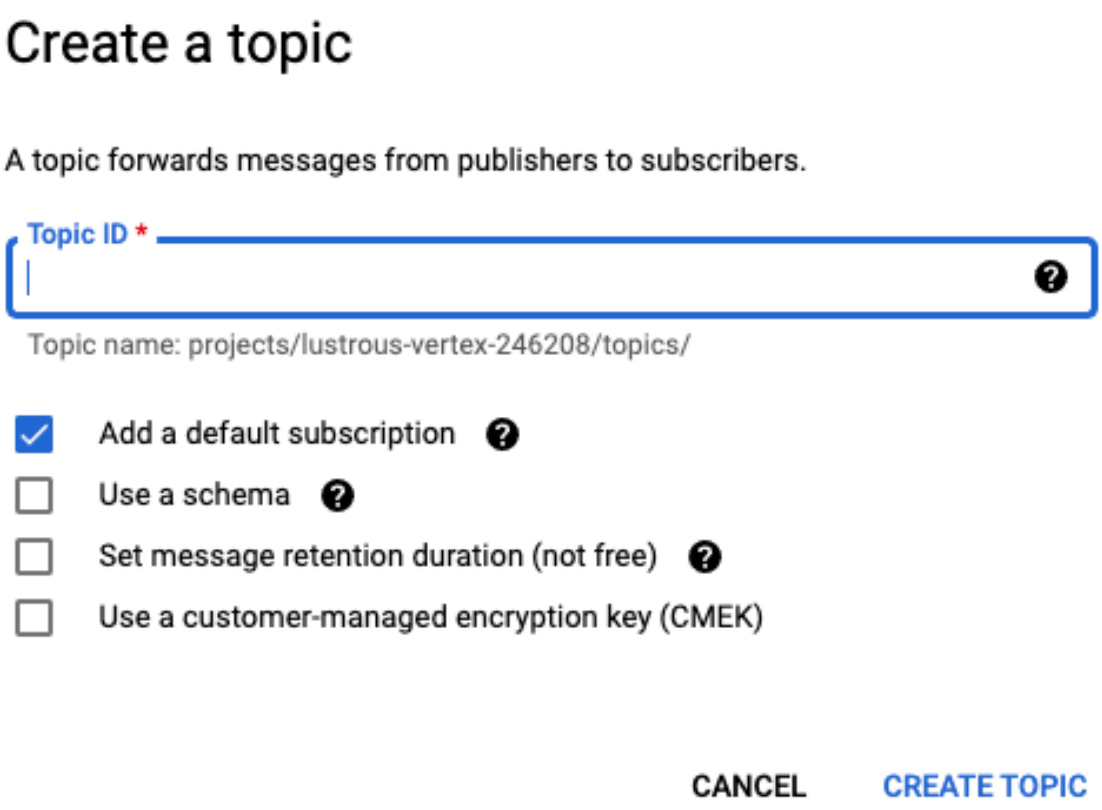

3. Give this topic a name.

   Note: Make note of this name, as you will need to reference it when creating your export sinks.

   By default, the *Add a default subscription* checkbox is selected. The subscription for this topic will be automatically created using the default settings. Its name is [topic]-sub (for example, MyTopic-sub).

   If you leave this checkbox deselected, you will need to create a subscription. See To create a subscription for a Cloud Pub/Sub topic for more information.

4. Click **Create Topic**.

Important: While your subscription is visible at this point, it will not begin reporting events until you have configured at least one export sink to publish to this topic. See Manually Create and Configure an Export Sink for Your GCP Sensor for more information.

To create a subscription for a Cloud Pub/Sub topic

1. Log in to your GCP environment and go to the Topics page under Pub/Sub.
2. Click the Topic ID from which you want to create the subscription to open the specific page of that Topic ID.
3. Go down the page and click **Create Subscription**.

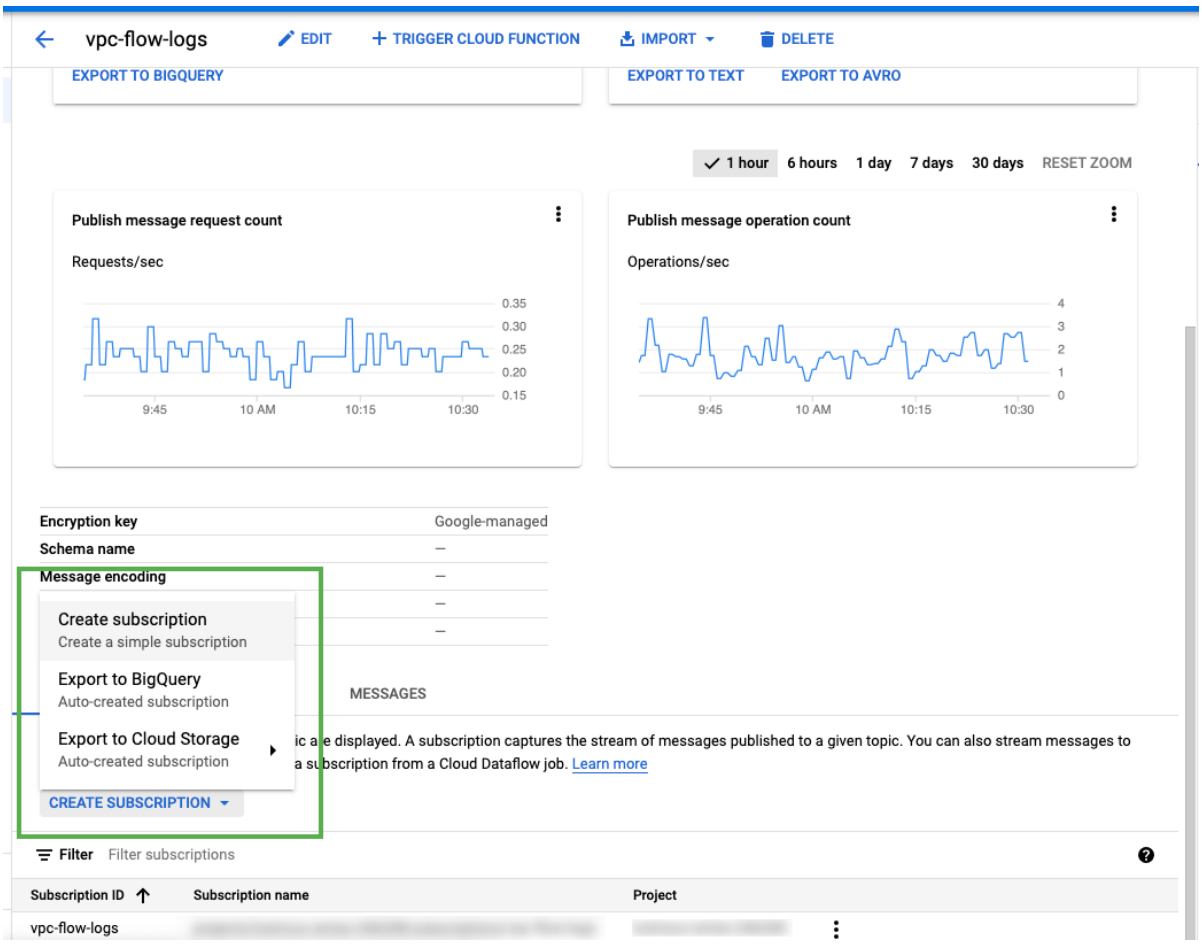

4. Click **Create Subscription**.

5. Name your subscription using the Subscription ID field.

   Note: This is the name that will appear in the user interface (UI) of your GCP Sensor under the Log Subscriptions tab.

6. Click **Create**.

   At this point, you may go to the Sensor Details within your USM Anywhere Sensor and review the Log Subscriptions tab to verify that this subscription appears as expected.

7. Click **Enable** to enable the subscription.

Manually Create and Configure an Export Sink for Your GCP Sensor

The export sink is what defines which logs are exported to a particular topic. You can create a single sink to export all the logs you want your Google Cloud Platform (GCP) Sensor to receive. Or you can create any number of individual sinks to group your exported logs by type, to maximize performance, or for any other reason that suits your specific implementation.

To create an export sink for a project or organization

1. Log in to your GCP environment and go to the organization or project for which you want to create this sink.
2. Go to the Logs Router page under Logging.
3. Click **Create Sink**.
4. Enter the following information:
   - **Sink details**: Enter an identifiable name for this export sink and a description, and then click **Next**.
   - **Sink Destination**: Using the drop-down list, select the topic you created for this sink. Select a Cloud Pub/Sub topic, and then click **Next**.

     Note: If you haven't yet created a topic for this sink, you can select **Create a topic** to create one from this page and immediately use it for your sink. If you do so, you must remember to go to that topic and create a subscription for it or your sensor will not receive any logs from it.

   - **Choose logs to include in sink**: Create an inclusion filter to determine which logs are included in logs routing sink, and then click **Next**.

- (Optional.) **Choose logs to filter out of sink**: Create exclusion filters to determine which logs are excluded from logs routing sink. And then configure a filter for this sink, following the guidelines in Configuring Export Sink Filters.

5. Click **Create Sink**.

Important: If your sink and topic are in different GCP projects, or if you are exporting organization-level logs to a Google Cloud Pub/Sub topic in a project, you must complete some additional steps. See the following sections for detailed instructions regarding those two cases.

To create a sink that publishes to a Cloud Pub/Sub topic in a different project

Note: If you have not already granted your service account permission to this second project, first use the instructions in Preparing Your GCP Environment for Sensor Deployment to grant permission to this project now. Be sure to restart the sensor app before proceeding on to step one.

1. Log in to your GCP environment and go to the project for which you want to create this sink.
2. Go to the Logs Router page under Logging.
3. Click **Create Sink**.
4. Enter the following information:
   - **Sink details**: Enter an identifiable name for this export sink and a description, and then click **Next**.
   - **Sink Destination**: Using the drop-down list, select **Cloud Pub/Sub topic**, and select a Cloud Pub/Sub topic, and then click **Next**.

     Note: If you haven't yet created a topic for this sink, you can select **Create a topic** to create one from this page and immediately use it for your sink. If you do so, you must remember to go to that topic and create a subscription for it or your sensor will not receive any logs from it.

     When you make your selection in Sink Destination, the menu item transforms into a text field. Use that field to enter the following, substituting your relevant information where there are variables:

     ```
     pubsub.googleapis.com/projects/<project-id>/topics/<topic_name>
     ```

     Where the <project-id> you reference is the project your topic resides in.

- **Choose logs to include in sink**: Create an inclusion filter to determine which logs are included in logs routing sink, and then click **Next**.
- (Optional.) **Choose logs to filter out of sink**: Create exclusion filters to determine which logs are excluded from logs routing sink. And then configure a filter for this sink, following the guidelines in Configuring Export Sink Filters.

5. Click **Create Sink**.

To create a sink to publish from an organization to a topic in a project

Important: Unlike the previous methods, it is not possible to use the web user interface (UI) to create an export sink to publish from the organization level to a topic at the project level. Instead, use the Google Cloud Shell Editor native to your GCP environment to enter the following commands.

1. Access the Cloud Shell editor in your GCP environment by clicking the **Activate Cloud Shell** button.
   This opens a new window at the bottom of your screen, which may take a few minutes to finish loading.

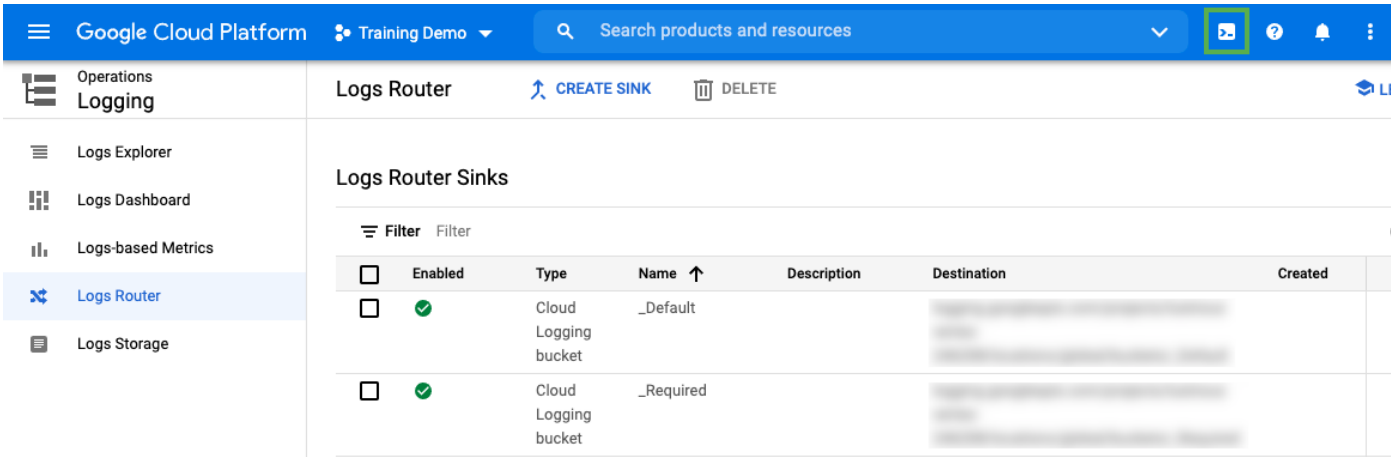

2. Use the following command to create a new sink for your organization:

```
gcloud logging sinks create \
       <sink-name> \
       --organization=<organization-id> \
       --include-children \
       pubsub.googleapis.com/projects/<project-name>/topics/<topic-name> \
       --log-filter "logName=(\"organizations/<organization-
id>/logs/cloudaudit.
               googleapis.com%2Factivity\" OR \"organizations/<organization-
```

```
id>/logs/
                cloudaudit.googleapis.com%2Fdata_access\" OR
\"organizations/<organization-id>
                /logs/cloudaudit.googleapis.com%2Fsystem_event\")"
```

This returns the following message. Make note of the service account name (highlighted here in bold) to enter in the next step.

```
Created
[https://logging.googleapis.com/v2/organizations/<organization_id/
       sinks/<sink_name>].
Please remember to grant `serviceAccount:<name-of-sensor-service-
account>@
       <name-of-project>.iam.gserviceaccount.com.com` the Pub/Sub
Publisher
       role on the topic.
More information about sinks can be found at
https://cloud.google.com/logging/docs/
       export/configure_export
```

3. Use the following command to grant the service account the permissions it requires:

```
gcloud organizations add-iam-policy-binding <organization_id> \
       --member=<name-of-sensor-service-account>@<name-of-project>.iam.
               gserviceaccount.com> \
       --role=roles/pubsub.publisher
```

Configuring Export Sink Filters

The filter configured for your export sink determines which logs that sink exports to your topic.

To configure the filters for your sink

1. Go to the export sink for which you wish to create a filter.

   You can do this either when you first created the export sink or by opening it any time after that for editing.

2. Click the carrot in the text box of your export filter and select **Convert to advanced filter**.

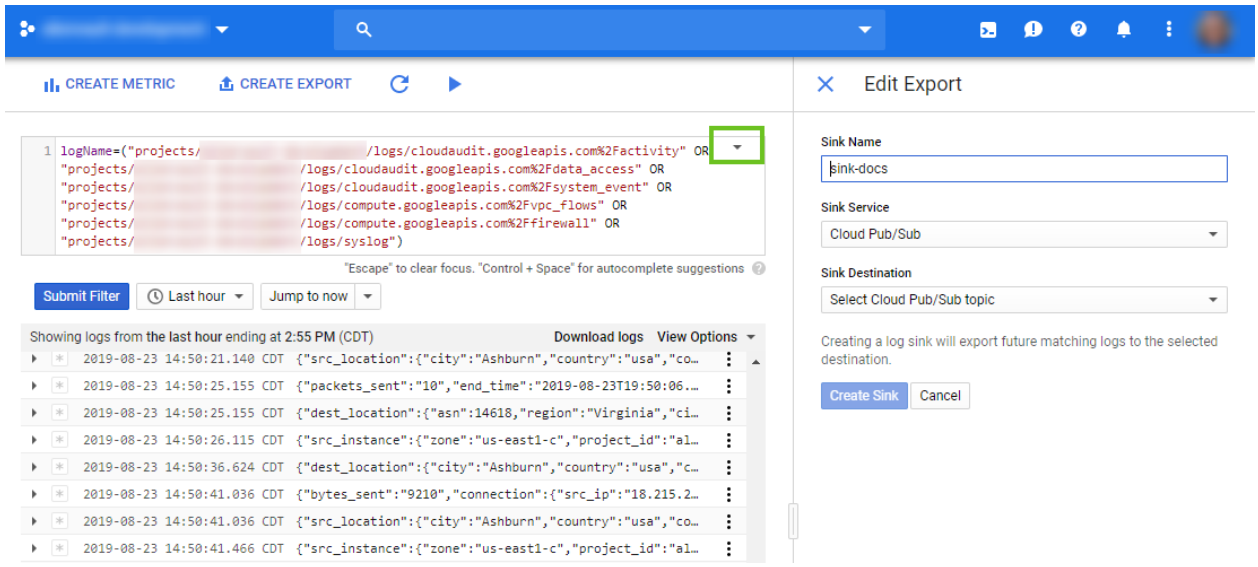

3. Use the specifications described in the following table to define which filters will be exported by this sink, separating each filter specification with "OR" (as seen in the preceding image).

Note: Any logs included in your filter but not supported by the GCP Sensor will be ignored by the sensor. AT&T Cybersecurity recommends including syslog in your filter to collect these unsupported logs.

The GCP Sensor relies on hints to parse syslog logs, meaning that any logs that can be assigned to a plugin will be, while the remainder will be parsed as generic events. See AlienApps and Data Sources for more information about how hints help USM Anywhere parse logs to plugins.

Log Types Supported by the GCP Sensor

Log Type	Filter to Capture This Log	Notes
Audit Logs at the Organization Level	organizations/<organization-id>/logs/cloudaudit.googleapis.com	To filter these logs further, append the following: • **%2Factivity**: For activity logs • **%2Fdata_access**: For data access logs • **%2Fsystem_event**: For system events
Audit Logs at the Project Level	projects/<project-id>/logs/cloudaudit.googleapis.com	To filter these logs further, append the following: • **%2Factivity**: For activity logs • **%2Fdata_access**: For data access logs • **%2Fsystem_event**: For system events
VPC Flow Logs	projects/<project-id>/logs/compute.googleapis.com%2Fvpc_flows	
Firewall Logs	projects/<project-id>/logs/compute.googleapis.com%2Ffirewall	
Syslog	projects/<project-id>/logs/syslog	These logs are delivered via the Stackdriver logging agent

Log Types Supported by the GCP Sensor (Continued)

Log Type	Filter to Capture This Log	Notes
Apache Logs	projects/<project-id>/logs/apache	To filter these logs further, append the following: • **-access**: For access logs • **-error**: For error logs
Nginx Logs	projects/<project-id>/logs/nginx	To filter these logs further, append the following: • **-access**: For access logs • **-error**: For error logs

The AWS Cloud Connector Deployment in USM Anywhere

The Amazon Web Services (AWS) Cloud Connector provides operational visibility into the security of your AWS environment. Based on the collected log information, USM Anywhere receives the data stored in your Amazon Simple Storage Service (S3) buckets, generates the related events for that data within USM Anywhere, and provides real-time alerting to identify malicious activity.

Important: USM Anywhere starts processing the files contained within Amazon S3 buckets after enabling the AWS Cloud Connector. Any files contained within Amazon S3 buckets before setting up a Cloud Connector will not be processed.

Differences Between an AWS Cloud Connector and a Sensor

Before choosing between an AWS Cloud Connector and a USM Anywhere sensor, you need to know how they work and the existing differences between them.

This table includes a summary of the main differences between an AWS Cloud Connector and a sensor.

Differences Between an AWS Cloud Connector and a Sensor

Item	AWS Cloud Connector	Sensor
Deploy a sensor	✗	✓
Create a virtual machine (VM)	✗	✓
Inventory data detection (users and assets)	✗	✓
NIDS	✗	✓
AlienApps	✗	✓

Differences Between an AWS Cloud Connector and a Sensor (Continued)

Item	AWS Cloud Connector	Sensor
Maintenance, updates, upgrades	✗	✓
Upload an AWS CloudFormation template into the AWS account	✓	✗
Monitor multiple AWS accounts	✓ (one connector per account)	✓ (one sensor per account)
Receive Amazon S3 events	✓	✓
Log aggregation	✓	✓

Warning: You will have duplicate events if your sensor is monitoring buckets from an AWS account and you configure an AWS Cloud Connector in the same account monitoring the same buckets.

Keep in mind these points when you are going to choose between an AWS Cloud Connector and a USM Anywhere sensor:

- A sensor requires a deployment. An AWS Cloud Connector doesn't need to deploy a sensor on a VM; instead, it requires an upload of an AWS CloudFormation template that you generate within the USM Anywhere user interface (UI). See Adding an AWS Cloud Connector for more information. This process is much easier and, unlike a sensor, it doesn't require ongoing maintenance.
- A sensor detects inventory data automatically in your account, such as users and assets. An AWS Cloud Connector receives Amazon S3 events but doesn't detect users and assets. Deploying a sensor is the best choice if you a specific account that needs to automatically detect users or assets in the AWS environment that are you monitoring.
- An AWS Cloud Connector receives Amazon S3 events, but no events from network-based intrusion detection systems (NIDS) nor AlienApps. Deploying a sensor is the best choice if you have a specific account that needs either NIDS or AlienApps, and that are critical for an AWS environment that are you monitoring.

> **Important:** If you have multiple AWS accounts, you can configure some of them with sensors and the rest with AWS Cloud Connectors. You can have a mix of deployments, but best practice is to only deploy one connector or one sensor per AWS account.

- An AWS Cloud Connector is easier to maintain. For example, a sensor often requires upgrades.

Activating an AWS Cloud Connector

To activate an AWS Cloud Connector, you must follow these steps:

1. Add a new connector.

   See Adding an AWS Cloud Connector for more information.

2. Download the AWS CloudFormation template.

   See Downloading an AWS Cloud Connector Template for more information.

3. Create a stack to upload the AWS CloudFormation template.

   See Uploading AWS CloudFormation Templates for more information.

4. Go to USM Anywhere to enable the AWS Cloud Connector.

   See AWS Cloud Connector List View for more information.

Uploading AWS CloudFormation Templates

Role Availability	Read-Only	Investigator	Analyst	Manager
	✗	✗	✓	✓

Through USM Anywhere you can generate the Amazon Web Service (AWS) CloudFormation templates that you need to begin gathering data from your Amazon Simple Storage Service (S3) buckets. See Downloading an AWS Cloud Connector Template for more information.

To upload an AWS CloudFormation template

1. Open you AWS Management Console page and go to CloudFormation.
2. Click **Create stack > With new resources (standard)**.
3. Select **Upload a template file** and then click **Choose file** to select the template you have downloaded from USM Anywhere.

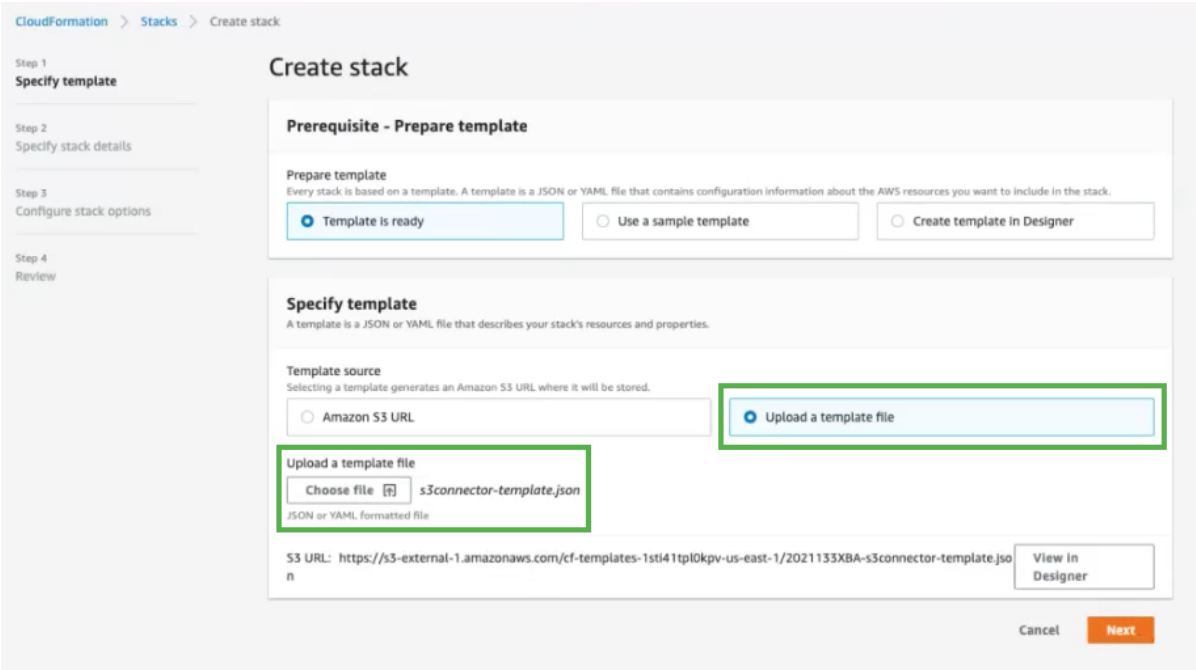

4. Click **Next**.
5. In the stack name, enter a name for your stack.
6. Use the bucketArns field to enter the Amazon S3 buckets names where you currently store security logs.

   There are two options:

   - Enter the existing Amazon S3 bucket Amazon Resource Names (ARNs) that contain the logs you would like to monitor. You can enter several ARNs separated by commas.

     Important: If you choose this option, you must enable the event notifications. See To enable the event notifications for Amazon S3 buckets you have selected for more information.

   - Leave this field empty and a new Amazon S3 bucket will be automatically created.

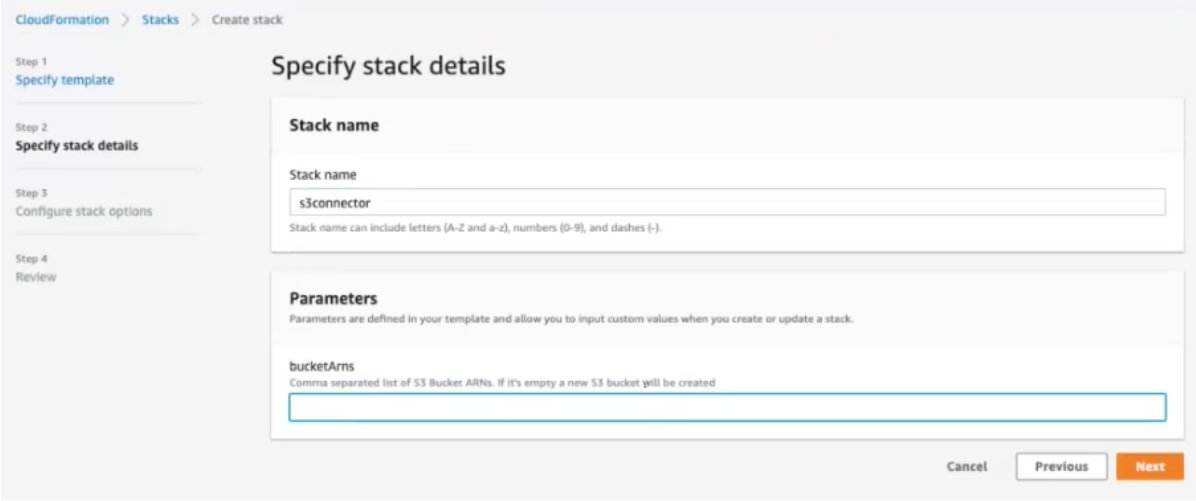

Note: If you later decide to enter new or additional Amazon S3 buckets, you can come here and add them.

7. Click **Next**.
8. (Optional.) If your organization requires tags, you may enter them at this point. You can also leave them blank.
9. Click **I acknowledge that AWS CloudFormation might create IAM resources with custom names**.
10. Click **Create stack**.

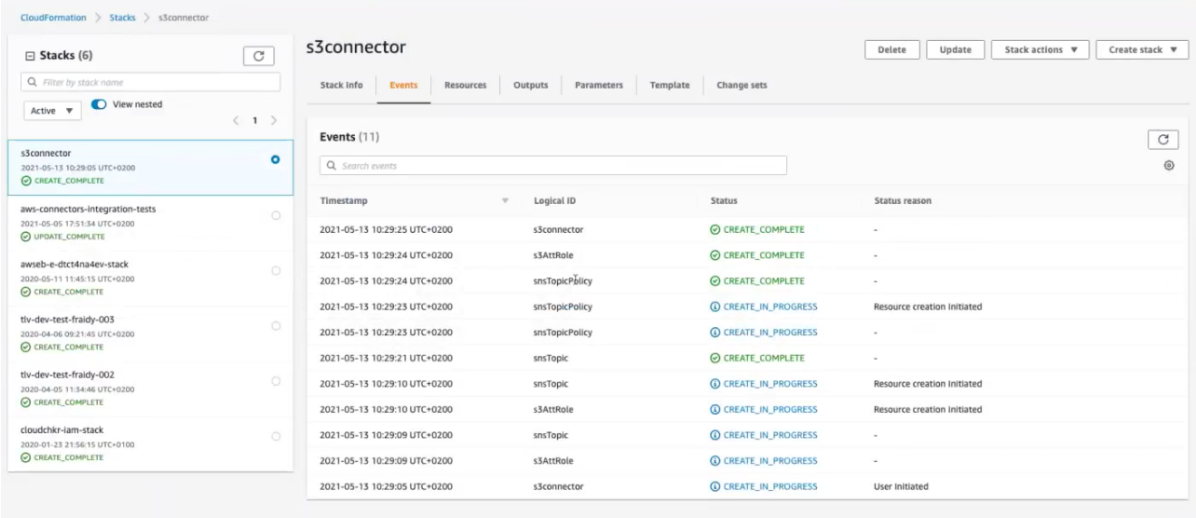

To enable the event notifications for Amazon S3 buckets you have selected

Important: You must enable event notifications if you have entered your Amazon S3 bucket Amazon Resource Names (ARNs) in the bucketArns field. If you left this field empty, it is not necessary to enable the event notifications. If you don't do these instructions, USM Anywhere will not receive events from your Amazon S3 buckets.

1. Open you AWS Management Console page and go to CloudFormation.
2. Select the buckets you have entered previously in the bucketArns field.
3. Click the **Properties** tab.
4. Click **Create event notification** to enable the event notifications.

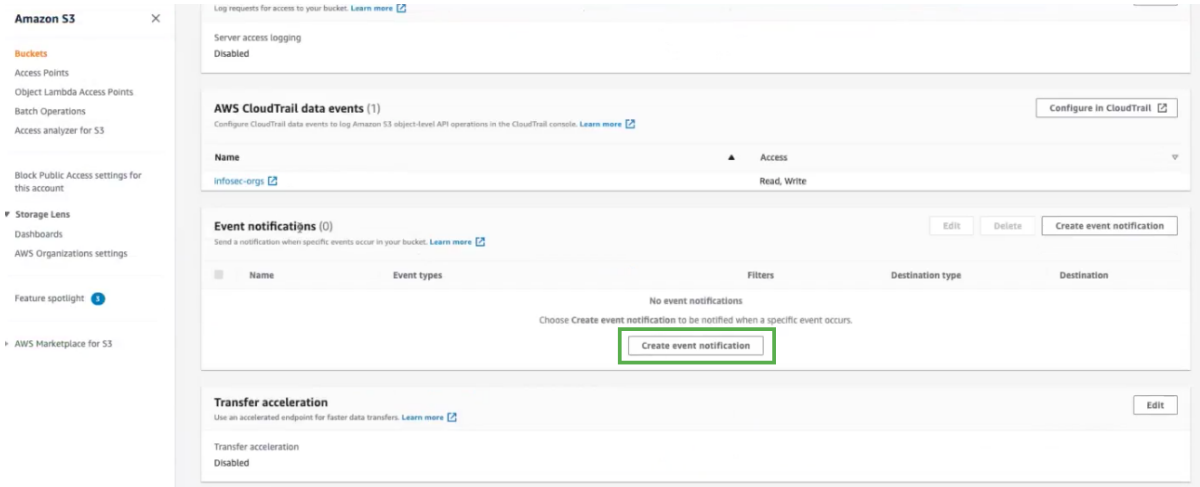

5. In the general configuration section, enter an event name.
6. In the event types section, select the **All object create events** option.
7. In the destination section, click the **SNS topic** option, and then select **attcs-s3-connector**.

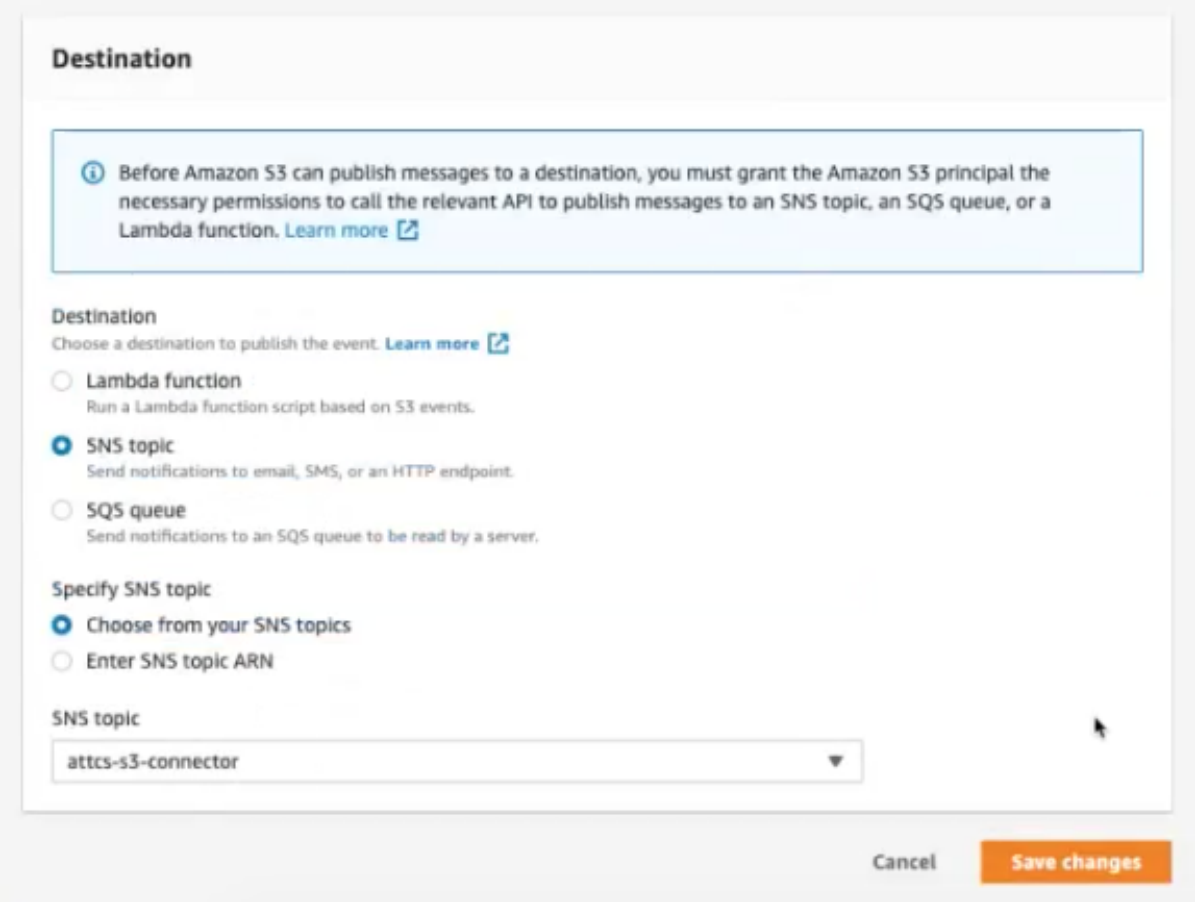

8. Click **Save changes**.
9. Go to USM Anywhere to enable the AWS Cloud Connector.

   See AWS Cloud Connector List View for more information.

AWS Cloud Connector Resources

Through USM Anywhere you can generate the Amazon Web Service (AWS) CloudFormation template that you need to begin gathering data from your Amazon Simple Storage Service (S3) buckets. See Downloading an AWS Cloud Connector Template for more information about how to download a CloudFormation template.

After uploading a CloudFormation template, some resources have been created.

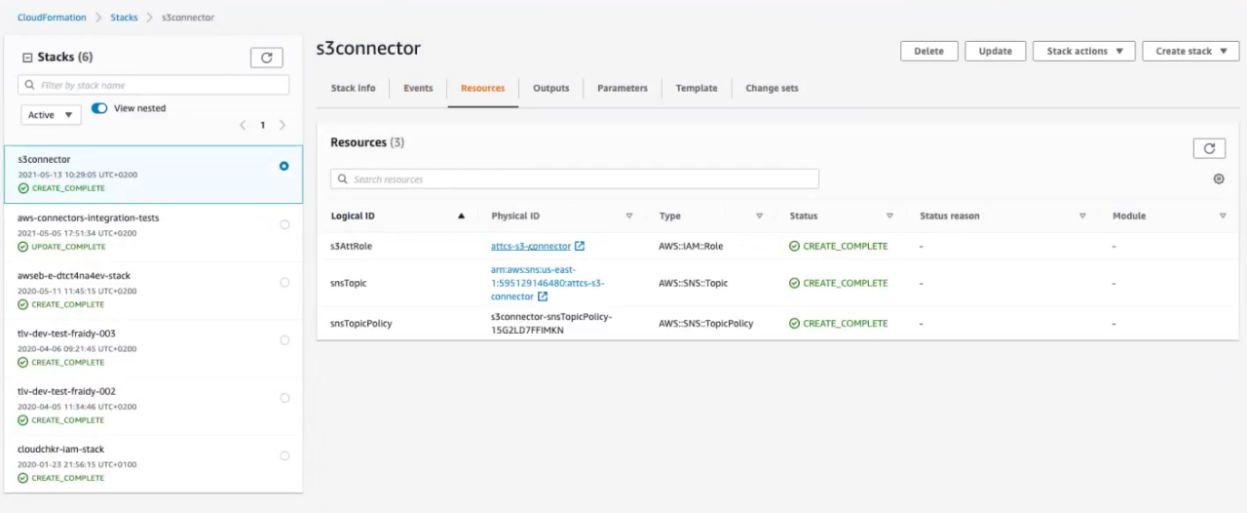

These are the created resources:

- **s3AttRole**: The role enables USM Anywhere the access to your account to read the configured buckets.
- **snsTopic**: It enables you to receive notifications when there are new files in the bucket.
- **snsTopicPolicy**: The access policy that enables the configured Amazon S3 bucket to publish notifications, and the USM Anywhere account that has permissions to subscribe to that amazon Simple Notification Service (SNS).

Network Setup and Configuration

The Setup Wizard helps first-time users configure USM Anywhere Sensor capabilities within minutes. Its simple, step-by-step workflow guides you through the basic elements for getting your USM Anywhere environment up and running.

- Asset discovery and authenticated scans
- Network monitoring
- Log management

But, before you can start seeing all of your events, there are certain configurations within your network that must be in place to ensure that USM Anywhere is receiving all of your relevant event logs:

- Network firewall allows communication with USM Anywhere
- Operating systems forward all relevant logs to USM Anywhere

This section includes the following topics:

Configure Network Interfaces for On-Premises Sensors 309

Configure USM Anywhere to Receive ERSPAN Traffic 320

Port Mirroring Configuration on Network Devices 321

Granting Access to Active Directory for USM Anywhere 332

Proxy Configuration on the USM Anywhere Sensor 334

Configure Network Interfaces for On-Premises Sensors

A USM Anywhere Sensor deployed on VMware or Hyper-V uses five network interfaces. These network interfaces have a predefined role that cannot be changed. The USM Anywhere management interface is required for many essential functions, including the following:

- Connection to USM Anywhere
- Updates to the system
- Log collection within the monitored network
- Vulnerability scans
- Asset discovery

The management interface needs an IP address with permissions to access the following:

- Inbound packets containing syslog data sent from other hosts on that network
- Outbound connections made to perform authenticated scans

The other interfaces passively monitor network traffic in promiscuous mode; the system does allow the configuration of an IP address on them. These interfaces should be plugged into a port in the switch where port mirroring is configured. The following table summarizes each interface's usage.

Network Interfaces Available in On-Premises Sensors

Interface Name	Network Configuration Required
Management Interface	Internet connectivity and IP address routed to provide the access to USM Anywhere. This IP address also allows connections to assets in a monitored network for log collection and asset scans.
Network Monitoring Interface 1	Interface connected to a mirrored port in the network switch 1.
Network Monitoring Interface 2	Interface connected to a mirrored port in the network switch 2.
Network Monitoring Interface 3	Interface connected to a mirrored port in the network switch 3.
Network Monitoring Interface 4	Interface connected to a mirrored port in the network switch 4.

> The VMware Sensor and Hyper-V Sensor require *all five network interface cards (NICs)* to be enabled; otherwise, the USM Anywhere update will fail. The NICs can remain disconnected.
>
> You should only connect the other NICs to any additional network you want to monitor. Don't connect the NICs to the same Switched Port Analyzer (SPAN) port because it'll produce duplicate events in USM Anywhere.

Use the functions provided by the sensor console to configure the management interface and your Domain Name System (DNS).

Setting Up the Management Interface

By default, USM Anywhere has Dynamic Host Configuration Protocal (DHCP) and log collection enabled.

To configure the management interface automatically using DHCP

During the installation, your system sets an IP address assigned by a DHCP server. You can check the IP address afterwards:

1. Connect to the USM Anywhere Sensor console.
2. Go to **Network Configuration > View Network Configuration**.

To manually configure the management interface

1. Connect to the USM Anywhere Sensor console.
2. Go to **Network Configuration > Configure Management Interface > Set a Static Management IP Address**.

   > **Note:** The Configure Management Interface option is only available on VMware and Hyper-V Sensors.

3. Enter the IP address.
4. Press **Enter**.

Defining the DNS nameservers

The DNS nameserver is part of the DNS that maintains a directory of domain names and translates them to IP addresses.

Important: If you specify two servers for DNS resolution, USM Anywhere determines their priority by their order. Configure your local DNS in the first position to have DNS name resolution in your internal network.

To define the DNS Nameservers

1. Connect to the USM Anywhere Sensor console.
2. Go to **Network Configuration > Configure DNS**.

   Note: The Configure DNS option is only available on VMware and Hyper-V Sensors.

3. Enter the primary DNS and press **Enter**.

   A confirmation screen opens to apply changes.

4. Select **Yes**.
5. Optionally, you can provide the secondary DNS and press **Enter**.

   When the confirmation screen appears to apply changes, select **Yes**.

Creating a Firewall Rule for Communication Between USM Anywhere Sensor and Cloud Service

USM Anywhere is hosted as a cloud service with an IP address that is not statically assigned and may change periodically. For this reason, you must set up a firewall rule that uses the URL of the cloud service to allow incoming and outgoing traffic between the USM Anywhere Sensor and the cloud service.

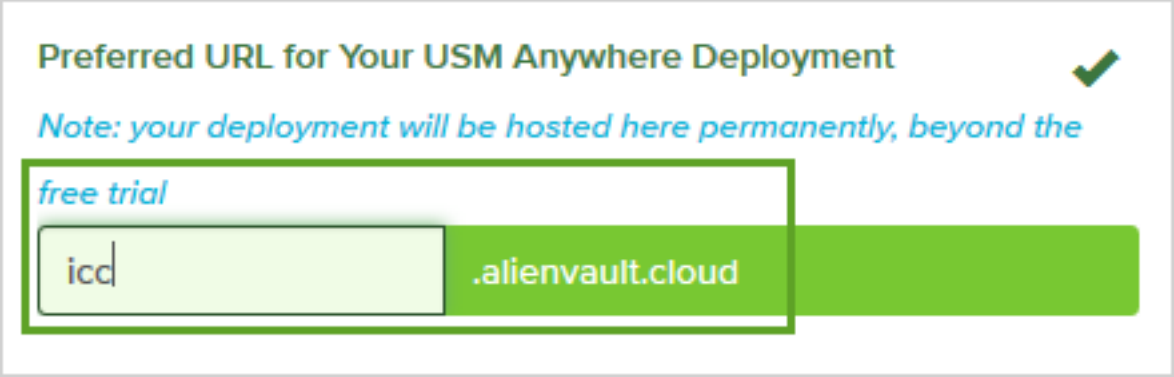

In this example, the URL for the USM Anywhere instance is displayed within the green box.

Checking Your Settings

You can verify your network settings in the USM Anywhere Sensor Setup wizard or through the sensor console.

To verify the network settings in the USM Anywhere web user interface (UI)

1. Go to **Data Sources > Sensors** and click the USM Anywhere Sensor name.

   At the bottom of the USM Anywhere Sensor page, click the **Network IDS** tab, where you can view the traffic in your network over various interfaces.

   > **Important:** The interface will only show as receiving data if it is receiving more than 1000 packets over a 30-second period.

   You can configure a new interface as well as port mirroring here. See the following documentation for more information:

   - Direct Traffic from Your Physical Network to the VMware Sensor
   - Direct Traffic from Your Physical Network to the Hyper-V Sensor

The **Network IDS** tab also allows you to configure your Classless Inter-Domain Routing (CIDR) blocks by clicking the **Configure CIDR Blocks** button. Your CIDR blocks are automatically populated by the setup wizard during the initial USM Anywhere Sensor

deployment. By default, the system will scan all internal IPv4 addresses and assign their names based on those designated in your asset groups. If you want to remove a block or change the subnet range of the block, click the **x** button next to the CIDR block to remove it, and click **Add Another CIDR Block** to input a new CIDR block with the desired subnet range. Be aware, however, that removing part of a subnet range or deleting a block completely will result in the sensor no longer monitoring that portion of your internal network.

To verify the network settings in the USM Anywhere Sensor console

1. Connect to the USM Anywhere Sensor console.
2. Go to **Network Configuration > View Network Configuration**.

Direct Traffic from Your Physical Network to the VMware Sensor

For USM Anywhere to monitor traffic from your physical network, you need to allocate a spare network interface card (NIC) to pass the mirrored traffic, or in Cisco terms the Switched Port Analyzer (SPAN) traffic, to the virtual network. AT&T Cybersecurity recommends that you implement SPAN on your internal firewall ports, connect the SPAN port to the spare NIC, and then associate the spare NIC with a virtual switch (vSwitch) on your VMware server, as illustrated in the following diagram:

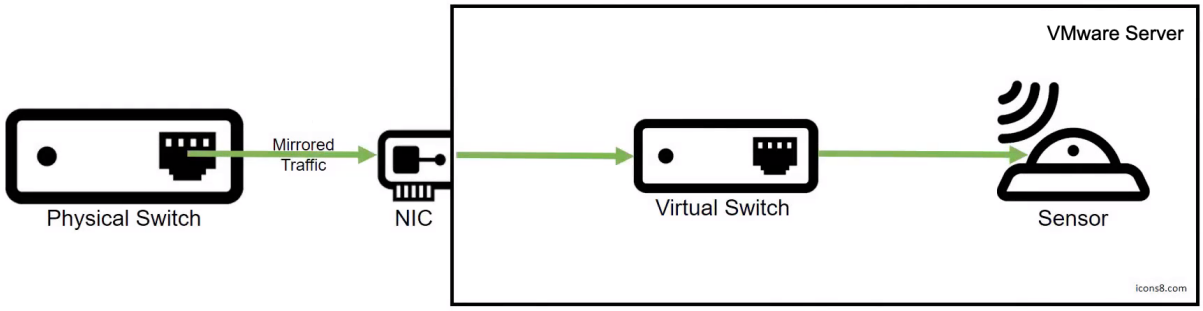

Illustration for Directing Network Traffic from Your Physical Network to the VMware Sensor

Important: USM Anywhere provides multiple network interfaces to monitor your network. To avoid duplicating data, you should not connect them all to the same vSwitch. Instead, you can connect each interface to a different vSwitch dedicated to a different subnet within your network or a different virtual local area network (VLAN).

In the following procedure, you will create a new standard vSwitch in VMware vSphere, configure it to allow promiscuous mode, and then assign it to one of the network adapters on the USM Anywhere VMware Sensor virtual machine (VM). It is important to create a new vSwitch dedicated to the mirrored traffic. Adding a promiscuous port group to an existing vSwitch may cause instability in the hypervisor.

This procedure assumes that you have completed the following tasks:

- Enabled port mirroring on the network you want USM Anywhere to monitor.
- Allocated a spare NIC on the VMware ESXi server to receive the mirrored traffic.

To direct the mirrored traffic to the VMware Sensor

1. Configure a new standard vSwitch specifically for the mirrored traffic (see VMware Documentation for detailed instructions):
   a. For the connection type, select **Virtual Machine Port Group for a Standard Switch**.
   b. Add the spare NIC as the network adapter for the new switch.
   c. For the connection settings, enter a new network label for the port group, for example, *SPAN Target*.
   d. Enter a VLAN number or select **All (4095)**, which enables the switch to capture traffic from all the VLANs connected to the spare NIC.
2. Configure the port group to allow promiscuous mode so that connected devices can view traffic on the entire switch:
   a. Next to the new vSwitch, click **Properties**.
   b. Select the vSwitch and click **Edit**.
   c. Set Promiscuous Mode to **Accept**, and click **OK**.
   d. Select the port group and make sure that the default security policy permits promiscuous mode there as well.

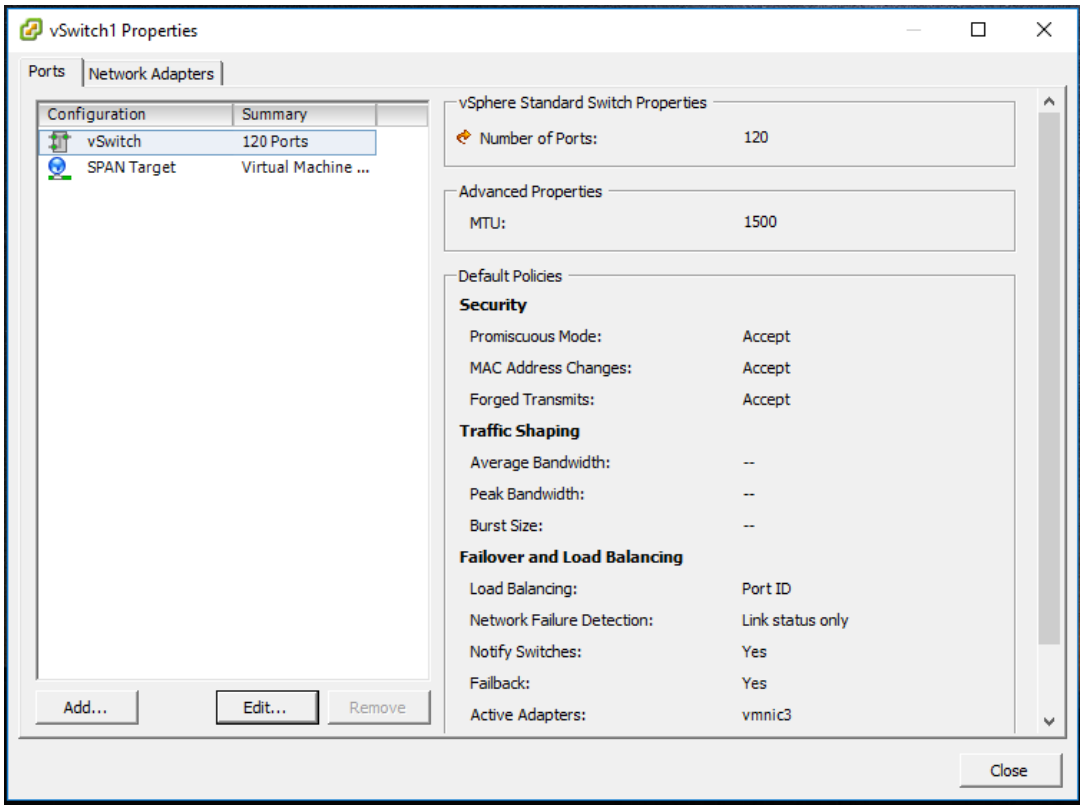

e. Select the **Network Adapters** tab and make sure that your spare NIC is associated with the vSwitch.

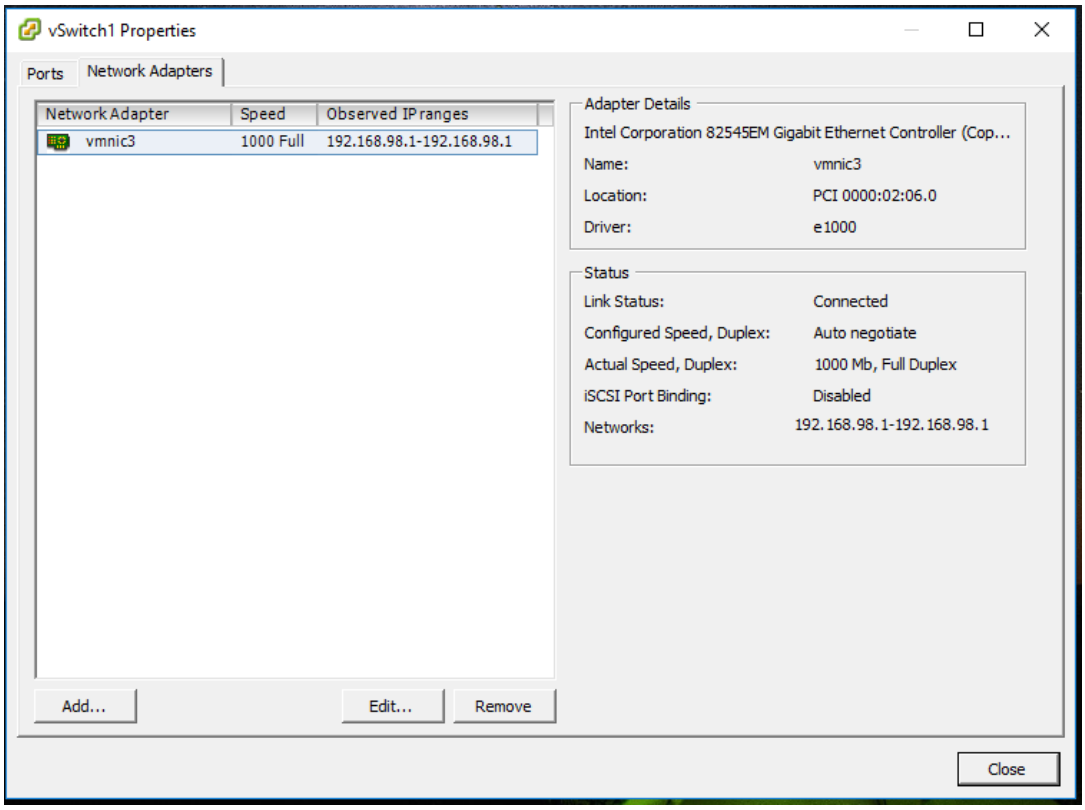

f. In the dialog box, click **Close**.

3. Connect the vSwitch to your VMware Sensor.

   a. Edit the VMware Sensor VM and select an available network adapter.

   > **Note:** Network adapter 1 is reserved for the management interface. See Configure Network Interfaces for On-Premises Sensors for more information.

   b. Associate the adapter with the vSwitch and save your changes.

   c. Restart the VM if changes are not automatically applied.

4. Repeat the steps for every vSwitch you want to monitor.

Direct Traffic from Your Physical Network to the Hyper-V Sensor

For USM Anywhere to monitor traffic from your physical network, you need to send a copy of the network traffic to the Hyper-V Sensor, as illustrated in the following diagram:

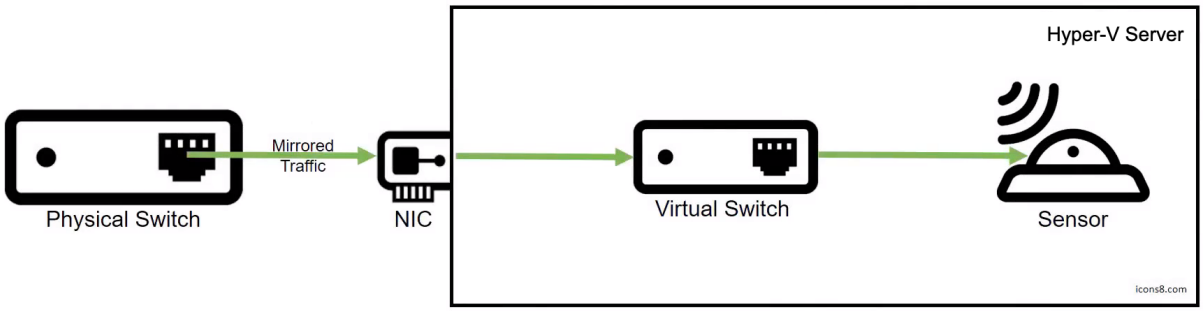

Illustration for Directing Network Traffic from Your Physical Network to the Hyper-V Sensor

Important: USM Anywhere provides multiple network interfaces to monitor your network. To avoid duplicating data, you should not connect them all to the same virtual switch. Instead, you can connect each interface to a different virtual switch dedicated to a different subnet within your network or a different virtual local area network (VLAN).

In the following procedure, you will configure the Hyper-V Sensor virtual machine (VM) as the destination and the virtual switch as the source of the mirrored traffic.

This procedure assumes that you have completed the following tasks:

- Enabled port mirroring on the network you want USM Anywhere to monitor.
- Allocated a spare network interface card (NIC) to direct the mirrored traffic to the virtual switch.

Configure a Port Mirroring Destination

When configuring your Hyper-V Sensor to capture mirrored traffic, you need to set the appropriate network adapter on the Hyper-V Sensor VM as the destination for port mirroring.

To configure your Hyper-V Sensor to capture mirrored traffic

1. Open the Microsoft Hyper-V Manager and right-click the Hyper-V Sensor VM.
2. Select **Settings**.
3. Expand the available network adapter and select **Advanced Features**.

   Note: Network adapter 1 is reserved for the management interface. See Configure Network Interfaces for On-Premises Sensors for more information.

4. Scroll to the *Port mirroring* section and set the *Mirroring mode* to **Destination**.

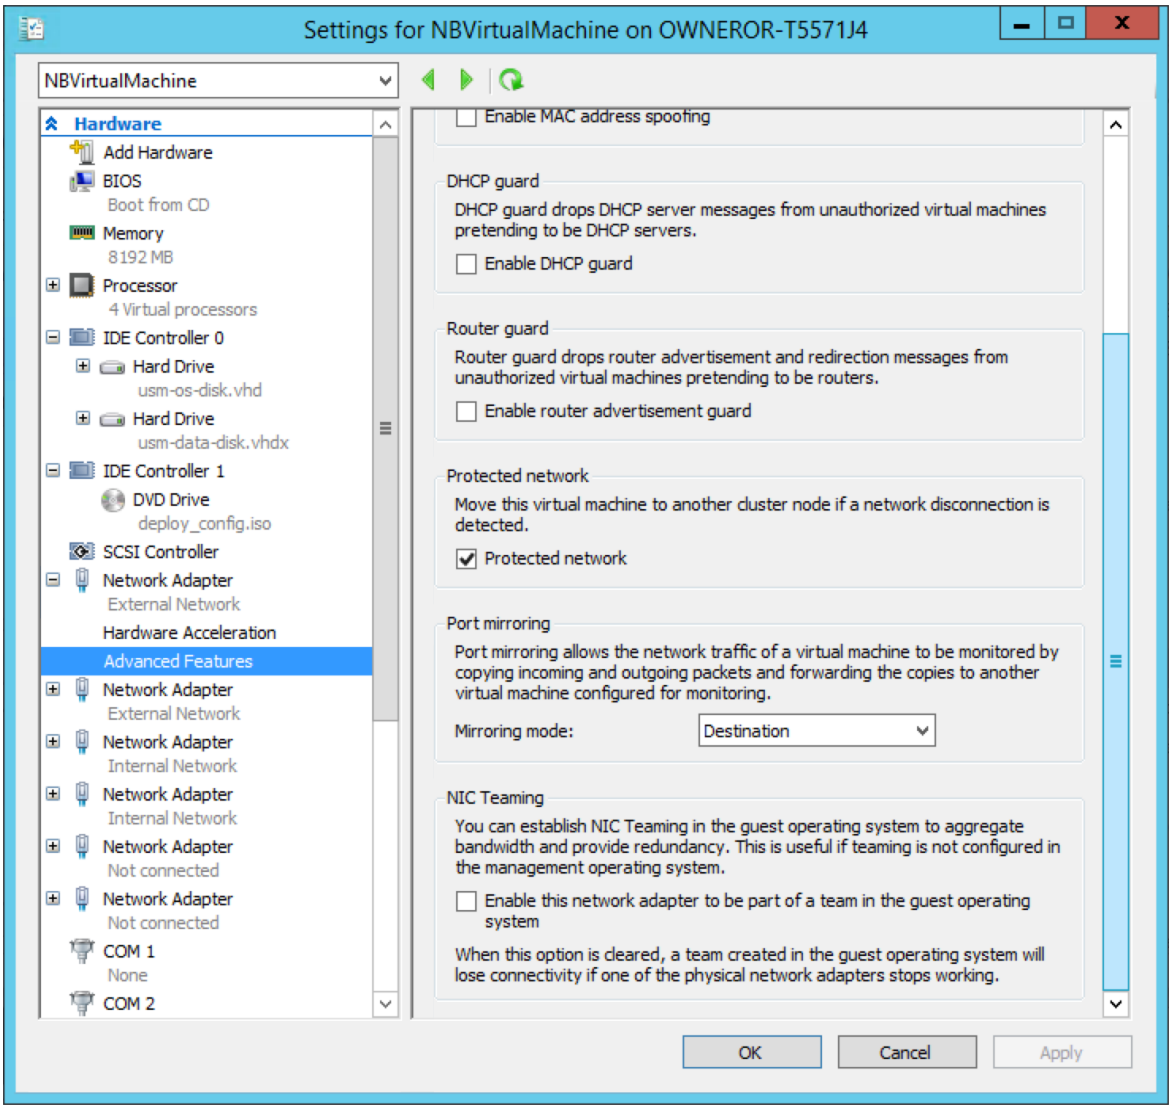

5. Click **Apply**.
6. In the same network adapter, select **Hardware Acceleration** and uncheck *Enable virtual machine queue*.

   If virtual machine queue (VMQ) is enabled on the associated network adapter, the Hyper-V Sensor is not able to detect any mirrored traffic.

7. Click **Apply** and then **OK**.

Configure the Port Mirroring Source

To complete the port mirroring configuration in Hyper-V, you need to specify the source of the mirrored traffic. In our case, this is the virtual switch connected to your spare NIC.

Note: For the virtual switch to monitor network traffic, you must enable the Microsoft Network Driver Interface Specification (NDIS) Capture extension on the switch.

To set up the virtual switch as the port mirroring source

1. Open the Microsoft Windows PowerShell console.
2. Enter the following:

```
$a = Get-VMSystemSwitchExtensionPortFeature -FeatureId 776e0ba7-94a1-41c8-
8f28-951f524251b5

$a.SettingData.MonitorMode = 2

add-VMSwitchExtensionPortFeature -ExternalPort -SwitchName <virtual_
switch_name> -VMSwitchExtensionFeature $a
```

Important: Be aware that if you enable promiscuous mode for a physical port, it directs all the traffic received on that port toward the VM destination.

Additional Configurations for Port Mirroring Setup from VLAN Traffic

If your environment uses a VLAN to route traffic, you will also need to configure Microsoft Hyper-V to accept packets from the designated VLAN identifier (ID) range.

To set up VLAN port mirroring

1. In Hyper-V Guest, create a NIC designated as "management" using the following PowerShell command:

```
Add-VMNetworkAdapter -VMName <VirtualMachineName> -Name "Management"
```

2. Add the port you will use as a mirror. For example:

```
Add-VMNetworkAdapter -Vmname <VirtualMachineName> -name "Mirror"
```

   If you have multiple NICs you are mirroring, repeat this step for each NIC.

3. Add the VLAN ID ranges to be mirrored. For example:

```
Set-VMNetworkAdapterVlan -VMName VIRTUALMACHINENAME -VMNetworkAdapterName
"mirror" -trunk -allowedvlanidlist <VLAN-ID-Range> -nativevlanid <VLAN-ID-
Range>
```

Important: The NIC needs to be created, named, and tagged with VLAN ID ranges as a guest in Microsoft Hyper-V. If the NIC is not named and tagged properly, it can create errors in the guest system.

Configure USM Anywhere to Receive ERSPAN Traffic

Encapsulated Remote Switched Port Analyzer (ERSPAN) is a traffic mirroring method that enables the mirrored traffic to be encapsulated in Generic Routing Encapsulation (GRE). USM Anywhere supports ERSPAN on its Hyper-V Sensor and VMware Sensor, and although successful testing was only done on newer Cisco devices, it should work with other modern ERSPAN device manufacturers.

To enable ERSPAN in your Hyper-V or VMware Sensor

1. Open your virtualization management console and connect to the USM Anywhere Sensor virtual machine (VM).

   Alternatively, you can open an SSH session to the sensor VM. When using an SSH session, the default username is *sysadmin*.

2. From the system menu, select **Network Configuration** and press **Enter**.
3. Select **Configure SPAN Interface** and press **Enter**.
4. Select **Enable SPAN Interface** and press **Enter**.
5. Enter the IP address for this interface.

   Note: ERSPAN must use your sensor's eth1 IP address for its interface. If your sensor's eth1 is already used by another resource, you must reconfigure that resource to use eth2 or eth3.

6. Select **OK** and press **Enter**.
7. Enter a netmask range for this configuration.

   Important: When setting up this netmask, ensure that it does not conflict with the settings for eth0 and the admin interface netmask. If both interfaces are in the same subnet, AT&T Cybersecurity recommends that you use 255.255.255.255 for this netmask instead.

8. Select **OK** and press **Enter**.
9. After you receive confirmation that the ERSPAN interface has been enabled, refer to the

Configuring ERSPAN section of the vendor website to continue the configuration.

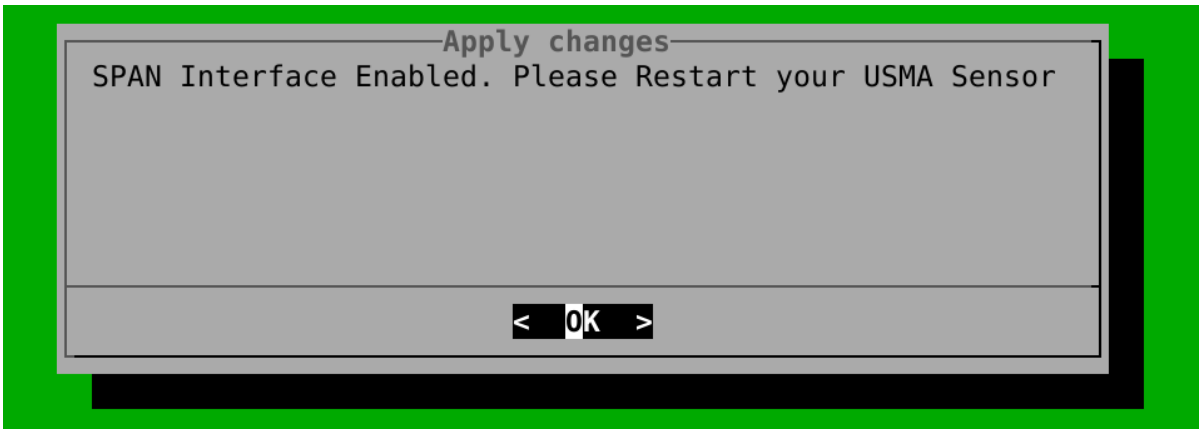

Port Mirroring Configuration on Network Devices

With a deployed on-premises USM Anywhere Sensor, you can implement Network Intrusion Detection (NIDS) by monitoring the network traffic. You can implement this by enabling promiscuous mode on the port that the Sensor network interface(s) are connected to so they can see the traffic on the networks you wish to monitor, and through the use of port mirroring. This allows USM Anywhere to perform analysis on the network traffic, which aids in the detection of threats in your environment.

By configuring a mirror port on your virtual switch or physical network device, you can clone all traffic to a single port. After configuration, the switch sends a copy of all network packets seen on one port (or an entire VLAN) to another port. The USM Anywhere Sensor immediately starts receiving events from the device through the port and begins its analysis.

Important: AT&T Cybersecurity recommends that you send packets *untagged* through the SPAN/mirror port. This is because VLAN trunking is currently not supported. Therefore, Bridge Protocol Data Units (BPDUs) or packets sent through the other Layer 2 protocols are dropped. The Layer 2 protocols include, but are not limited to, Cisco Discovery Protocol (CDP), Dynamic Trunking Protocol (DTP), Link Aggregation Control Protocol (LACP), Port Aggregation Protocol (PAgP), Spanning Tree Protocol (STP), and VLAN Trunk Protocol (VTP).

Virtual Switches

- **VMware**: This is configured by attaching one of the Sensor network interfaces to a port configured in Promiscuous mode on a Virtual Switch. See Direct Traffic from Your Physical Network to the VMware Sensor for more information.

  In addition, the upstream physical switch that the ESXi host is connected to must have Port Mirroring enabled.

- **Hyper-V**: This is configured by attaching one of the Sensor network interfaces to a port configured in Promiscuous mode on the Virtual Network. See Direct Traffic from Your Physical Network to the Hyper-V Sensor for more information.

  In addition, the upstream physical switch that the Hyper-V Server is connected to must have Port Mirroring enabled.

Physical Devices

See the following for detailed information about port mirroring on a number of third-party network devices.

Configuring the ADTRAN (AOS) Switch for Port Mirroring

Configuring the Check Point Gateway for Port Mirroring

Configuring the Cisco ASA 5505 for Port Mirroring

Configuring the Cisco Nexus 5000 Series for Port Mirroring

Configuring the Cisco SGxxx Series for Port Mirroring

Configuring the Dell Networking Force10 Switch for Port Mirroring

Configuring Dell SonicWALL Port Mirroring

Configuring the Fortinet FortiGate Switch for Port Mirroring

Note: Cisco switches support a feature known as a Switched Port Analyzer (SPAN) which enables traffic received on an interface or virtual local area network (VLAN) to be sent to a single physical port. SPAN technically implies that the source and destination ports are local to the same switch. If the traffic destination is on another remote switch, it uses Remote SPAN (RSPAN). If the destination requires crossing one or more IP networks, some switches can use Encapsulated Remote SPAN (ERSPAN).

USM Anywhere supports SPAN, RSPAN, ERSPAN, and VMware Encapsulated Remote Mirroring (L3) Source, which is an ERSPAN-like feature.

Configuring the ADTRAN (AOS) Switch for Port Mirroring

Complete one of these tasks to configure the ADTRAN Operating System (AOS) Switch for port mirroring, either through the switch CLI or the switch web UI.

Configuration Through the CLI

To configure the device through the CLI

1. Open a monitor session.
2. Specify the port to be mirrored (the source port).

```
(config)#monitor session 1 source interface ethernet 0/<source_port>
```

3. Specify the port that is going to mirror the traffic (the destination port).

```
(config)#monitor session 1 destination interface ethernet 0/<destination_
port>
```

> **Note:** There can be only one monitor session. Therefore, the only available monitor session is "1."

4. Verify the configuration of the mirrored ports.

```
Switch#show monitor session all
```

Configuration Through the Web UI

To configure the device in the web UI

1. Go to **Utilities > Port Mirroring**.
2. Choose **Destination Port** and select the proper port from this menu.
3. (Optional.) Select the **No-Tag** option so as not to tag virtual local area network (VLAN) traffic.
4. Select the **Source Port** to mirror from the menu.
5. Click **Add**.

> **Note:** If you want to add additional source ports to monitor, select another port from the **Source Port** menu and click **Add**.

6. In the command-line interface (CLI), you can verify the configuration of the mirrored

ports.

```
Switch#show monitor session all
```

To learn more about configuring port mirroring on AOS switches, refer to the Configuring Port Mirroring on AOS document on the vendor website.

Configuring the Check Point Gateway for Port Mirroring

You can configure a mirror port for a Check Point deployment that includes a Security Management Server, a gateway, and a SmartDashboard. The mirror port duplicates the network traffic and records the activity in logs.

Use these procedures to configure a Check Point Gateway Switch for port mirroring.

Connecting the Device

To configure the device

1. Open the VMware Security Gateway.
2. From the command line, run

   ```
   sysconfig
   ```

3. Select **Network Connections**.
4. Select **Configure Connections**.
5. Select the interface to configure as the mirror port.

   This is the one that you connected.

6. Select **Define as connected to a mirror port**.
7. Enable the **Application Control** blade in the SmartDashboard.

   You can also enable the IPS blade to see IPS traffic.

   > **Note:** If you only want to enable the IPS blade, you must activate at least one HTTP protection.

8. Install the Policy.

Verifying the Configuration

To verify the configuration

1. Browse to any website, such as Google.
2. Open SmartView Tracker.
3. Verify that you see traffic from the blade you enabled.

To learn more about configuring a mirror port on a Check Point gateway, refer to the Check Point documentation on the vendor website.

Configuring the Cisco ASA 5505 for Port Mirroring

The Cisco ASA 5505 Adaptive Security Appliance supports SPAN, also known as switch port monitoring, to monitor traffic that enters or exits one or more switch ports. The port where you enable SPAN (destination port) receives a copy of every packet transmitted or received on a specified source port. You can only enable SPAN for one destination port.

Note: USM Anywhere supports SPAN, RSPAN, ERSPAN, and VMware Encapsulated Remote Mirroring (L3) Source, which is an ERSPAN-like feature.

To configure the device

1. Open a monitoring session.
2. Configure the interface.

```
#interface <port>
```

3. Specify the destination port.

```
#switchport monitor<destination_port>
```

4. Specify the source port.

```
#switchport monitor<source_port>
```

To learn more about configuring port mirroring in the Cisco ASA 5505 device, refer to the Cisco ASA 5500-X Series Firewalls - Configuration Guides on the vendor website.

Configuring the Cisco Nexus 5000 Series for Port Mirroring

The Cisco Nexus 5000 Series switch supports the switched port analyzer (SPAN) feature, which allows an administrator to analyze all traffic between ports by nonintrusively directing the SPAN session traffic to a SPAN destination port that has an external analyzer attached to it. A source port, also called a monitored port, is a switched interface that you monitor for network traffic analysis. The switch supports any number of ingress source ports (up to the maximum number of available ports on the switch) and any number of source VLANs or VSANs.

Note: USM Anywhere supports SPAN, RSPAN, ERSPAN, and VMware Encapsulated Remote Mirroring (L3) Source, which is an ERSPAN-like feature.

To configure the device

1. Open a monitor session.
2. Enter global configuration mode.

```
#configure terminal
```

3. Enter interface configuration mode for the specified Ethernet interface selected by the port values.

```
#interface ethernet [port]
```

4. Set the interface to monitor mode.

```
#switchport monitor
```

Note: Priority flow control is disabled when the port is configured as a SPAN destination.

5. Revert the global configuration mode.

```
#exit
```

6. Enter monitor configuration mode.

```
#monitor session [session-number]
```

7. Configure the Ethernet destination port.

```
#destination interface ethernet [port]
```

To learn more about configuring port mirroring for the Cisco Nexus device, refer to the Configuring SPAN section of the *Cisco Nexus 5000 Series NX-OS Software Configuration Guide* on the vendor website.

Configuring the Cisco SGxxx Series for Port Mirroring

Cisco switches support a feature known as a Switched Port Analyzer (SPAN) which enables traffic received on an interface or virtual local area network (VLAN) to be sent to a single physical port. SPAN technically implies that the source and destination ports are local to the same switch. If the traffic destination is on another remote switch, it uses Remote SPAN (RSPAN). If the destination requires crossing one or more IP networks, some switches can use Encapsulated Remote SPAN (ERSPAN).

Important: USM Anywhere supports SPAN, RSPAN, ERSPAN, and VMware Encapsulated Remote Mirroring (L3) Source, which is an ERSPAN-like feature.

To configure port and VLAN mirroring

1. On the device, select **Administration > Diagnostics > Port and VLAN Mirroring**.
2. If your switch supports RSPAN, complete these steps:
   - **RSPAN VLAN**: Select **Enable** to enable RSPAN VLAN mirroring.
   - **RSPAN VLAN ID**: Select the VLAN to be mirrored.

     Note: When you configure a RSPAN mirroring session, you should select this VLAN as the RSPAN VLAN.
3. Click **Add** to add a SPAN or RSPAN mirroring session.
4. Provide the mirror session information:
   - **Session ID**: Select the identifier for the mirroring session.
   - **Session Type**: Select the appropriate option:
     - *Local Port Based*: Copies Tx, Rx, or both Tx and Rx traffic from each port to the destination port.
     - *Local VLAN Based*: Copies traffic from the local VLAN to the destination port.
     - *RSPAN Source Session*: Uses a VLAN to copy traffic from a source port or a source VLAN to another device.
     - *RSPAN Destination Session*: Uses a VLAN to copy traffic from a destination port to another device.
5. Based on the selected session type, specify the parameters for the session.

Local Port Based

- **Destination Port**: Select the analyzer port as the destination for the copied packets.

  A network analyzer, such as a PC running Wireshark, is connected to this port.

  > **Note:** Any port identified as an analyzer destination remains such until all the entries have been removed.

- **Allow Ingress Packets**: Select **Enable** to enable the destination port to receive uncopied ingress packets.
- **Source Port**: Select the source ports for the mirrored traffic and the type of traffic to be mirrored to the analyzer port:
  - *Rx Only*: Port mirroring on incoming packets.
  - *Tx Only*: Port mirroring on outgoing packets.
  - *Tx and Rx*: Port mirroring on both incoming and outgoing packets.
  - *N/A*: Traffic from this port is not mirrored.

Local VLAN Based

- **Destination Port**: Select the analyzer port to where packets are copied.
- **Allow Ingress Packets**: Select **Enable** to enable the destination port to receive ingress packets that are not copied.
- **VLAN**: Select the source VLAN from where traffic is mirrored.

RSPAN Source Session

- **RSPAN VLAN**: Select the VLAN to be used to copy traffic to another device.

  This VLAN should be the same as the VLAN defined in the RSPAN VLAN ID field.

- **Reflector Port**: Select the port or Link Aggregation Group (LAG) to be connected to another device.
- **Source Type**: Select **Port** or **VLAN** as the source port or source VLAN.

If Port is selected, set the source ports for the mirrored traffic and the type of traffic to be mirrored to the analyzer port.

- *Rx Only*: Port mirroring on incoming packets.
- *Tx Only*: Port mirroring on outgoing packets.
- *Tx and Rx*: Port mirroring on both incoming and outgoing packets.
- *N/A*: Traffic from this port is not mirrored.

If VLAN is selected, select a source VLAN.

- *VLAN*: Select a VLAN as the source VLAN.

RSPAN Destination Session

- **RSPAN VLAN**: Select the VLAN to be used to copy traffic to another device.

  This VLAN should be same as the VLAN defined in the **RSPAN VLAN ID** field.

- **Destination Port**: Select the analyzer port as the destination for the copied packets.
- **Allow Ingress Packets**: Select **Enable** to enable the destination port to receive ingress packets that are not copied.

6. Click **Apply**.

   This updates the running configuration.

See SG220-50P Switch documentation on the vendor website to learn more about configuring port mirroring on the Cisco SGxxx Series devices.

Configuring the Dell Networking Force10 Switch for Port Mirroring

The Dell Networking Force10 Switches support port monitoring on both physical and logical interfaces, such as a virtual local area network (VLAN) and port channel. The monitored (the source) and monitoring ports (the destination) must be on the same switch.

To configure the device

1. Enter configuration mode:

```
#configure
```

2. Enter the destination port to use for the monitoring session, and confirm that it has no configuration:

```
#interface te 0/2
```

3. Remove any IP addresses that may have previously been configured:

```
#no ip address
```

4. Enable the port:

```
#no shutdown
```

5. Exit the destination port interface:

```
#exit
```

6. Set up and identify the session number (range is from 0 - 65535):

```
#monitor session 0
```

7. Configure the source, the port you want to monitor, the destination port you want to send the monitored packets to, and the direction (`both`, `Rx`, or `Tx`):

```
#source te 0/1 destination te 0/2 direction both
```

8. Verify that port monitoring is active:

```
#show monitor session 0
```

See the Knowledge Base article on the vendor website to learn more about configuring port mirroring on the Dell Networking Force 10 Switch.

Configuring Dell SonicWALL Port Mirroring

You can configure port mirroring on the SonicWALL NSA 2400MX to send a copy of network packets seen on one or more switch ports (or on a virtual local area network [VLAN]) to another switch port, called the mirror port. By connecting to the mirror port, you can monitor the traffic passing through the mirrored ports.

Note: A VLAN trunk port can be mirrored, but cannot act as a mirror port itself.

To create a new port mirroring group

1. Go to **Switching > Port Mirroring**.
2. Click **New Group**.
3. In the Edit Mirror Group dialog box, enter a descriptive name for the group into the Interface Group Name field.

4. For **Direction**, select one of the following:
   - **ingress** — Monitors traffic arriving on the mirrored ports.
   - **egress** — Monitors traffic being sent from the mirrored ports.
   - **both** — Monitors traffic in both directions on the mirrored ports.
5. In the All Interfaces list, select the port to use to mirror the traffic, then click the upper right-arrow button to move it to the Mirror Port field.

   You must use an unassigned port as the mirror port.
6. In the All Interfaces list, select one or more ports to monitor, and click the lower right-arrow button to move them to the Mirrored Ports field.

   You will be able to monitor traffic on the mirrored ports by connecting to the mirror port.
7. Select the **Enable** checkbox to enable port mirroring for these ports.
8. Click **OK**.

See the Knowledge Base article on the vendor website to learn more about configuring port mirroring on SonicWall devices.

Configuring the Fortinet FortiGate Switch for Port Mirroring

This procedure explains how to configure Fortinet FortiGate switches for port mirroring on models with built-in hardware switches (for example, the FortiGate-100D, 140D, and 200D), using the Switch Port Analyzer (SPAN) feature.

Configuration Through the CLI

To configure SPAN through the CLI

1. Enter the following:

```
config system virtual-switch
edit <port>
set span enable
set span-source-port <port>
set span-dest-port <port>
set span-direction {both | Tx | Rx}
end
end
```

Configuration Through the Web UI

To configure SPAN through the web UI

1. Go to **System > Network > Interfaces**.
2. Edit a hardware switch interface.

   By default, the system may have a hardware switch interface called a LAN. You can also create a new hardware switch interface.

   a. Select the SPAN checkbox, then select a source port from which you want traffic mirrored.

   b. Select one of the following:

      - Traffic received
      - Traffic sent
      - Both

See the Knowledge Base article on the vendor website to learn more about configuring port mirroring on Fortinet-FortiGate Switches.

Granting Access to Active Directory for USM Anywhere

If you want to run Active Directory (AD) scans in USM Anywhere, you need to configure your AD server assets to grant access to the USM Anywhere Sensor. You also need to configure credentials in USM Anywhere to make an authenticated connection.

This process contains three tasks:

- Create a dedicated administrator account in AD on all the hosts you want to scan. This is used by USM Anywhere to log into that host system to perform a scan.
- Activate Windows Remote Management (WinRM) in the domain controller and in all the hosts you want to scan.
- Apply the AD account credentials for those assets in USM Anywhere.

Note: See Microsoft's guide on authentication for remote connections for more information on Microsoft Windows authentication permissions.

Create a Dedicated AD Account

When configuring your VMware Sensor, Hyper-V Sensor, or Azure Sensor, you can define AD credentials that USM Anywhere uses to perform an AD scan through the sensor. These are the credentials that you define in the Credentials page and assign to the asset to support a scheduled Active Directory scan job. It is a best practice to use a dedicated account for this purpose.

To create a new dedicated account in AD

1. Log in to your domain controller administrator account.
2. Open **Active Directory Users and Computers**.
3. Create a new user called either `alienvault_usm_anywhere` or any other name that's easy to associate with USM Anywhere.
4. Add the user you've just created to the Domain Admins group.

Activate WinRM to Enable Windows PowerShell Remoting

For Microsoft Windows systems, USM Anywhere uses the WinRM framework to execute the corresponding commands. Therefore, if WinRM is unavailable on a target Windows system through the account credentials, USM Anywhere won't be able to connect. You must satisfy the following requirements:

- WinRM version 2.0 or later.
- PowerShell version 5.1 or later. The Active Directory Scanner runs a PowerShell command through WinRM, which requires PowerShell 5.1 or later to be installed on your machine.

To activate WinRM, you can use a group policy to combine the domain controller and all the hosts in your AD. (For reference, see this How to enable PowerShell Remoting via Group Policy article.)

Alternatively, if you prefer to activate WinRM manually in each system you want to scan, use this procedure to activate a Windows RM listener on port 5985.

To start the WinRM service

1. Open the Windows Command Prompt using administrator privledges and run the command `winrm qc`.

   Important: Only the members of the Remote Management Users and Administrators groups can log in through WS-Management.

2. Accept the default settings.

   The command starts the WinRM service and configures a listener for the port 5985.

3. Create a firewall rule to allow incoming connections to port 5985.

For more information about WinRM, you can refer to these Microsoft articles:

- Installation and configuration for Windows Remote Management
- WinRM (Windows Remote Management) Troubleshooting

Manage Credentials for Your AD Servers

Before you run an AD scan from USM Anywhere, you should make sure that each of the assets has assigned credentials that are able to connect to the system. In USM Anywhere, you can assign credentials for an individual asset or for an asset group. See Creating Credentials on how to create credentials and Assigning Credentials to Assets on how to assign them to assets.

Note: Credentials assigned directly to an asset have higher priority than those assigned to an asset group.

When USM Anywhere runs a scan or executes a system-level action, it uses the credential set assigned directly to the asset, if there is one. If those credentials don't connect or the asset doesn't have an assigned credential set, it uses the credential set assigned to the group where the asset is a member, if that asset is a member of an asset group.

Proxy Configuration on the USM Anywhere Sensor

USM Anywhere enables you to configure a proxy server on the USM Anywhere Sensor.

Important: With this proxy configuration, you will still need to open a direct connection to your USM Anywhere instance via TCP port 7100. This connection is required to forward all events and data, and to distribute configuration changes to and from your sensor. See the deployment requirements for your sensor to read more about required port configurations.

Warning: Even if a proxy server is configured, some scripts corresponding with netcat and traceroute tools will direct through port 443.

To configure a proxy server

1. Open your virtualization management console and connect to the USM Anywhere Sensor virtual machine (VM).

   Alternatively, you can open an SSH session to the sensor VM. When using an SSH session, the default username is *sysadmin*.

   Important: If you are accessing a Microsoft Azure Sensor through SSH and you specified a username other than the default (*sysadmin*) for your SSH access, you must use the following commands at the command line to "sudo up" and access the sensor console:

   ```
   # sudo su - sysadmin
   ```

2. From the USM Anywhere Sensor console System Menu, select **Network Configuration**, and then press **Enter**.

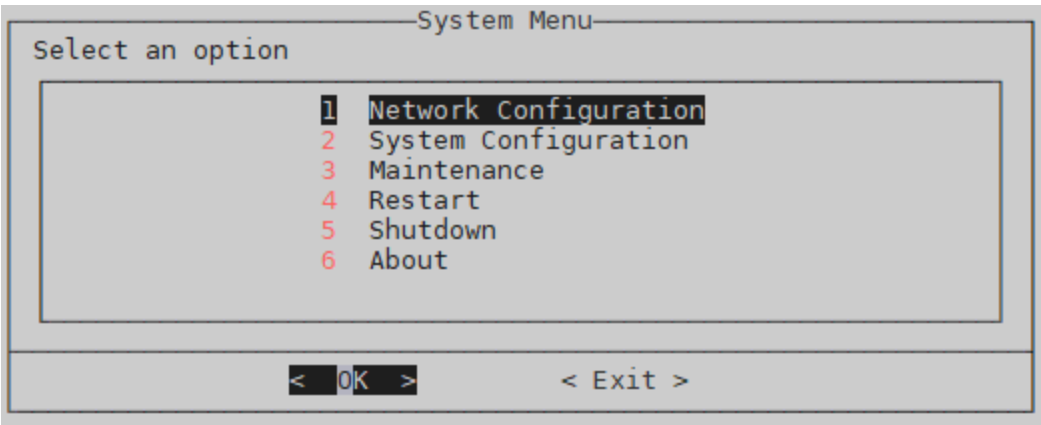

3. From the Network Configuration menu, select **Configure Proxy**, and then press **Enter**.

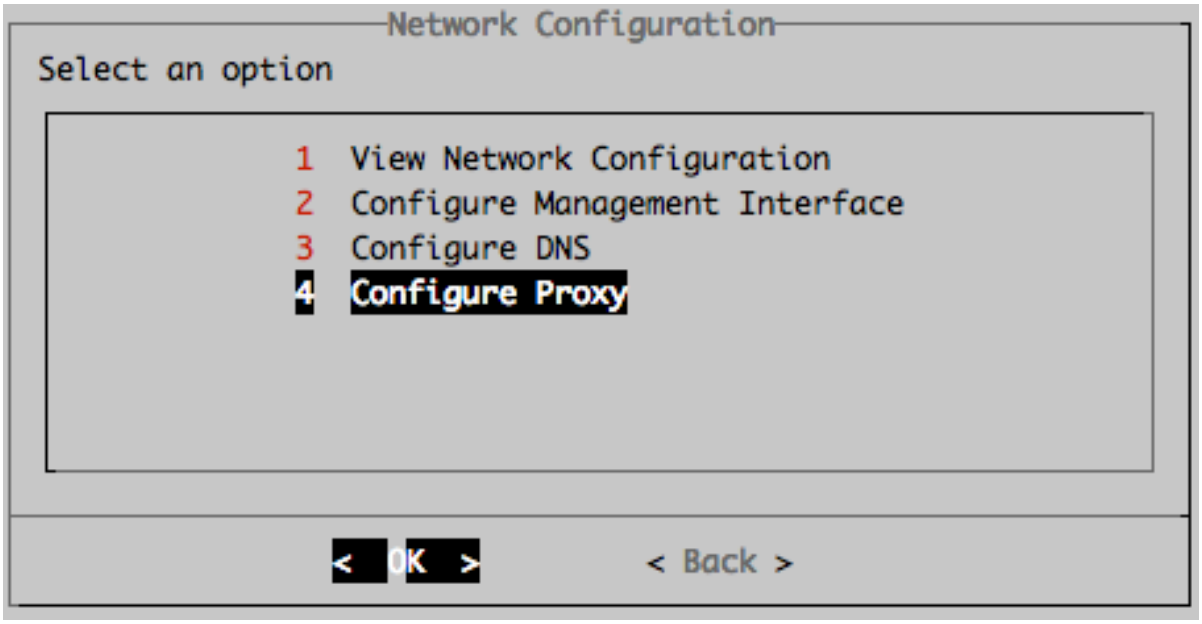

4. From the Configure Proxy menu, select **Set Proxy IP and Port**, and then press **Enter**.

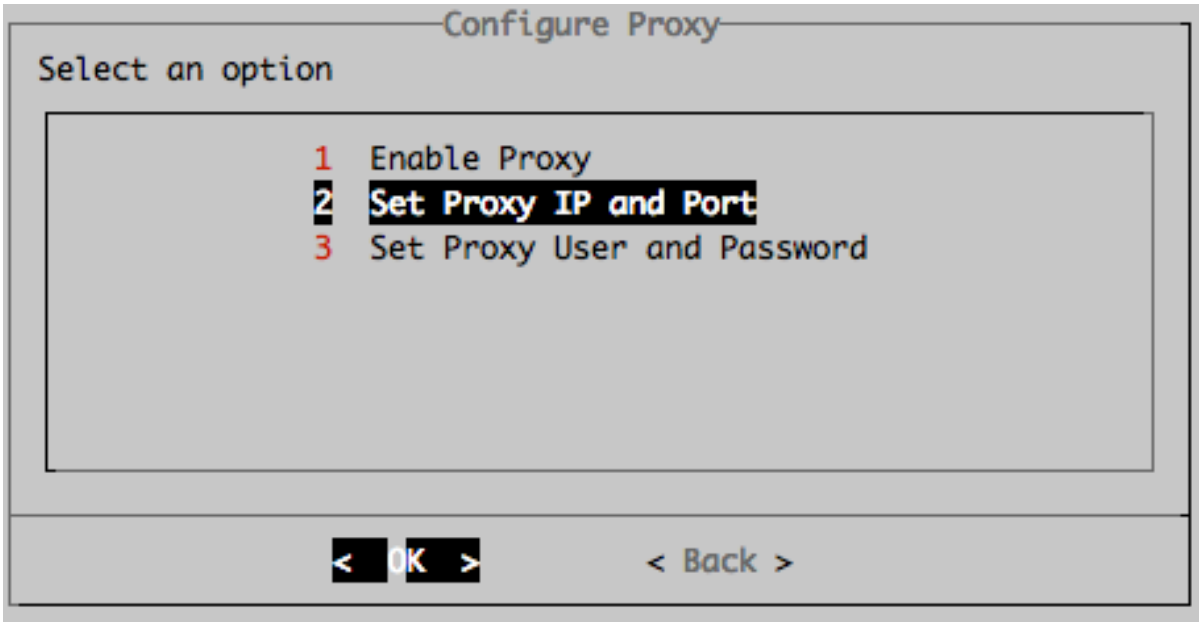

5. Enter the proxy port, and then press **Enter**.

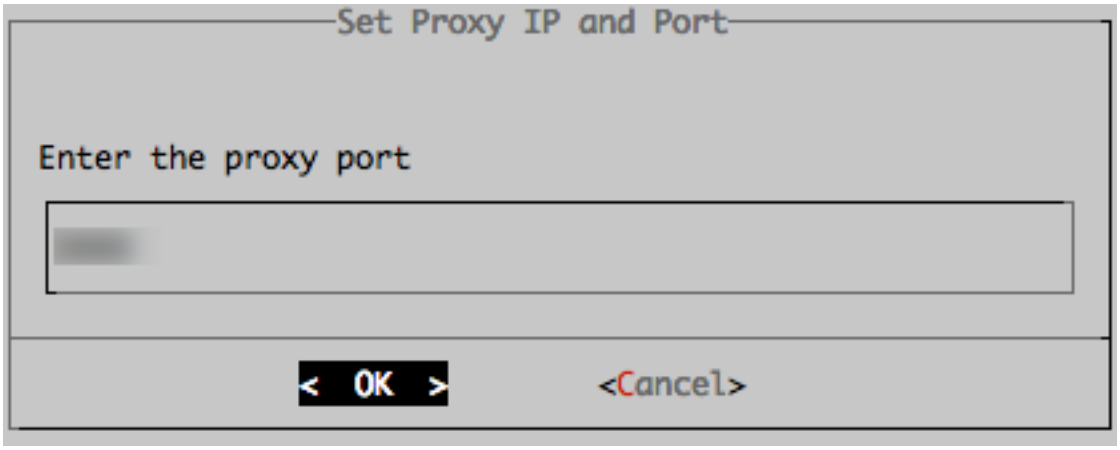

6. Enter the proxy hostname or IP address, and then press **Enter**.

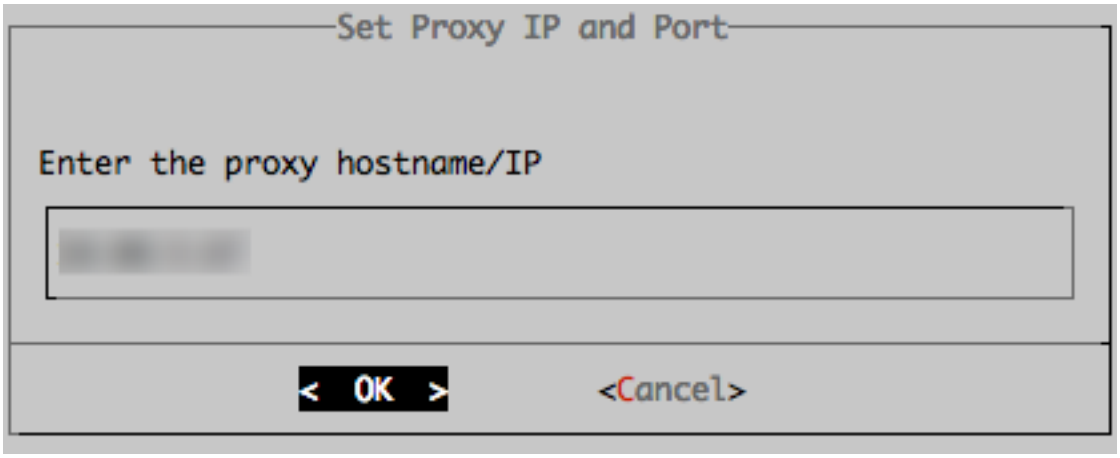

7. (Optional.) Set a proxy user and password by selecting **Set Proxy User and Password**, and then pressing **Enter** from the Configure Proxy menu.

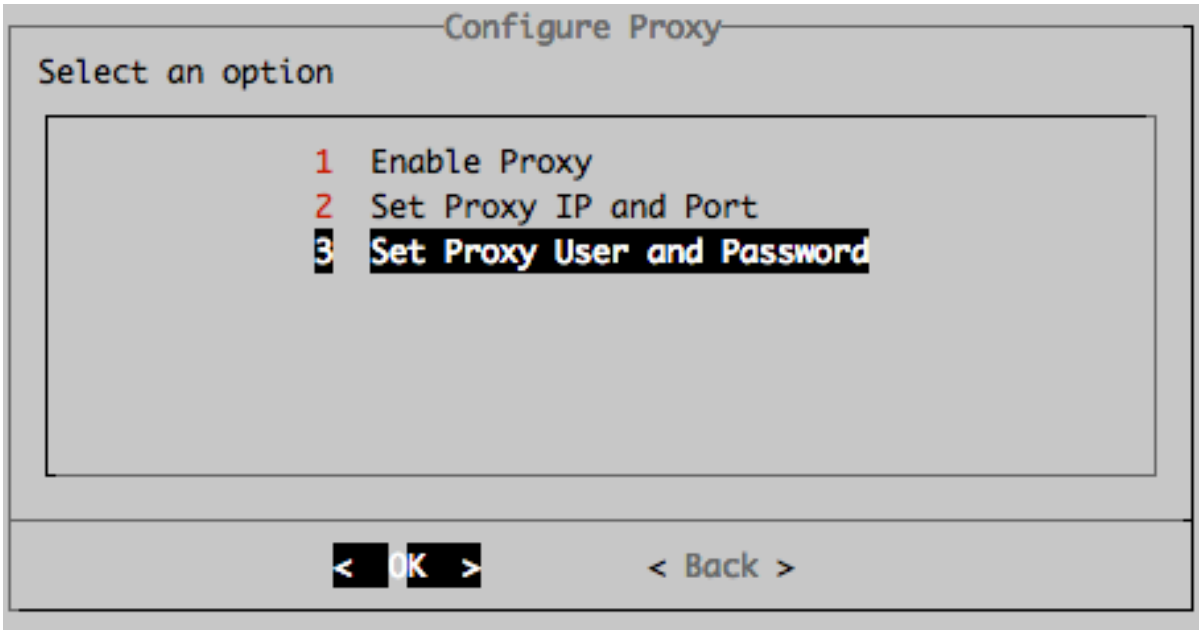

8. Enter the proxy user, and then press **Enter**.

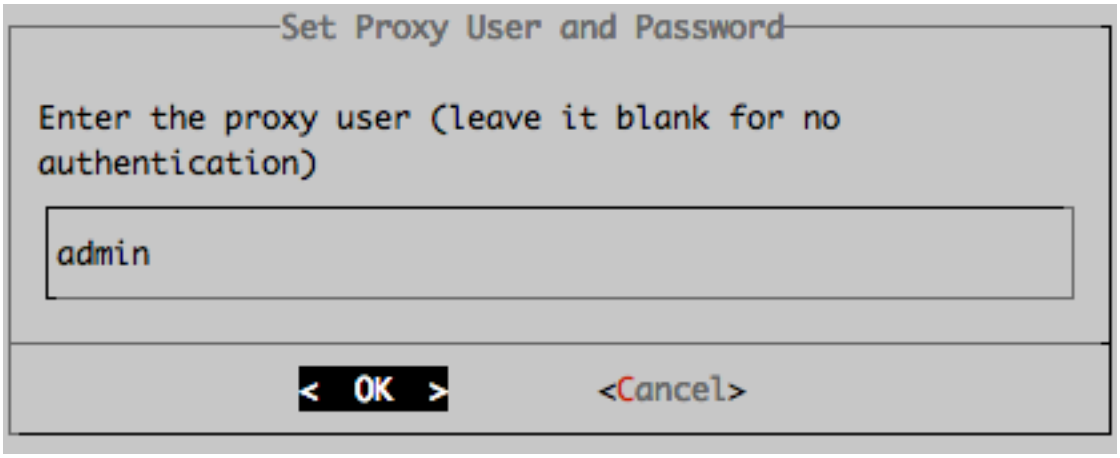

9. Enter the proxy password, and then press **Enter**.

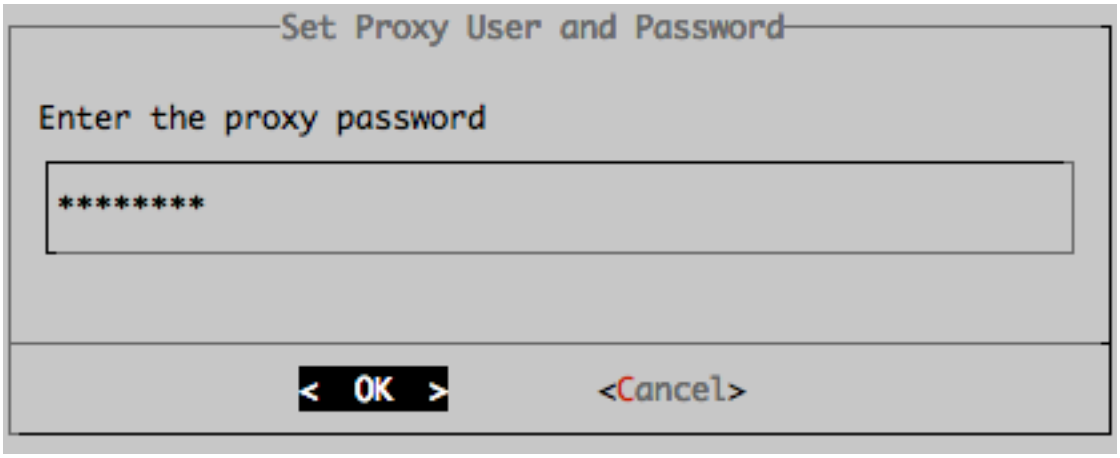

10. From the Configure Proxy menu, select **Enable Proxy**, and then press **Enter**.

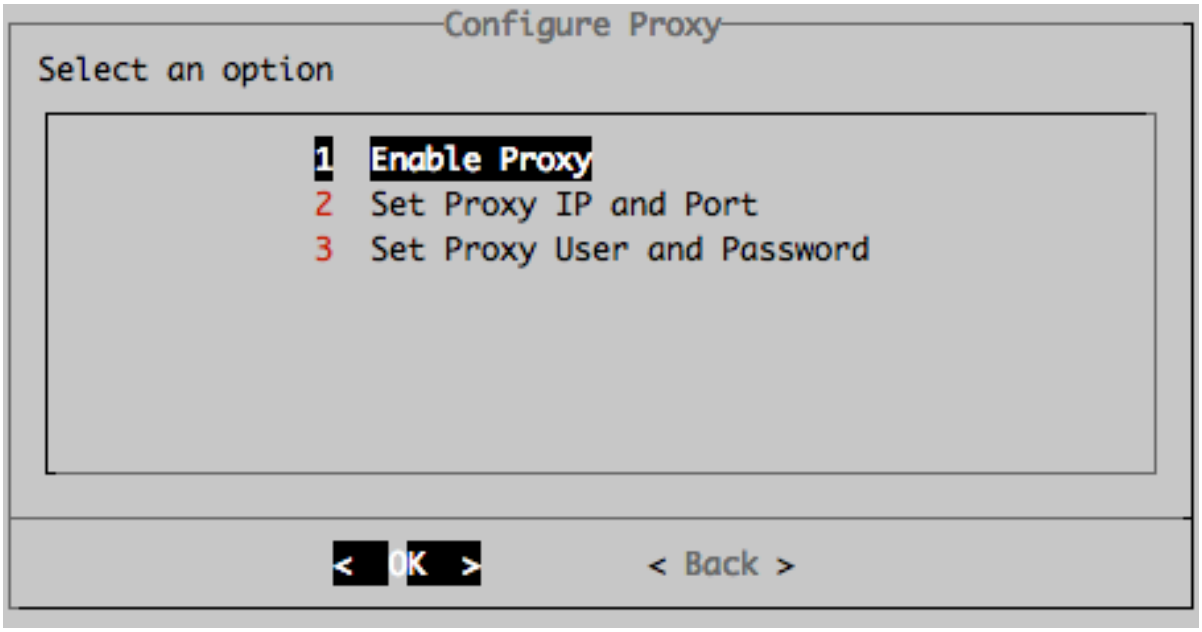

11. From the Network Configuration menu, you can check your proxy configuration by selecting **View Network Configuration**, and then pressing **Enter**.

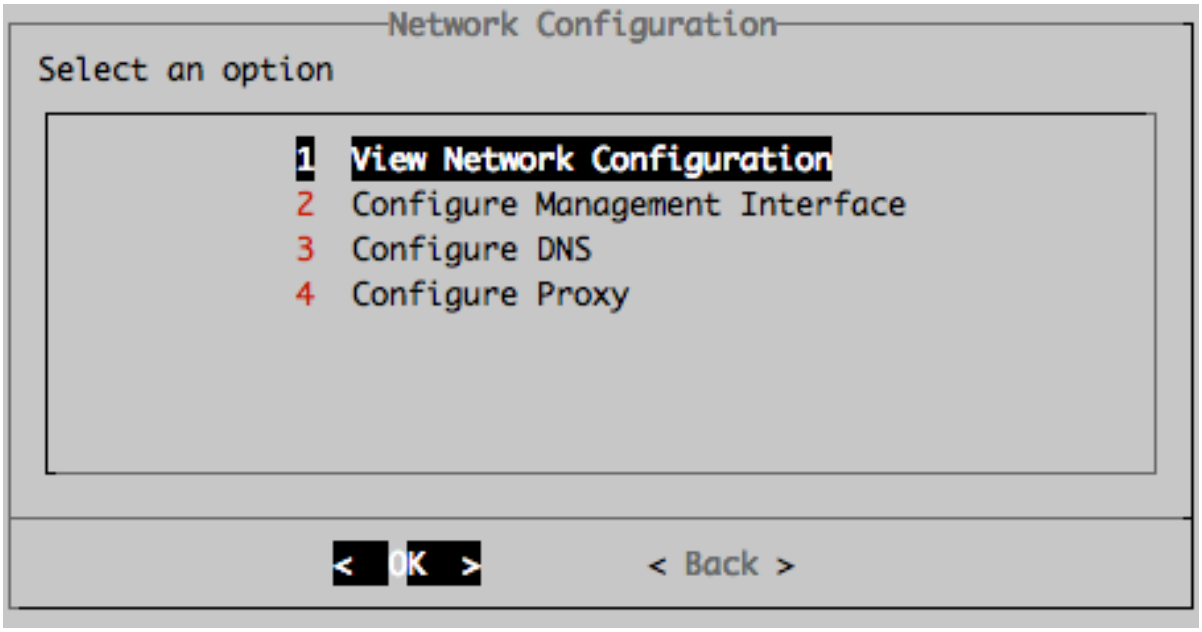

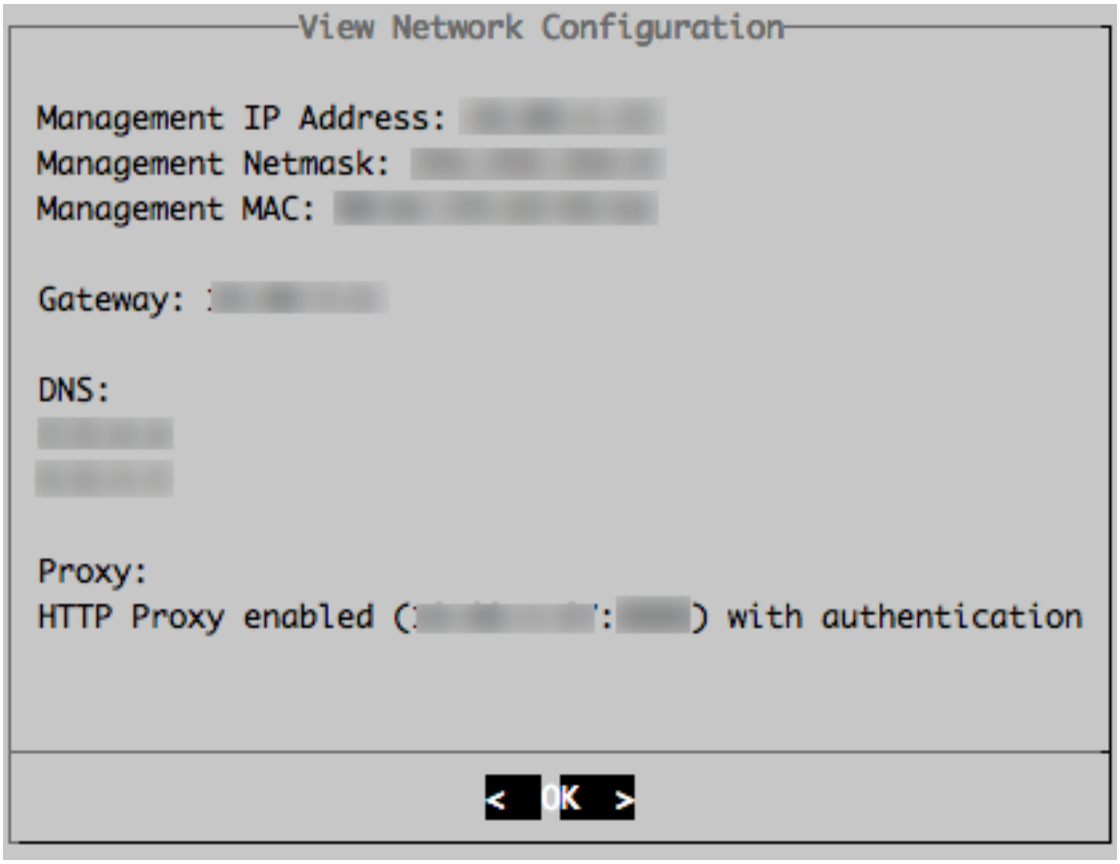

Data Sources and Log Processing

Some data sources, such as those that support the syslog protocol, can send their logs directly to the USM Anywhere Sensor. For other data sources, USM Anywhere retrieves the logs through scheduled log collection jobs, queries through registered AlienVault Agents, and queries through configured AlienApp integrations. In each of these cases, USM Anywhere uses an AlienApp for normalizing the collected data to extract and store information in common data fields that define an event.

USM Anywhere Sensors securely transfer the event data from your network environment to your single-tenant USM Anywhere instance for centralized collection, security analysis, threat detection, and compliance-ready log management. Installed AlienVault Agents communicate over an encrypted channel to send data directly to USM Anywhere.

This section provides detailed information about collecting data from your devices, systems, and applications.

Data Sources and Log Collection 342

File Integrity Monitoring 394

Scheduling Active Directory Scans from the Job Scheduler Page 401

Data Sources and Log Collection

Some data sources, such as those that support the syslog protocol, can send their logs directly to the USM Anywhere Sensor. For other data sources, USM Anywhere retrieves the logs through scheduled log collection jobs, queries through registered AlienVault Agents, and queries through configured AlienApp integrations. In each of these cases, USM Anywhere uses an AlienApp for normalizing the collected data to extract and store information in common data fields that define an event.

USM Anywhere Sensors securely transfer the event data from your network environment to your single-tenant USM Anywhere instance for centralized collection, security analysis, threat detection, and compliance-ready log management. Installed AlienVault Agents communicate over an encrypted channel to send data directly to USM Anywhere.

You configure your third-party devices, systems, and applications to transmit generated log data to your USM Anywhere Sensor, to a location that the sensor can query, or directly to USM Anywhere from a registered AlienVault Agent. Your data sources can produce the data using various formats that are compatible with AlienApps, see AlienApps Supported Log Formats for more information.

Data Collection by Sensor Apps

When log data is transmitted directly to a USM Anywhere Sensor, a Sensor App collects this data according to the identified log message protocol. The following table shows the data collection by sensor apps.

Data Collection by Sensor Apps

Sensor App	Functional support
Syslog Server	Passively collects syslog data transmitted to the USM Anywhere Sensor. For more information, see The Syslog Server Sensor App. The Syslog Server app is supported on all USM Anywhere Sensor types.
Graylog (GELF)	Passively collects GELF data transmitted to the USM Anywhere Sensor. For more information, see The Graylog (GELF) Sensor App. The Graylog app is supported on all USM Anywhere Sensor types.

Data Collection by Sensor Apps (Continued)

Sensor App	Functional support
Amazon Web Services	Collects data from AWS logging services and performs queries to collect log data stored in an S3 repository within your AWS environment. For more information about built-in support for AWS logs, see AWS Log Discovery and Collection in USM Anywhere. The AWS app is supported only on the AWS Sensor.
Azure	Collects data from Azure logging services configured within your Azure environment. For more information about built-in support for Azure logs, see Azure Log Discovery and Collection in USM Anywhere. The Azure app is supported only on the Azure Sensor.

Host-Based Log Collection

USM Anywhere provides the AlienVault Agent, which you can install on your endpoints to centralize the collection and analysis of event logs from remote servers and desktops, making it easier to track the health and security of these systems. It also supports host-based log collection through manual installation and configuration of NXLog and osquery.

Note: With the addition of the AlienVault Agent, USM Anywhere provides an easier implementation of HIDS, FIM, and endpoint log collection across your Windows and Linux environments in the cloud and on premises. If you already have NXLog or osquery installed and configured on your endpoints to forward events to a USM Anywhere Sensor, these methods are still supported and you do not need to replace them.

Refer to the following topics for detailed information about sending log data from your host systems:

- Log collection from a **Linux System** — Collecting Linux System Logs
- Log collection from a **Windows System** — Collecting Windows System Logs

Log Collection by Advanced AlienApps

Advanced AlienApps use API and system integrations to actively collect data directly from a third-party device or service. The *USM Anywhere AlienApps Guide* provides detailed information about this method of data collection.

Log Collection from Various Third-Party Devices and Systems

To support the wide array of third-party devices and systems you may have in your environments, AT&T Cybersecurity provides instructions in the AlienApps UI to assist you with configuration of the most commonly-used external data sources to send log data to a USM Anywhere Sensor.

Syslog Parsing

It is important for the date and time listed in the header of the syslog files to be formatted correctly from the data source for USM Anywhere to properly parse the information when generating event details. Some formats for date and time, such as the ISO format, may create conflicts in the way event information is parsed. Instead, it is recommended you follow the practice of using the IETF BSD specifications for syslog formatting, resulting in the following timestamp format in the syslog headers: `Mmm dd hh:mm:ss`. Per the BSD protocol, the header should contain a TIMESTAMP field and HOSTNAME field, and the MSG portion of the log should contain a TAG field and a CONTENT field.

Note that the use of an intermediary log collection agent can cause parsing errors by adding extra, unformatted context to the syslog messages.

The Syslog Server Sensor App

Role Availability	Read-Only	Investigator	Analyst	Manager
	✗	✗	✗	✓

Syslog is a message-logging standard supported by most devices and operating systems. USM Anywhere uses Syslog-ng, which supports IETF-syslog protocol, as described in RFC 5424 and RFC 5426; and BSD-syslog-formatted messages, as described in RFC 3164. While RFC 5424 and RFC 3164 define the format and rules for each data element within the syslog header, there can be a great deal of variance in the message content received from your data sources. Although Syslog-ng fixes some missing or incorrect headers, USM Anywhere doesn’t support syslog-formatted messages other than the ones previously mentioned.

> **Note:** You can send syslog messages to USM Anywhere directly from the data source or use log-forwarding software such as Splunk or Loggly. USM Anywhere accepts most log-forwarding software that doesn't alter the raw log messages.

The USM Anywhere Sensors use the syslog server app to collect syslog messages for processing. The USM Anywhere Sensor passively listens on the syslog ports.

The following tables list the ports that require the syslog server app for the RFC 3164 and RFC 5424 protocols.

Ports the Syslog Server App Requires for Specific Protocols (RFC 3164)

Protocol	Port	BSD – Syslog Protocol Support
UDP	514	USM Anywhere collects data through syslog over UDP on port 514 by default.
TCP	601	USM Anywhere collects data through syslog over TCP on port 601 by default.
TLS/TCP	6514	USM Anywhere collects Transport Layer Security (TLS)-encrypted data through syslog over TCP on port 6514 by default. **Important:** USM Anywhere requires the use of the TLS 1.2 protocol to ensure security.

Ports the Syslog Server App Requires for Specific Protocols (RFC 5424)

Protocol	Port	IETF – Syslog Protocol Support
TCP	602	USM Anywhere collects data through syslog over TCP on port 602 by default.
TLS	6515	USM Anywhere collects data through syslog over TLS on port 6515 by default. **Important:** USM Anywhere requires the use of the TLS 1.2 protocol to ensure security.

Important: Make sure that the required ports are open for these protocols within your security groups and firewalls.

Configure Syslog on Your Data Sources

For each of the data sources in your network where you want to collect syslog data, you must forward the logs to a USM Anywhere Sensor. Use the following configuration information to use rsyslog to collect and send syslog to your USM Anywhere Sensor. Many third-party systems and devices support other methods for sending syslog messages. Go to the specific AlienApp in USM Anywhere for instructions about syslog forwarding.

Note: The `*.*` configuration enables you to forward all syslog messages. However, AT&T Cybersecurity strongly recommends that you use any of the rsyslog filtering capabilities to forward only the logs that need to be monitored by USM Anywhere.

Important: You have to use the following syntax in the */etc/rsyslog.conf* with older version of rsyslog:

```
*.* @remote_server:port
```

Standard Syslog over UDP

To configure syslog over UDP, you need to configure rsyslog on your data source to forward the logs to your USM Anywhere Sensor over the UDP port (the default is 514).

```
*.* action(type="omfwd" target="<IP>" port="514" protocol="udp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000") #
send (all) messages - Forward to the USM Anywhere Sensor IP address
```

Where `<IP>` is the IP address for the USM Anywhere Sensor.

Standard Syslog over TCP

To configure syslog over TCP, you need to configure rsyslog on your data source to forward the logs to your USM Anywhere Sensor over the TCP port (default 601).

```
*.* action(type="omfwd" target="<IP>" port="601" protocol="tcp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000") #
send (all) messages - Forward to the USM Anywhere Sensor IP address
```

Where `<IP>` is the IP address for the USM Anywhere Sensor.

TLS-Encrypted Syslog over TCP

If you want to enable encrypted syslog communications between a host and the USM Anywhere Sensor to comply with your organization's security policies that require encryption

of log data in transit, you can configure syslog TLS/TCP forwarding. TLS uses certificates to encrypt the communication between a client (the data source) and server (the USM Anywhere Sensor).

To configure syslog for TLS over TCP, you need to configure rsyslog on your data source to use TLS encryption and forward the logs to your USM Anywhere Sensor over the default port (6514 or 6515). The following configuration information is tested on Ubuntu 16.04 using rsyslog 8. For Red Hat Linux distributions, use rpm or yum in place of apt-get. For other systems supporting rsyslog TLS configuration, you can extrapolate from this information.

Note: For devices such as Trend Micro and Palo Alto Networks, AT&T Cybersecurity requires you to upload your own certificates to both the device and the USM Anywhere Sensor. See Upload Your Own Certificate for more information.

Install the rsyslog-gnutls Package

The network stream driver implements a TLS-protected transport through the GnuTLS library.

```
sudo apt-get install rsyslog-gnutls
```

You can download the library from the GnuTLS site: https://gnutls.org/download.html.

Download and Install the Certificate

USM Anywhere provides a digital certificate that enables the SSL operation and the crypto keys used to secure the connection. This certificate is automatically renewed or updated with no maintenance requirement on your part.

Note: AT&T Cybersecurity recommends you follow these steps to configure remote forwarders.

To download and install the certificate

1. In USM Anywhere, go to **Data Sources > Sensors**.
2. Click the **Sensor Apps** tab.
3. In the left navigation, select **Syslog Server**.
4. Select the sensor where you want to forward the logs.

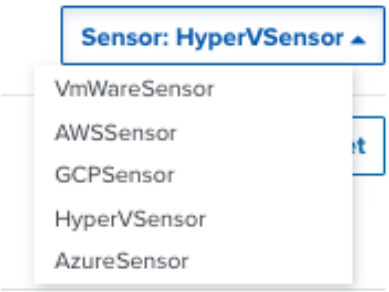

5. Download the certificate by clicking the **Download AlienVault CA Certificate** link.

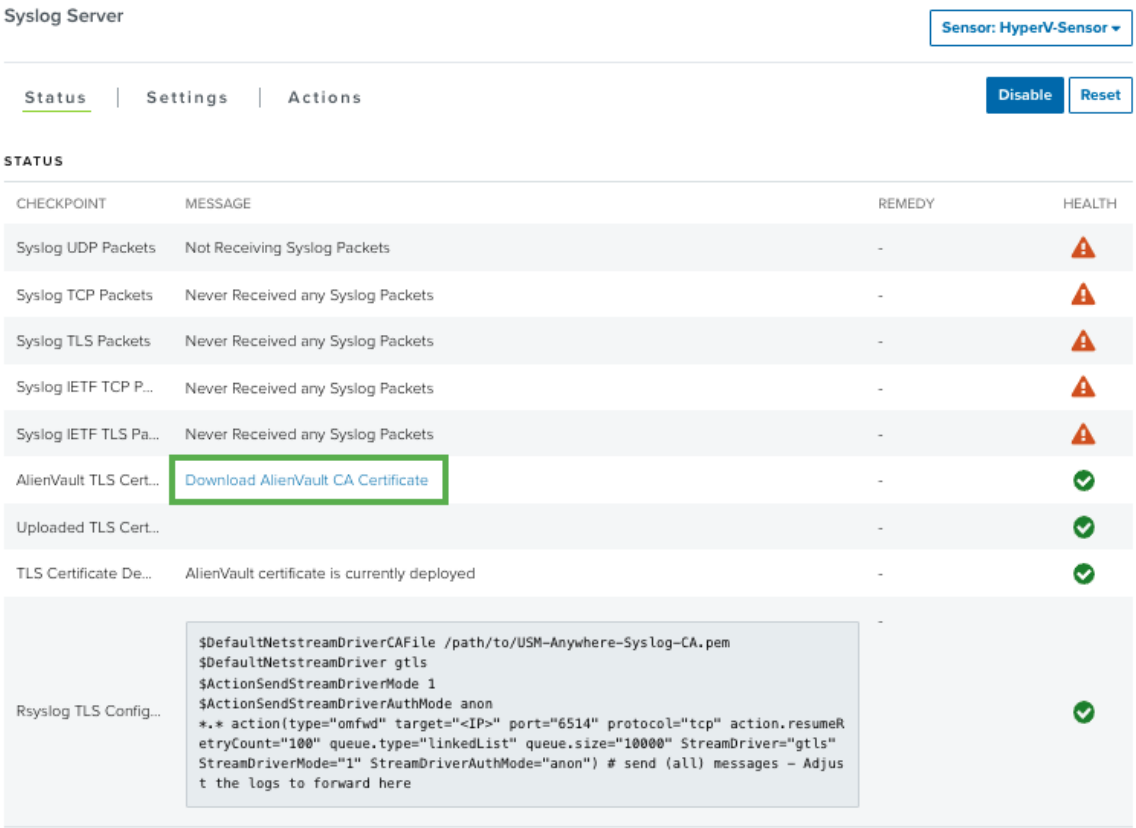

6. Copy the certificate file to the client system.
7. Configure the rsyslog service. See Configure the rsyslog Service for more information.

Upload Your Own Certificate

You can upload your own TLS certificate and deploy it on the syslog server app through port 6514 or 6515. If you choose to generate and upload your own TLS certificate, you must provide the Certificate Authority (CA) certificate, the Server Certificate (signed by the root certificate

authority, not an intermediate certificate), and the Server Private Key. USM Anywhere supports intermediate certificates for both self-signed or not self-signed.

To upload your own TLS certificate

1. In USM Anywhere, go to **Data Sources > Sensors**.
2. Click the **Sensor Apps** tab.
3. In the left navigation, select **Syslog Server**.
4. Under the Syslog Server header, click the **Settings** tab.
5. Paste the raw text of your PEM files to each of the fields individually.
   - Both the CA Certificate and Server Certificate must contain the PEM delimited `-----BEGIN CERTIFICATE-----` header and `-----END CERTIFICATE-----` footer.
   - The Server Private Key must be formatted in Public-Key Cryptography Standards (PKCS) #8 standard syntax and contain the PEM delimited `-----BEGIN PRIVATE KEY-----` header and `-----END PRIVATE KEY-----` footer.
   - If you want to use intermediate certificates, paste this text:

     ```
     -----BEGIN CERTIFICATE-----
     (Your Intermediate2 certificate if exists)
     -----END CERTIFICATE-----
     -----BEGIN CERTIFICATE-----
     (Your Intermediate1 certificate if exists)
     -----END CERTIFICATE-----
     -----BEGIN CERTIFICATE-----
     (Your Root certificate)
     -----END CERTIFICATE-----
     ```

   - If you don't want to use intermediate certificates, paste this text:

     ```
     -----BEGIN CERTIFICATE-----
     (Your Root certificate)
     ```

```
-----END CERTIFICATE-----
```

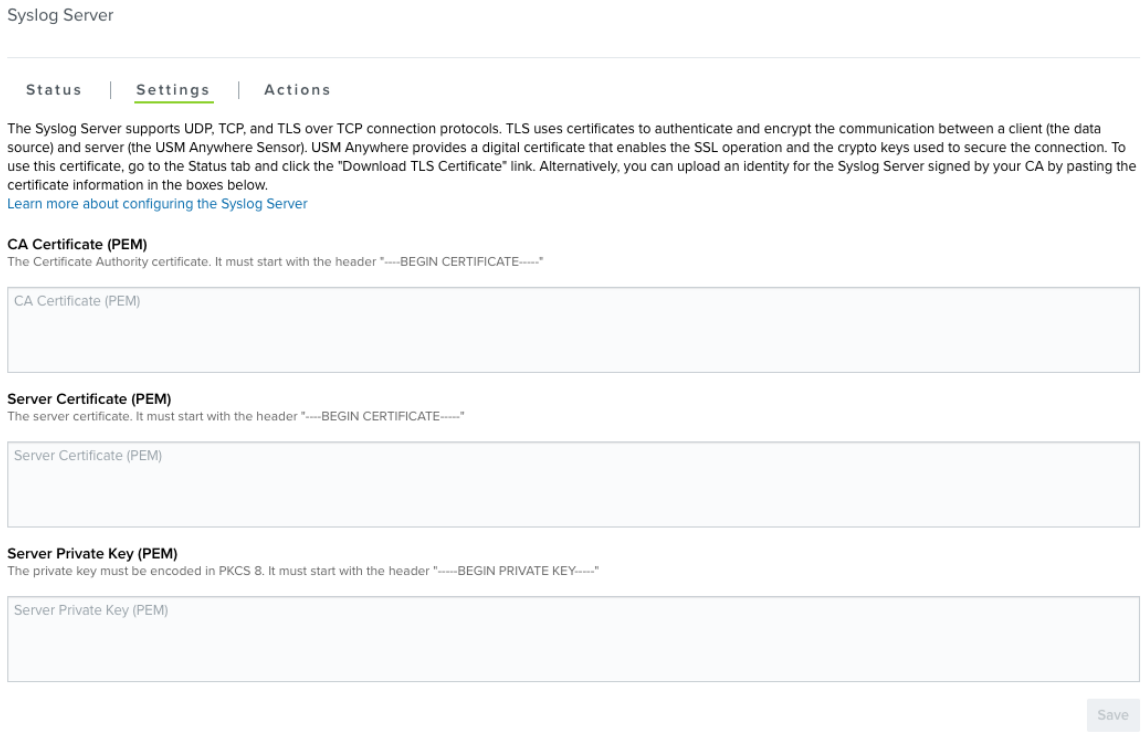

Warning: The full chain of certificates must be contained if you want to upload your own TLS certificate. This includes the root certificate.

6. Click **Save** to deploy your custom certificates.

After you have uploaded your certificates, you can confirm that they have been saved by clicking the **Status** tab of the Syslog Server page. The Status for Uploaded TLS Certificate will now display the message *Uploaded Successfully*, and the Health displays the green checkmark icon.

Thereafter, if you want to use the default TLS certificates instead, you need to remove the custom certificates first.

To remove your custom certificates and use the default TLS certificates

1. In USM Anywhere, go to **Data Sources > Sensors**.
2. Click the **Sensor Apps** tab.
3. In the left navigation, select **Syslog Server**.
4. Under the Syslog Server header, click the **Settings** tab.

5. Delete the text in the PEM certificate fields and click **Save**.
6. Under the Syslog Server header, click the **Actions** tab.
7. In the Regenerate AlienVault TLS Certificates action row, click **Run**.

   The Select Action window displays the Action Type as Syslog Server, and the App Action as Regenerate AlienVault TLS Certificates.

8. Click the **Run** button to regenerate the TLS certificates.

Configure the rsyslog Service

Edit the rsyslog configuration file to send logs to the USM Anywhere Sensor. The configuration file is located at /etc/rsyslog.conf by default.

```
$DefaultNetstreamDriverCAFile /PATH/TO/CERTIFICATE
$DefaultNetstreamDriver gtls
$ActionSendStreamDriverMode 1
$ActionSendStreamDriverAuthMode anon
*.* action(type="omfwd" target="<IP>" port="6514" protocol="tcp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000") #
send (all) messages - Forward to the USM Anywhere Sensor IP address
```

Where `<IP>` is the IP address for the USM Anywhere Sensor.

Restart the rsyslog Service

You must restart the rsyslog service for the configuration change to take effect.

```
sudo service rsyslog restart
```

Note: When redeploying a sensor with TLS syslog encryption enabled, the new sensor will not maintain your previous encryption configuration. You must configure your TLS syslog encryption again for the redeployed sensor:

- If you have used the default sensor certificates, the redeployed sensor will have generated new certifications. Be sure to use the new sensor certifications when reconfiguring your TLS syslog encryption.
- If you have used your own TLS certificates you must reupload your PEM files again, but they can be the same PEM files you used originally.

Check the Syslog Collection Status

After you have configured the syslog forwarding policy on the required data sources, you can verify the log forwarding in USM Anywhere. When you select the sensor on the Syslog Server page, the Health column displays ✓ for each of the syslog protocols where the sensor has received a packet within the last 10 minutes.

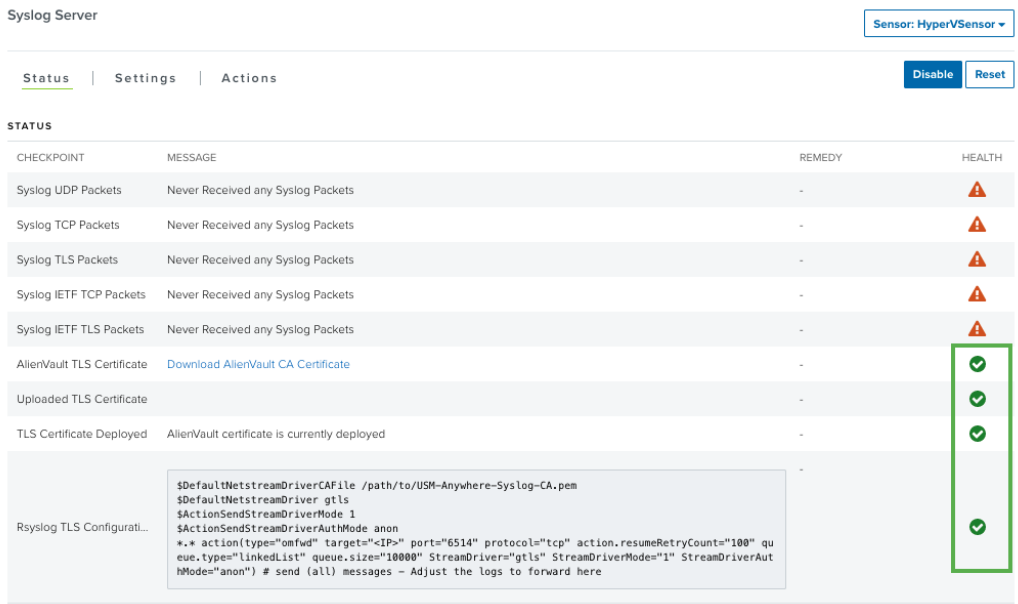

Scroll down to the Stats section to review more detailed information about the syslog activity on the sensor.

STATS

SYSLOG UDP	
Listening on port	**514**
Number of Syslog Packets Received	**1**
Received Syslog from the following IPs	
SYSLOG TCP	
Listening on port	**601**
Number of Syslog Packets Received	**0**
Received Syslog from the following IPs	
SYSLOG TLS	
Listening on port	**6514**
Number of Syslog Packets Received	**2**
Received Syslog from the following IPs	
SYSLOG IETF TCP	
Listening on port	**602**
Number of Syslog Packets Received	**0**
Received Syslog from the following IPs	
SYSLOG IETF TLS	
Listening on port	**6515**
Number of Syslog Packets Received	**0**
Received Syslog from the following IPs	

- **Number of Syslog Packets Received**: Number of packets received by the sensor since it has been up and running. (Restarting the sensor resets this counter.)
- **Received Syslog from the following IPs**: List of IP addresses forwarding logs to the sensor. There is a maximum of 100 IPs, and IPs not sending logs in the last 24 hours are discarded. (Restarting the sensor resets this list.)

Disable Syslog Collection on a USM Anywhere Sensor

The syslog server app is enabled for log collection by default for each deployed USM Anywhere Sensor. If you want to disable the app for a particular sensor, complete the following procedure.

To disable syslog data collection on a sensor

1. In USM Anywhere, go to **Data Sources > Sensors**.
2. Click the **Sensor Apps** tab.
3. In the left navigation menu, click **Syslog Server**.
4. Select the sensor where you want to disable the app.

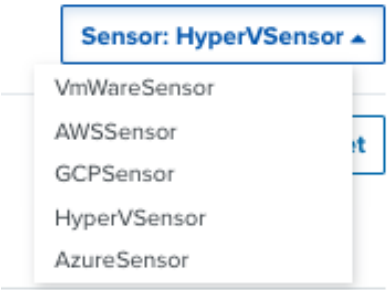

5. Click **Disable**.

The Graylog (GELF) Sensor App

Role Availability	Read-Only	Investigator	Analyst	Manager

The Graylog Extended Log Format (GELF) is a log format designed to overcome many of the limitations of standard syslog. It is a great solution for applications because it provides more robust logging support — larger payloads, compression, and chunking — and developers can leverage libraries and appenders for many programming languages and logging frameworks.

All of the USM Anywhere Sensors use the Graylog (GELF) app, which passively listens to the Graylog UDP port 12201 and collects the GELF log data for processing. To configure your applications to send data to USM Anywhere, you must specify the IP address of your USM Anywhere Sensor and the port number as the Graylog host.

Important: When you configure GELF for your applications, you must use UDP as the transport layer. The Graylog Sensor App *does not support TCP/TLS or HTTP* transport.

For more information, see the following vendor documentation:

- GELF via UDP
- GELF Libraries

The Graylog app is enabled by default for each deployed USM Anywhere Sensor. If you want to disable the app for a particular Sensor, follow this procedure.

To disable GELF data collection on a Sensor

1. In USM Anywhere, go to **Data Sources > Sensors**.
2. Click the **Sensor Apps** tab.
3. In the left navigation menu, click **Graylog**.
4. Select the Sensor where you want to enable the app.

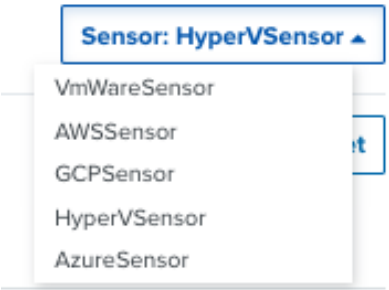

5. Click **Disable**.

Collecting Linux System Logs

In USM Anywhere, you can centralize the collection and analysis of Linux event logs from your servers, making it easier to track the health and security of these systems.

Using the AlienVault Agent

The AlienVault Agent provides simple installation, configuration, and management for host monitoring in USM Anywhere. When you install the AlienVault Agent on a Linux host, it communicates over an encrypted channel to send data directly to USM Anywhere. The agent installation script configures a default set of folders and files to automatically support file integrity monitoring (FIM). You can set the configuration profile to manage the queries that USM Anywhere runs for an asset associated with a deployed agent.

Using AlienVault Agents is the best choice for monitoring endpoints outside of the network, in remote locations, or where deploying a sensor is impractical. Additionally, it provides the ability to query the asset for additional forensic data as part of your investigation activities. See AlienVault Agents for more information about the AlienVault Agent and how you can use it to simplify your endpoint detection and response (EDR), FIM, and rich endpoint telemetry capabilities.

Collecting Logs from Cloud Environments

USM Anywhere provides USM Anywhere Sensors for different cloud environments and collect logs using their native tools:

- AWS Log Discovery and Collection in USM Anywhere
- Azure Log Discovery and Collection in USM Anywhere
- GCP Log Discovery and Collection in USM Anywhere

Sending Logs Directly to a USM Anywhere Sensor

Supplementary to using the AlienVault Agents, you can configure syslog or manually install osquery on your hosts to forward logs to a USM Anywhere Sensor:

- Linux Log Collection with Syslog
- Linux Log Collection with Osquery
- Linux Log Collection with NXLog

Linux Log Collection with Syslog

The use of syslog is required to send log data from Linux systems to the USM Anywhere Sensor IP address over UDP on port 514, over TCP on port 601 or 602, or Transport Layer Security (TLS)-encrypted data over TCP on port 6514 or 6515.

Using Syslog to Send Logs from a Linux System

General InformationSyslog is an industry-standard message logging protocol that is used on many devices and platforms. It provides a mechanism for network devices to send event messages to a logging server, also known as a syslog server. In this case, a USM Anywhere Sensor is acting as the syslog server. USM Anywhere supports both the BSD syslog protocol (RFC 3164) and the syslog protocol (RFC 5424). For RFC 3164, USM Anywhere listens for syslog over UDP on port 514, TCP on port 601, or Transport Layer Security (TLS) on port 6514. For RFC 5424, USM Anywhere listens for syslog over TCP on port 602 or TLS on port 6515. For example, a router might send messages about users logging on to console sessions, while a web server might log access-denied events.

Follow the procedure that corresponds to the Linux distribution you use.

Fedora Linux Distribution

You must have sudo privileges to complete this procedure.

To send logs from Fedora Linux using syslog

1. On your Linux machine, open `/etc/rsyslog.conf` and add the following line:

```
*.* action(type="omfwd" target="<IP>" port="514" protocol="udp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000")
# send (all) messages - Forward to the USM Anywhere Sensor IP address
```

2. Restart rsyslog:

```
sudo service rsyslog restart
```

Red Hat Enterprise Linux Distribution

You must have sudo privileges to complete this procedure.

To send logs from Red Hat Enterprise Linux using syslog

1. On your Linux machine, install rsyslog for RHEL-5 (installed by default for RHEL-6 and 7):

```
sudo yum install rsyslog
```

2. Open `/etc/rsyslog.conf` and add the following line to the start of the file:

```
*.* action(type="omfwd" target="<IP>" port="514" protocol="udp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000")
```

Where `<IP>` is the IP address for the USM Anywhere Sensor.

3. Restart rsyslog:

```
sudo service syslog stop (only for RHEL-5)
sudo service rsyslog restart
```

openSUSE Distributions

You must have sudo privileges to complete this procedure.

To send logs from openSUSE Distributions

1. Install rsyslogd:

```
sudo yast -i rsyslog
```

2. Set rsyslog as syslog server:

   a. Open `/etc/sysconfig/syslog`.

   b. Add the following lines:

```
SYSLOG_DAEMON=”rsyslogd”
RSYSLOGD_COMPAT_VERSION=”4”
```

   c. Save it and run `SuSEconfig`.

3. On your Linux machine, open `/etc/rsyslog.d/remote.conf` and add the following line:

```
*.* action(type="omfwd" target="<IP>" port="514" protocol="udp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000"
```

   Where `<IP>` is the IP address for the USM Anywhere Sensor.

4. Restart rsyslog:

```
sudo service rsyslog restart
```

Debian GNU/Linux and Ubuntu Distributions

You must have sudo privileges to complete this procedure.

To send logs from Debian GNU/Linux and Ubuntu Distributions

1. On your Linux machine, open the appropriate configuration file:

   - (debian) `/etc/rsyslog.conf`
   - (ubuntu) `/etc/rsyslog.d/50-default.conf`

2. Add one of these lines:

- (UDP) `*.* action(type="omfwd" target="<IP>" port="514" protocol="udp" action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000")`
- (TCP) `*.* action(type="omfwd" target="<IP>" port="601" protocol="tcp" action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000")`

Where `<IP>` is the IP address for the USM Anywhere Sensor.

3. Restart rsyslog:

```
sudo service rsyslog restart
```

SUSE Linux Enterprise 11 SP4 - 12 SP1Server Distribution

You must have sudo privileges to complete this procedure.

To send logs from SUSE Linux Enterprise Server Distribution

1. Install the `rsyslogd` package:

```
sudo yast -i rsyslog
```

2. Set rsyslog as syslog server by editing the following parameters in `/etc/sysconfig/syslog`:

```
SYSLOG_DAEMON=”rsyslogd”
RSYSLOGD_COMPAT_VERSION=”4”
```

3. Save the file and run `SuSEconfig`.
4. On your Linux machine, open `rsyslog.d/remote.conf` and add the following line:

```
*.* action(type="omfwd" target="<IP>" port="514" protocol="udp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000")
```

5. Restart rsyslog:

```
sudo rcsyslog restart
```

Solaris Distribution

You must have sudo privileges to complete this procedure.

To send logs from Solaris distributions

1. On your Linux machine, open `/etc/syslog.conf` and add the following line:

```
*.notice @<USM-Anwhere-Sensor-IP-address>
```

> **Important:** In the foregoing command, you must **tab** from auth.notice to `@<USM-Anwhere-Sensor-IP-address>`; if you type a space the command will fail.

2. Stop, then restart syslog:

   Solaris 5.9 and earlier

```
sudo /etc/init.d/syslog stop
sudo /etc/init.d/syslog start
```

   Solaris 5.10 and above

```
#sudo svcadm refresh svc:/system/system-log
```

FreeBSD Distributions

You must have sudo privileges to complete this procedure.

To send logs from FreeBSD Distributions

1. On your Linux machine, open `/etc/syslog.conf` and add the following line:

```
*.* action(type="omfwd" target="<IP>" port="514" protocol="udp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000")
```

2. Restart rsyslog:

```
sudo service syslogd restart
```

Gentoo Distributions

You must have sudo privileges to complete this procedure.

To send logs from Gentoo Distribution

1. On your Linux machine, open `/etc/rsyslog.conf` and add the following line:

```
*.* action(type="omfwd" target="<IP>" port="514" protocol="udp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000")
```

2. Restart rsyslog:

```
sudo /etc/init.d rsyslog restart
```

Arch Linux Distribution

You must have sudo privileges to complete this procedure.

To send logs from Arch Distribution

1. On your Linux machine, open `/etc/syslog-ng/syslog-ng.conf` and add the following line:

```
*.* action(type="omfwd" target="<IP>" port="514" protocol="udp"
action.resumeRetryCount="100" queue.type="linkedList" queue.size="10000")
```

2. Restart rsyslog:

```
sudo systemctl start rsyslog
```

Linux Log Collection with Osquery

Osquery is an operating system instrumentation framework for Linux that exposes this operating system as a high-performance relational database so that SQL queries can explore the operating system data. With osquery, SQL tables represent abstract concepts such as running processes, loaded kernel modules, open network connections, browser plugins, hardware events, or file hashes.

AT&T Cybersecurity recommends that you use osquery to collect data and send them to USM Anywhere through syslog. Alternatively, you can install the AlienVault Agent on your Linux hosts outside of the network to monitor endpoints and collect logs.

Note: Do not run osquery in parallel with the AlienVault Agent because it will interfere with the agent and cause USM Anywhere not to parse the data it receives.

You must have sudo privileges to complete the following procedure.

To collect logs from Linux using osquery

1. If you do not yet have osquery, download it and follow the instructions appropriate for your operating system.
2. Create a text file called `osquery.conf` and copy-paste the contents of this file into it.

   Important: After you copy-paste the text, make sure to edit it so that all strings with equals signs (=) in them remain on the same line. Otherwise, this procedure will fail.

3. Save `osquery.conf` and copy it to `/etc/osquery/`.

   Note: We recommend leaving the queries created by default, but you can create your own osquery configuration.

4. Start the osquery daemon:

```
osqueryd --daemonize --config_path /etc/osquery/osquery.conf
```

Linux Log Collection with NXLog

NXLog is a universal log collection and forwarding agent for various platforms, including Linux. With the NXLog Enterprise Edition, you can scan files and directories to report detected change, known as file integrity monitoring (FIM). USM Anywhere provides an AlienApp for Linux NXLog and the configuration file to collect FIM data.

According to the vendor documentation, FIM is only available in the NXLog Enterprise Edition. In addition, NXLog must have permission to read the files you want to monitor. You can run NXLog as *root*, or make sure the *nxlog* user or group has permission to read the files.

To download the configuration file from USM Anywhere

1. Log in to USM Anywhere and go to **Settings > System**.
2. In the left navigation panel, click **NXLog Configuration** to open the page.
3. Select **Linux Systems**.

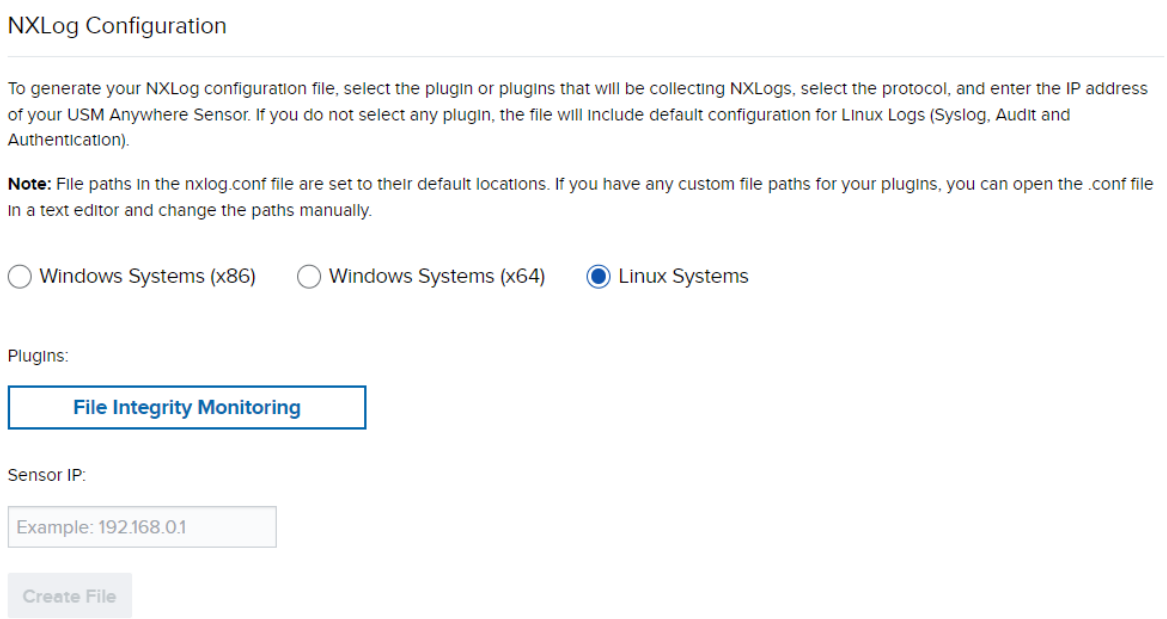

4. Click **File Integrity Monitoring**.
5. Enter the IP address of your USM Anywhere Sensor.

Note: USM Anywhere uses UDP port 514 to forward the logs.

6. Click **Create File** to generate the new `nxlog.conf` file and save it to your NXLog installation directory.

Note: AT&T Cybersecurity recommends you to save a copy of the original `nxlog.conf` file first.

7. Restart NXLog.

The AlienApp for Linux NXLog is auto-discovered. No additional configuration is needed. Following is an example of the downloaded file:

```
#
# Configuration for converting and sending Linux logs
# to AlienVault USM Anywhere.
#
# Version: 0.0.1
# Last modification: 2020-12-02
#

define OUTPUT_DESTINATION_ADDRESS_AND_PORT 1.1.1.1:514

# Default values:
# Note: These values can change depending on the Linux flavour

define INSTALLDIR /opt/nxlog
define LOGDIR %INSTALLDIR%/var/log/nxlog
define MYLOGFILE %LOGDIR%/nxlog.log
LogFile     %MYLOGFILE%

# Load extension common to all inputs
<Extension _json>
        Module  xm_json
</Extension>

<Extension _syslog>
        Module  xm_syslog
</Extension>

# Set all inputs
<Input rsyslog_in>
        Module      im_file
        File        "/var/log/syslog"
        Exec        parse_syslog();
</Input>

<Input secure_in>
```

```
        Module  im_file
        File        "/var/log/auth.log"
        Exec        parse_syslog();
</Input>

<Input audit_in>
        Module  im_file
        File        "/var/log/audit/audit.log"
        Exec        parse_syslog();
</Input>

<Output out>
        Module      om_udp
        Host        %OUTPUT_DESTINATION_ADDRESS_AND_PORT%
        Exec        $EventTime = integer($EventTime) / 1000000;
        Exec        $EventReceivedTime = integer($EventReceivedTime)/1000000;
        Exec        $Message = to_json(); to_syslog_bsd();
</Output>

# Set common route
<Route USM_Out>
        Path rsyslog_in, secure_in, audit_in  => out
</Route>

#######################################################################
####                        FIM-LINUX-NXLOG-EE                     #####
####    Uncomment the following lines for FIM LINUX NXLOG EE       #####
####                           log forwarding                      #####
#######################################################################

## This config includes just the paths specified in nxlog official
## documentation, but more paths can be added if needed

<Input fim_linux>
        Module          im_fim
        File            "/bin/*"
        File            "/etc/*"
        File            "/lib/*"
        File            "/opt/nxlog/bin/*"
        File            "/opt/nxlog/lib/*"
        File            "/sbin/*"
        File            "/usr/bin/*"
        File            "/usr/sbin/*"
        Exclude         "/etc/hosts.deny"
        Exclude         "/etc/mtab"
        Recursive       TRUE
        ScanInterval    1800
</Input>
```

```
<Output fim_out_linux>
        Module  om_udp
        Host    %OUTPUT_DESTINATION_ADDRESS_AND_PORT%
        Exec    $SourceName = "FIM-LINUX-NXLOG-EE";
        Exec    $Message = to_json(); to_syslog_bsd();
</Output>

<Route fim_route_linux>
        Path    fim_linux => fim_out_linux
</Route>

######################################################################
####                        FIM-LINUX-NXLOG-EE                    #####
######################################################################
```

Collecting Windows System Logs

In USM Anywhere, you can centralize the collection and analysis of Microsoft Windows event logs from your servers or desktops, making it easier to track the health and security of these systems. While the AlienVault Agent is ideal for most traditional end-user laptop or desktop environments, there are some situations for which alternative log collection options, such as NXLog, may be preferable. The following table compares some of the most common use cases between the AlienVault Agent and NXLog.

AlienVault Agent vs. NXLog Use Cases

Environmental Demands	Recommended Option
If you need to monitor endpoints outside of the network or in remote locations where it would be impractical to deploy a sensor	AlienVault Agent
If you want the ability to query assets for additional forensic data as part of your investigation activities	AlienVault Agent
If you want the benefits of AT&T Alien Labs actively monitoring endpoints with updated Alien Labs rules, including active process and network activity information	AlienVault Agent
If you a need to restrict off-premise connections for endpoints	NXLog

AlienVault Agent vs. NXLog Use Cases (Continued)

Environmental Demands	Recommended Option
If you need complete control over agent configuration and filtering rules	NXLog
If you have highly active servers that are required to maintain essential business functions where all or most of your resources are dedicated to the server	NXLog

Using the AlienVault Agent

The AlienVault Agent provides simple installation, configuration, and management for host monitoring in USM Anywhere without requiring a lot of manual configuration and setup tasks of a third-party agent. When installing the agent on a Windows host, it communicates over an encrypted channel to send data directly to USM Anywhere. The agent installation script configures a default set of folders, files, and registries to automatically support file integrity monitoring (FIM). You can set the configuration profile to manage the queries that USM Anywhere runs for an asset associated with a deployed agent.

Using AlienVault Agents is the best choice for monitoring endpoints outside of the network, in remote locations, or where deploying a sensor is impractical. Additionally, it provides the ability to query the asset for additional forensic data as part of your investigation activities. See AlienVault Agents for more information about the AlienVault Agent and how you can use it to simplify your endpoint detection and response (EDR), FIM, and rich endpoint telemetry capabilities.

Using NXLog

You can use NXLog to collect and forward Windows events to a USM Anywhere Sensor. NXLog is a universal log collection and forwarding agent for basic Windows event logs. But it's also useful in its own right for suppressing spurious events.

This is the best choice when you need complete control over agent configuration and filtering rules or must restrict cloud connections for the endpoint. There are two ways you can implement NXLog and integrate it with USM Anywhere to collect and forward events from your Windows systems:

- Install and configure NXLog Community Edition (CE) across your Windows hosts to capture events on your end servers and forward them to your USM Anywhere Sensor.
- Use the Windows Event Collector sensor app to manage the NXLog subscription and forward your Windows logs directly to a deployed USM Anywhere Sensor. When you use this method, the sensor acts as the collector and the Windows host will forward the logs directly to the sensor using a private IP address, not over the public Internet.

Note: NXLog provides an open source version and a paid, enterprise version. The USM Anywhere Sensor integration using the Windows Event Collector app is based on the enterprise version. And the custom configuration method is based on the open-source Community Edition.

NXLog CE for Windows Hosts

If you want to collect and forward Microsoft Windows events that are not supported by the Windows Event Collector sensor app or other types of non-Windows application events from a Windows host, you can install and configure NXLog Community Edition (CE) and customize your configuration file for integration with USM Anywhere. You can choose to set up NXLog on each Windows host to forward events directly to the USM Anywhere Sensor, or use a forwarding server as a central collection point.

The Windows NXLog plugin provided by USM Anywhere translates the raw log data into normalized events for analysis. This plugin automatically processes all messages forwarded to the USM Anywhere Sensor where the syslog tag matches the value `eventlog`.

You can choose to forward your NXLogs in one of two ways:

- Forward NXLog Messages Directly to a USM Anywhere Sensor
- Use Windows Server as an NXLog Collector

Note: See this Windows expert's guidance for more useful information about testing and debugging Windows events.

Forward NXLog Messages Directly to a USM Anywhere Sensor

The simplest method to receive NXLog messages is to install NXLog Community Edition (CE) on each Microsoft Windows host and configure it to forward messages to the USM Anywhere Sensor. In the event of a sensor disconnect, NXLog messages are cached locally and will be forwarded when the connection resumes.

To install NXLog and create your configuration file

1. On your Windows host, download and install the latest version of NXLog.

   Note: The NXLog Community Edition is open source and free of charge. But to use the File Integrity Monitoring plugin, you must download and install the NXLog Enterprise Edition instead. See vendor documentation for more information.

2. Make a backup copy of the original file and give it another name. Depending on the version, this file can be `C:\Program Files (x86)\nxlog\conf\nxlog.conf` (32-bit) or `C:\Program Files\nxlog\conf\nxlog.conf` (64-bit).
3. Log in to USM Anywhere and go to **Settings > System**.
4. In the left navigation panel, click **NXLog Configuration** to open the page.

   By default, USM Anywhere displays all the plugins available for the 32-bit Windows system.

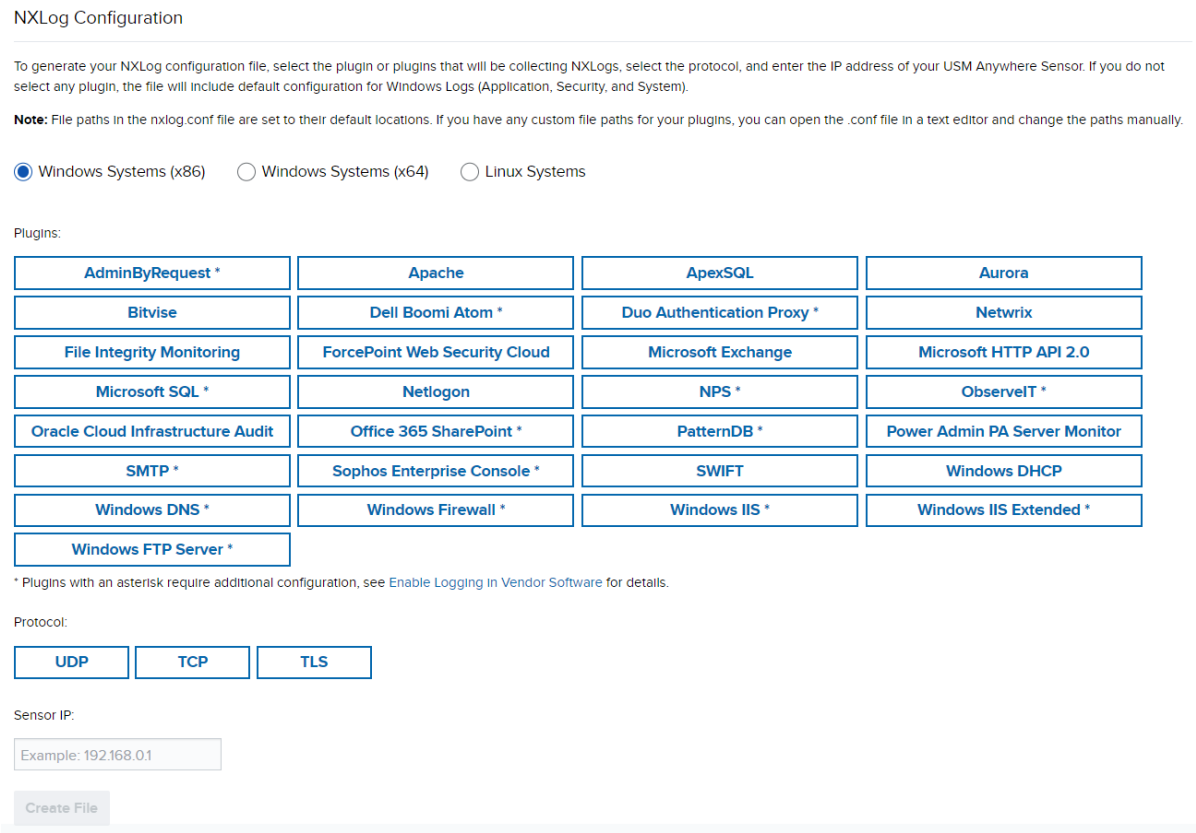

5. Select the desired Windows system and the plugin (or plugins) to collect NXLogs. You don't need to select any plugin to collect default Windows Logs or Sysmon logs.

> **Note:** Plugins with asterisk require additional configuration on the Windows host. See Enable Logging in Vendor Software for details.

6. Select the protocol you want to use.

Download the TLS Certificate

To use TLS, you need to download the certificate and save the file `USM-NXLog-Agent-TLS-CA.pem` in the `\nxlog\cert\` directory on your machine.

- Go to **Data Sources > Sensors**.
- Click **Sensor Apps** tab.
- In the left navigation pane, click **Windows Event Collector** to open the page.

- Download the certificate by clicking the **Download NXLog Agent TLS CA** link.

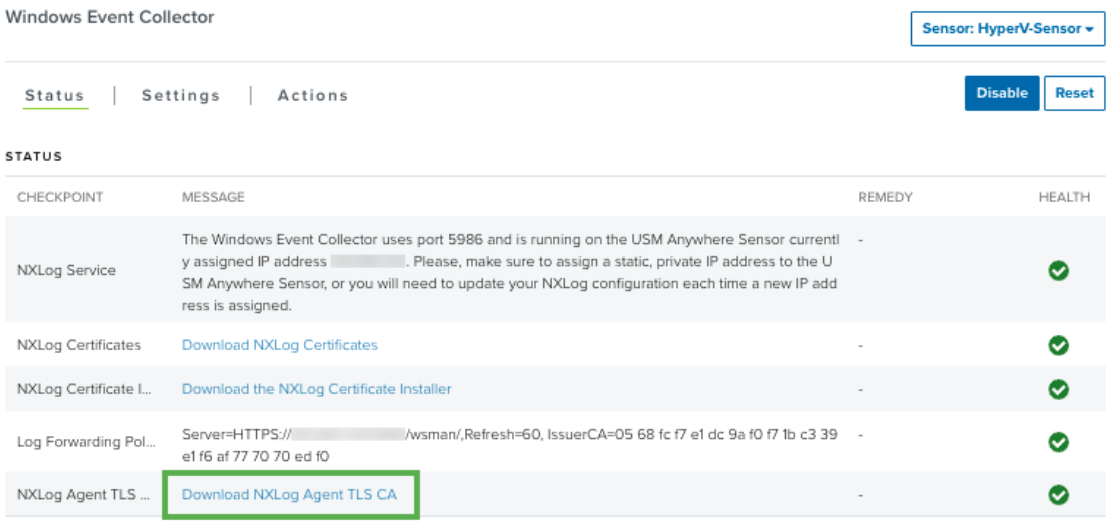

- Copy the certificate file to the client system.

7. Enter the IP address of your USM Anywhere Sensor.
8. Click **Create File** to generate the new `nxlog.conf` file and save it to the `\nxlog\conf\` directory on your machine.
9. Open Windows Services and restart the NXLog service.
10. In USM Anywhere, verify that you are receiving NXLog events.

If you decide not to use NXLog after the installation, you can uninstall the program using the Add or Remove Programs feature in the Windows Control Panel, or see How to Uninstall NXLog for detailed instructions from the vendor.

PatternDB

If you want to limit the events collected and sent to USM Anywhere, you can download the patterndb file provided by AT&T Cybersecurity and place it in the `\nxlog\conf\` directory on your machine. Follow the procedure above to download the NXLog configuration file and select the PatternDB plugin.

> **Important:** Windows Event IDs not present in `patterndb.xml` are not forwarded. Excluding events not relevant to security helps improve the overall performance of the plugin. Consequently, some correlation rules may not be triggered because they rely on those events.

Microsoft Sysmon

System Monitor (Sysmon) is a Windows system service and device driver that remains resident across system reboots to monitor and log system activity to the Windows Event Log. It provides detailed information about process creations, network connections, and changes to file creation time. Sysmon is a free Windows Sysinternals tool from Microsoft. Using NXLog, you can send Sysmon logs to USM Anywhere for event correlation.

To collect Sysmon logs

1. Download the NXLog configuration file from USM Anywhere. You do not need to select any plugin for Sysmon.
2. Open the NXLog configuration file, look for the `<Input eventlog>` tag and add this line under `<QueryList>`:

   ```
   <Select Path="Microsoft-Windows-Sysmon/Operational">*</Select>\
   ```

   With the line added, it should look like this example:

   ```
   <Input eventlog>
     Module im_msvistalog
     Query  <QueryList>\
              <Query Id="0">\
                <Select Path="Application">*</Select>\
                <Select Path="System">*</Select>\
                <Select Path="Security">*</Select>\
                <Select Path="Microsoft-Windows-Sysmon/Operational">*</Select>\
              </Query>\
            </QueryList>
   </Input>
   ```

3. Save the file.
4. Open Windows Services and restart the NXLog service.
5. In USM Anywhere, verify that you are receiving Sysmon events.

Enable Logging in Vendor Software

Some of the vendor plugins need to be configured to enable logging so that USM Anywhere can receive the logs. If you are using any of the plugins below, follow the described integration process to initiate system logging for the plugin.

Dell Boomi Atom

Before configuring NXLog, you must download and install the Dell Boomi Atom on your host. Follow the steps in the Boomi AtomSphere Documentation to download the local Atom installer.

Duo Authentication Proxy

Before configuring NXLog, you must export the Duo Authentication Proxy events to a secondary log file. See Enabling SIEM Logging in the Duo Authentication Proxy for detailed instructions from the vendor.

Forcepoint Web Security Cloud

Before configuring NXLog, you need to format data in Forcepoint (formerly Websense) Web Security Cloud for use by USM Anywhere. See Exporting data to a third-party SIEM tool for detailed instructions from the vendor.

Microsoft 365 SharePoint Server

Before configuring NXLog, you need to configure logging in the Microsoft 365 SharePoint Server. See Configure diagnostic logging in SharePoint Server for detailed instructions from the vendor.

Microsoft DNS Server

Before configuring NXLog, you need to enable debug logging on the DNS server.

To enable DNS server debug logging

1. From the Windows Start Menu, select **All Programs > Administrative Tools > DNS**.
2. From the console tree, right-click the applicable DNS server, and then click **Properties**.
3. Click the **Debug Logging** tab.

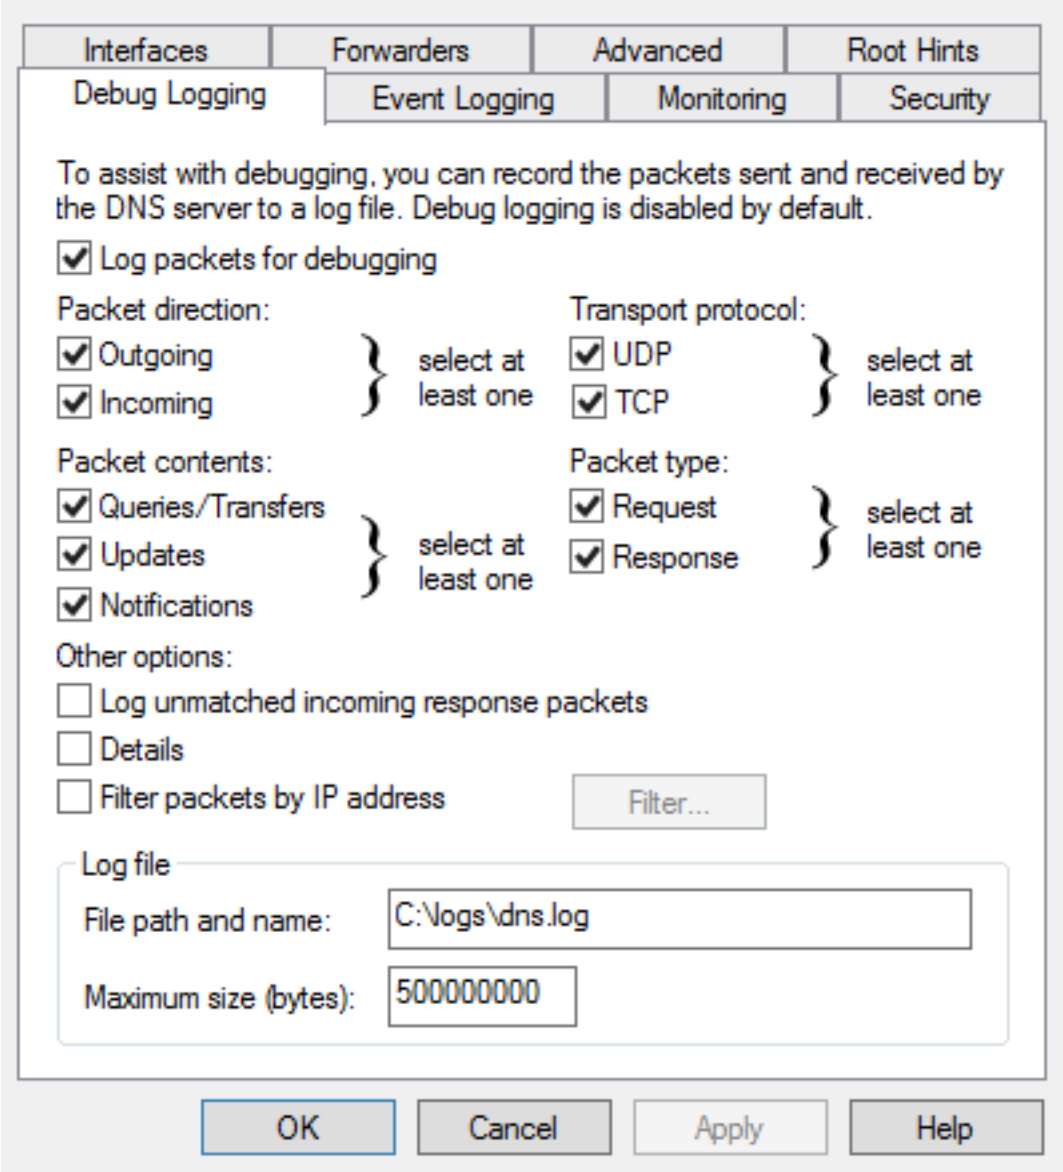

4. Select **Log packets for debugging**.

   The most useful debug logging output comes from selecting at least three options:

   - One option under Packet direction
   - One option under Transport protocol
   - At least one more option in another category

5. (Optional) Consider limiting the traffic captured by applying filters:

a. Select **Filter packets by IP address**.

b. Add the appropriate IP addresses by clicking **Filter**.

> **Warning:** Do not select the **Details** option, because it produces logs in multi-line format instead of single-line, which the USM Anywhere plugins cannot process.

6. Specify the name and location for the log file.
7. Click **Apply** to save and apply the settings.
8. Ensure that the log messages use the `MM/DD/YYYY` date format so USM Anywhere can parse the data correctly.

Microsoft FTP Server

The Microsoft Internet Information Services (IIS) Management Pack includes a FTP Server that you can configure. (See vendor website for documentation.)

Before configuring NXLog, you must configure FTP logging in IIS.

To configure FTP Logging in IIS Manager

1. Open IIS Manager.
2. In the Connections tab, select either the server or the site, and then click the **FTP Logging** icon.
3. Under Log File, click **Select W3C Fields**, and then select the information you want to log.

   Make sure to select the items checked in the screenshot below and click **OK**.

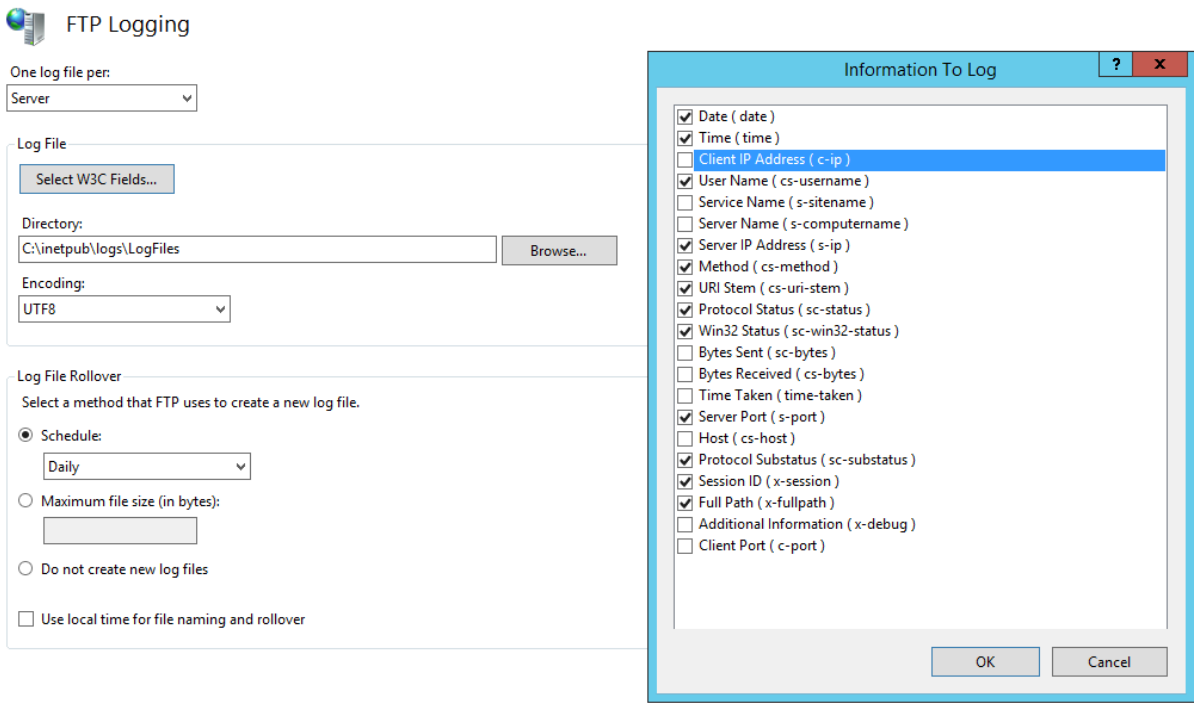

4. Select **UTF8** under Encoding and **Daily** under Schedule.
5. Click **Apply**.
6. Restart the FTP Server for the changes to take effect.

Microsoft HTTP Server API 2.0

Before configuring NXLog, you need to configure logging for the Microsoft HTTP Server API. See Configuring HTTP Server API Error Logging for detailed instructions from the vendor.

Microsoft IIS

Before configuring NXLog , you need to first configure logging on IIS.

To configure logging at the server level

1. Open the IIS Manager.
2. In the Connections tab, select the server and double-click the **Logging** icon.
3. Under One log file per, select **Site**.
4. Under Log File, click **Select Fields** to choose the information you want to log.

   Make sure to match the following screenshot because the *Windows IIS* plugin will look for these fields:

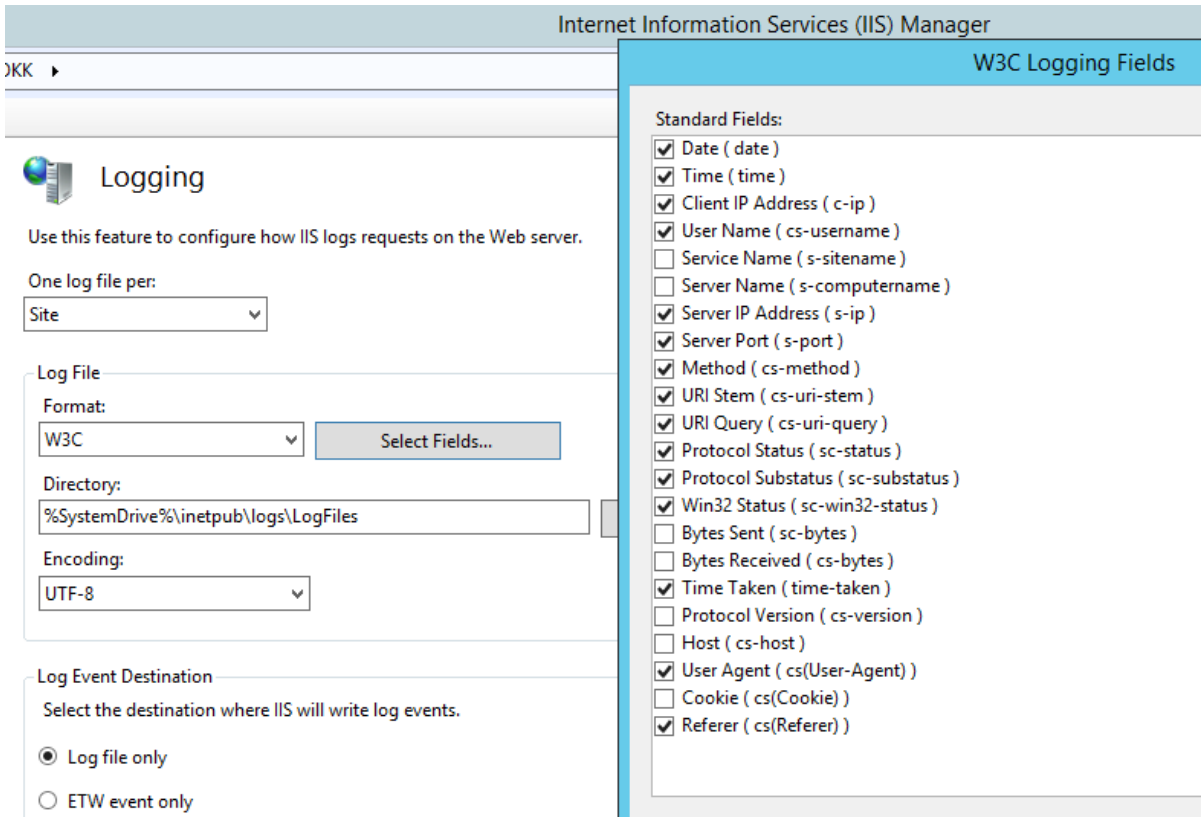

5. To use the *Windows IIS Extended* plugin, you must also enable Bytes Sent (sc-bytes), Bytes Received (cs-bytes), and Host (cs-host). In addition, you can add a custom field for X-Forwarded-For by clicking **Add Field** to fill in the information as shown:

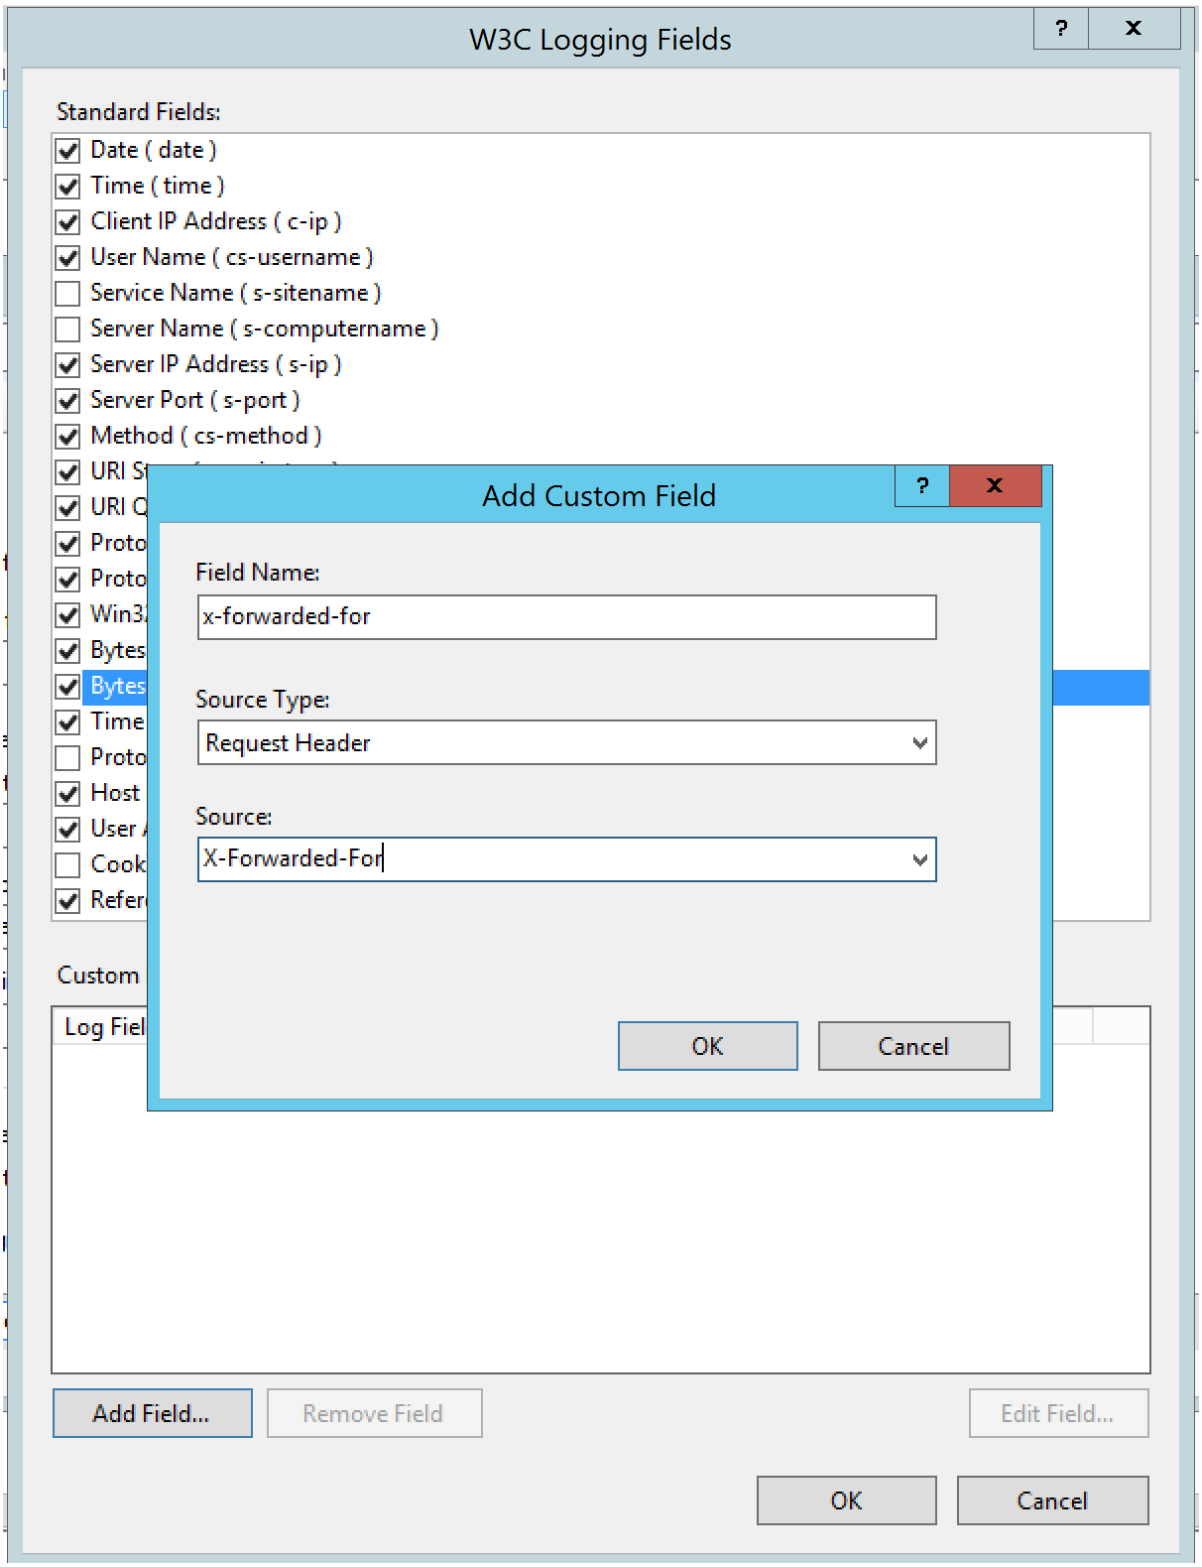
W3C Logging Fields
Standard Fields:
Date (date)
Time (time)
Client IP Address (c-ip)
User Name (cs-username)
Service Name (s-sitename)
Server Name (s-computername)
Server IP Address (s-ip)
Server Port (s-port)
Method (cs-method)
Add Custom Field
Field Name:
x-forwarded-for
Source Type:
Request Header
Source:
X-Forwarded-For
OK
Cancel
Custom
Log Fiel
Add Field...
Remove Field
Edit Field...
OK
Cancel

> **Note:** If you're using Microsoft Windows Server 2008, which doesn't provide an option to add more fields, the Window IIS Extended plugin won't work. Please use the Windows IIS plugin instead.

6. Click **Apply**.

Microsoft NPS

Before configuring NXLog, you need to configure logging for Network Policy Server (NPS) in Microsoft Windows Server 2016 and Server 2019. See Configure NPS Log File Properties for detailed instructions from the vendor. Make sure to select **DTS Compliant** as the log format.

Microsoft SQL Server

Before configuring NXLog, you must have enabled the SQL Server Audit feature and send audit results to the Windows Application Log.

> **Note:** You can use SQL Server Management Studio or Microsoft Transact-SQL (T-SQL) to perform this task. See the Microsoft documentation if you need detailed step-by-step assistance.

To use the SQL Server Management Studio

1. Create a new server audit:
   a. In Object Explorer, expand the **Security** folder, right-click the **Audits** folder, and select **New Audit**.
   b. In the Audit destination list, select **Application Log**.
   c. Select the other options as needed and click **OK**.
2. Create a database-level audit specification
   a. In Object Explorer, expand the database you want to send log to USM Anywhere.
   b. Expand the **Security** folder, right-click the **Database Audit Specifications** folder and select **New Database Audit Specification**.
   c. In the Audit list, select the audit you created in the previous step.
   d. Select the other options as needed and click **OK**.

Netwrix Auditor

Netwrix provides a free add-on for AT&T Cybersecurity to integrate Netwrix Auditor with USM Anywhere through the RESTful API. Before configuring NXLog, you must first use the add-on to generate special Windows event logs for USM Anywhere.

If you haven't already, download the add-on from the Netwrix website. Follow the detailed instructions in their Quick-Start Guide to install and properly configure the add-on. Make sure you have reviewed the events generated by the add-on, as documented in the Quick-Start Guide.

ObserveIT

The ObserveIT plugin leverages the integration support that ObserveIT provides for the HP ArcSight SIEM monitoring software. To configure the SIEM Log Integration, follow the ObserveIT documentation, Configuring CEF Log Integration.

Warning: ObserveIT has been renamed Proofpoint Insider Threat Management.

Important: Follow the steps closely and keep all the default values even though they contain the word "ArcSight". The NXLog configuration file that you download from AT&T Cybersecurity has been specified to parse logs in the default location.

Oracle Cloud Infrastructure

Before configuring NXLog, you need to request a bulk export of audit logs for Oracle Cloud Infrastructure. See Bulk Export of Audit Log Events for detailed instructions from the vendor.

Sophos Enterprise Console

You must download and install the Sophos Reporting Log Writer for this integration. Follow the instructions in the Sophos Reporting Log Writer user guide.

Windows Firewall

Before configuring NXLog, you must first enable logging in Windows Firewall with Advanced Security.

To enable logging in Windows Firewall

1. Open Windows Firewall with Advanced Security.
2. In the Actions panel, click **Properties**.
3. On the Domain Profile tab, click the **Customize** button in Logging.
4. In the new window, complete these steps:
   a. Keep the name as default.
   b. Increase the Size limit (KB) to 8192.

c. Change both Log dropped packets and Log successful connections to **Yes**.

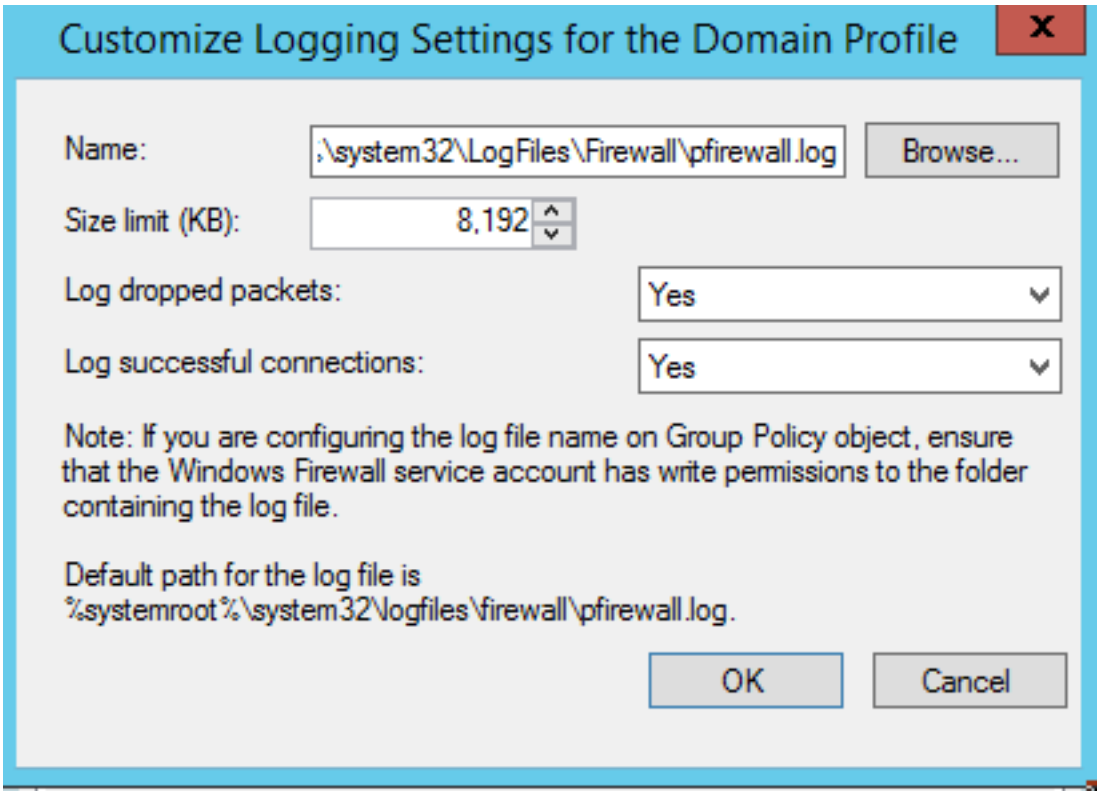

5. Click **OK**.

Manual File Creation and Installation Process

If you are unable to create the `nxlog.conf` file using the process above, or if you wish to edit it manually, you can use the manual process instead.

To install NXLog CE and configure forwarding

1. Download the latest stable version of NXLog.

> **Note:** The NXLog Community Edition is open source and free of charge. But to use the File Integrity Monitoring plugin, you must download and install the NXLog Enterprise Edition instead. See vendor documentation for more information.

2. Make a backup copy of the original file, `C:\Program Files (x86)\nxlog\conf\nxlog.conf` (32-bit) or `C:\Program Files\nxlog\conf\nxlog.conf` (64-bit), and give it another name.
3. Download the NXLog configuration for USM Anywhere and save it as your new `nxlog.conf` file.

4. Open the configuration file for editing and replace `usmsensoriphere` with the IP address of the USM Anywhere Sensor.

   USM Anywhere listens for syslog at UDP port 514, TCP port 601, or TLS/TCP port 6514.

5. Depending on the protocol you decide to use, edit the configuration file as detailed below. Make sure USM Anywhere allows inbound requests to the corresponding port.

   ### To Use UDP

   Keep the defaults and move to the next step.

   ### To Use TCP

   a. Replace `define OUTPUT_DESTINATION_PORT 514` with `define OUTPUT_DESTINATION_PORT 601`.

   b. In `<Output out>`, replace `Module om_udp` with `Module om_tcp`.

   > **Note:** If you are collecting logs from other products as detailed in the next step, you must modify the Output settings in the corresponding section as well. For example, if you are collecting logs from Microsoft Internet Information Services (IIS), you will need to update the `<Output IIS_OUT>` section too.

   ### To Use TLS

   a. First, you need to download the certificate from USM Anywhere, and place the file `USM-NXLog-Agent-TLS-CA.pem` in the `\nxlog\cert\` directory on your machine.

      - Go to **Data Sources > Sensors**.
      - Click **Sensor Apps** tab.
      - In the left navigation pane, click **Windows Event Collector** to open the page.

- Download the certificate by clicking the **Download NXLog Agent TLS CA** link.

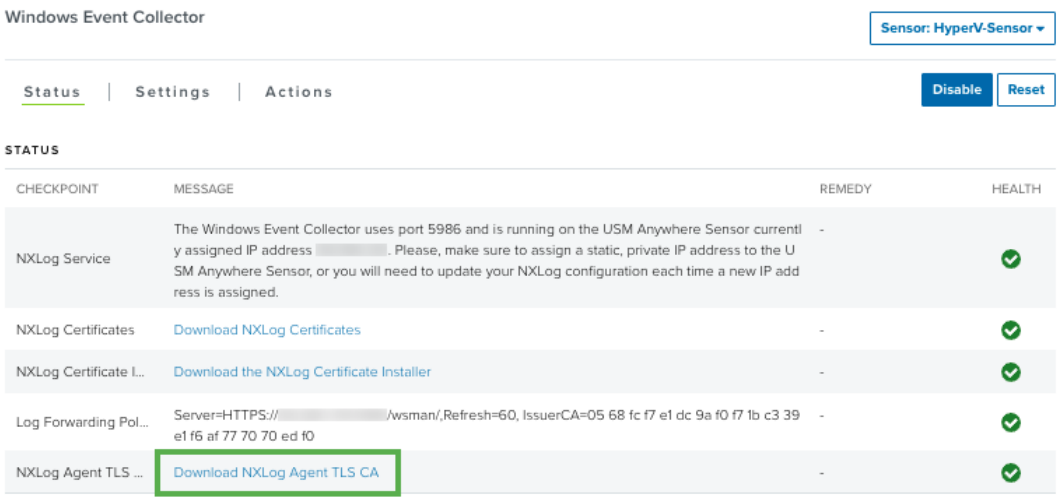

- Copy the certificate file to the client system.

b. In the `nxlog.conf` file, add `define CERTDIR %ROOT%\cert` after the last `define` statement.

c. Replace `define OUTPUT_DESTINATION_PORT 514` with `define OUTPUT_DESTINATION_PORT 6514`.

d. In `<Output out>`, complete these steps:

   - Replace `Module om_udp` with `Module om_ssl`.
   - Add these two lines:

```
CAFile %CERTDIR%\USM-NXLog-Agent-TLS-CA.pem
AllowUntrusted TRUE
```

> **Note:** If you are collecting logs from other products as detailed in the next step, you must modify the Output settings in the corresponding section as well. For example, if you are collecting logs from IIS, you will need to update the `<Output IIS_OUT>` section too.

6. Some sections in the `nxlog.conf` file have been commented out to improve performance. Depending on which product you want to collect logs from, you need to uncomment the corresponding section or sections.
7. Save the file.
8. Open Windows Services and restart the NXLog service.

9. Log in to USM Anywhere and verify that you are receiving NXLog events.

Note: If you need to debug NXLog, open `\nxlog\data\nxlog.log`.

Use Windows Server as an NXLog Collector

You can choose an implementation where you set up each Microsoft Windows source machine to forward its events to a subscribing server that acts as a collector. In this scenario, the collector server acts as a central repository for Windows logs from other servers in the network. With this method, you must set up Windows Event Forwarding (WEF) on *each* Windows source.

Using Windows Server as a means of collecting Windows event logs is intended for use in these USM Anywhere environments:

- On-premises (VMware or Hyper-V Sensors)
- Amazon Web Service (AWS), where the Windows source machines are deployed within one of the following configurations:
  - The Windows source machines, the NXLog agent server, and USM Anywhere Sensor are located in the same Amazon Virtual Private Cloud (VPC).
  - The Windows source machines, the NXLog agent server, and USM Anywhere Sensor are *not* located in the same Amazon VPC, but you have VPC peering configured to allow the NXLog server to communicate with the sensor using UDP port 514.
- Azure, where the Windows source machines, the NXLog agent server, and USM Anywhere Sensor are located in the same virtual network.

Important: Because it does not require that you set up log forwarding on each source, the easiest and most straightforward method for Windows log collection in an Azure environment is to collect the Windows security events from the Azure storage account. However, if you need the additional logs forwarded by NXLog, you can use the following information to configure Windows log collection for this environment.

To set up your Windows Server to collect NXLogs, you need to perform the following two tasks:

- Forward NXLog Messages Directly to a USM Anywhere Sensor
- Configure NXLog Collection and Subscriptions

Configure NXLog Collection and Subscriptions

To have your Microsoft Windows server collect logs from other computers, you need to configure event forwarding for each of the source computers, and configure event collection and subscription on the Windows machine that is the designated collector of the events. This page describes the configuration steps required to set up event processing for all of your machines.

Event Collection and Forwarding

To configure domain computers to collect and forward events

1. Log on to all collector and source computers.

   Note: It is a best practice to use a domain account with administrative privileges.

2. On the collector computer, launch the Administration console and enter the following command:

   ```
   wecutil qc
   ```

3. On each source computer (every computer where you want to run logs), enter the following at an elevated command prompt:

   ```
   winrm quickconfig
   ```

4. Add the collector computer account to the Event Reader Group, and complete these steps:

   a. Edit the group configuration through Local Users and Group.

   b. Add the local computer NETWORK SERVICE account to the Event Log Readers Group.

   c. Change the search location for the NETWORK SERVICE account from the domain to local computer.

      This allows you to access the Security group channel.

   d. Reboot the machine.

      Note: If you don't want to reboot, you can read the Security Log without rebooting by entering `wevtutil sl security /ca:O:BAG:SYD:(A;;0xf0005;;;SY)(A;;0x5;;;BA)(A;;0x1;;;S-1-5-32-573)(A;;0x1;;;s-1-5-20)` from an Administration console.

Subscription Configuration

Set up the event subscription to receive forwarded events on the collector computer.

To add the subscription

1. Log in as administrator on the collector computer.
2. Go to **Administrator Tools** and run **Event Viewer**.
3. In the console tree, click **Subscriptions**.
4. From the Actions menu, click **Create Subscription**.
5. In the Subscriptions Name field, enter the name of the subscription.
6. (Optional.) In the Description field, enter a description of the subscription.
7. In the Destination Log list, select the log file in which you want to store collected events.

   By default, collected events are stored in the ForwardedEvents log.
8. Click **Add** and select the computers from which to collect events.
9. To test connectivity to the source computer, click **Test**.
10. Click **Select Events**.
11. In the Query Filter dialog box, use the controls to specify the criteria that events must meet to be collected.

    To take full advantage of USM Anywhere detection capabilities, AT&T Cybersecurity recommends the following minimum list of channels:

    - Windows Logs → Application
    - Windows Logs → Security
    - Windows Logs → System
    - Application and Services Logs → Microsoft → Windows → AppLocker
    - Application and Services Logs → Microsoft → Windows → PowerShell
    - Application and Services Logs → Microsoft → Windows → Sysmon
    - Application and Services Logs → Microsoft → Windows → Windows Defender
    - Application and Services Logs → Microsoft → Windows → Windows Firewall with Advanced Security
    - Application and Services Logs → Windows PowerShell

USM Anywhere supports a full list of channels, which allows it to detect a wide array of specific types of attacks on the Windows platform.

You can also enable Security Group auditing and Registry auditing on certain sensitive registry keys, such as `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\PowerShell\1\ShellIds\Microsoft.PowerShell`.

12. Under Advanced, select **Minimize Latency**.
13. In the Subscription Properties dialog box, click **OK**.

    This adds the subscription to the Subscriptions pane and, if the operation was successful, the status of the subscription becomes Active.

14. Right-click the new subscription and select **Runtime Status** to verify its status.

    If you have trouble connecting to the source computer, check that the Windows Firewall on the source computer allows inbound connections on TCP port 5985 from the collector.

15. Launch the Administration console and enter the following command to change the content format:

    ```
    wecutil ss <subscription-name> /cf:Events
    ```

    Important: By default, Windows subscriptions use rendered text to format all the events, which the USM Anywhere NXLog AlienApps cannot process. Your forwarded events will not be parsed correctly until this change is made.

16. To test forwarding, create test events using `eventcreate` on the source computer:

    ```
    eventcreate /t error /id 100 /l application /d "Custom event in
    application log"
    ```

Export the Subscriptions

If you are replacing a machine in your network, but you want to run both together for some time without having to reset Event Log Subscriptions manually on the new computer, you can export and re-import all the Event Log Subscriptions settings.

To export subscription configurations

1. From the command line, list the subscriptions:

   ```
   wecutil es
   ```

2. Export the subscriptions:

   ```
   wecutil gs "<subscriptionname>" /f:xml >>"C:\Temp\<subscriptionname>.xml"
   ```

3. Import the subscription:

```
wecutil cs "<subscriptionname>.xml"
```

Note: Importing a subscription with a custom QueryList doesn't work.

4. (Optional.) To use a custom query list, create a subscription as previously described or import a subscription that uses standard settings.
5. Open the subscription and click **Select Events**.
6. Click the **XML** tab, select **Edit query manually**, and paste it in your custom **QueryList**.
7. Click **OK**, then **OK** again.

Troubleshooting Subscription Configuration Exports

See https://www.itprotoday.com/strategy/q-what-are-some-simple-tips-testing-and-troubleshooting-windows-event-forwarding-and for basic troubleshooting help.

See https://technet.microsoft.com/en-us/itpro/windows/keep-secure/use-windows-event-forwarding-to-assist-in-instrusion-detection#how-frequently-are-wef-events-delivered for a more advanced configuration.

Windows Event Collector Sensor App

Role Availability	Read-Only	Investigator	Analyst	Manager

You can use the Windows Event Collector (WEC) sensor app to collect and store Windows events from the computers in your network. When you use the WEC sensor app, the Windows Server machines function as the sender, and the WEC sensor app itself functions as the collector for the events. However, for most instances AT&T Cybersecurity recommends that for enhanced performance and functionality, you should use the Windows Agent or the NXLogs plugin to monitor Windows event logs.

Installation of the WEC sensor app includes these prerequisites:

- Windows Server 2008, 2012, or 2019.
- PowerShell 3.0 or newer.
- A USM Anywhere Sensor with a private, static IP address, deployed in the same network forwarding logs to the WEC sensor app.
- USM Anywhere Sensors require TLS 1.2 for WEC. These are the accepted ciphers:

```
TLS_RSA_WITH_AES_256_GCM_SHA384
```

```
TLS_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
TLS_RSA_WITH_AES_256_CBC_SHA
TLS_RSA_WITH_AES_128_CBC_SHA
```

Installation and setup of the sensor requires:

- Microsoft Windows Event Collector Sensor App Setup.
- Windows Event Collector Sensor App Log Forwarding.
- Windows Event Collector Sysmon InstallationSystem Monitor (Sysmon) is a Windows system service and device driver that remains resident across system reboots to monitor and log system activity to the Windows Event Log. It provides detailed information about process creations, network connections, and changes to file creation time. Sysmon is a free Windows Sysinternals tool from Microsoft. Installation of Sysmon is optional, but highly recommended.To install SysmonDownload the Sysmon ZIP file and unzip it in the target system.Download the Sysmon configuration file to a folder and name the file sysmon_config.xml.https://cybersecurity.att.com/documentation/resources/downloads/sysmon_config_schema4_0.xmlInstall Sysmon in the Windows system and execute the following command:sysmon.exe -accepteula -h md5 -n -l -i sysmon_config.xmlSysmon starts logging the information to the Windows Event Log.Open USM Anywhere and verify that you are receiving Sysmon events..

Microsoft Windows Event Collector Sensor App Setup

To use the Windows Event Collector (WEC) sensor app, you need to download the certificate from USM Anywhere and install it to the Microsoft Windows Server machines on the network that will be forwarding the event logs. A PowerShell script for the installation is linked below, but you can use the Windows Event Collector Sensor App Manual Certificate Installation method if you need to install the certificate on an Active Directory (AD) domain controller.

> **Note:** See your vendor's documentation for headless deployments or more advanced configurations.

Download the Certificate

The Windows Server needs a certificate to establish a trusted connection between the USM Anywhere Sensor (collector) and Windows instances (sender). This certificate is available to download as a `USM-NXLog-client.pfx` file from USM Anywhere when you enable the WEC sensor app.

To download the certificate for the WEC sensor app

1. In USM Anywhere, go to **Data Sources > Sensors**.
2. Click the **Sensor Apps** tab.
3. In the left navigation list, click **Windows Event Collector**.
4. Click the Sensor drop-down list and select the deployed USM Anywhere Sensor you want the app to be installed on.

   If you have more than one deployed USM Anywhere Sensor, choose the sensor that is deployed in the same network as the Windows Server and client systems where you plan to configure a subscription and log forwarding to USM Anywhere.

5. In the Status tab, click the **Download NXLog Certificates** link and save the certificate.

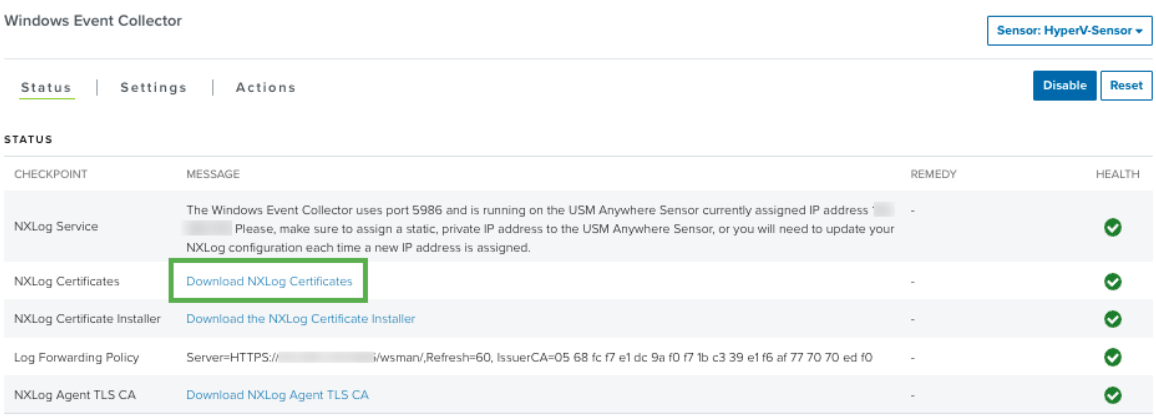

Install and Configure the Certificate on the Windows Server

AT&T Cybersecurity provides a PowerShell installer script that you can use to automatically install the certificates. However, if you need to manually perform the installation, you can follow the Windows Event Collector Sensor App Manual Certificate Installation to install the certificate on your Windows Server.

Using the Certificate Installer Script

The PowerShell installer script is the easiest method for installing the NXLog certificates on your Windows Server so that you can configure Windows event forwarding for a USM Anywhere Sensor.

To use the installer script

1. In USM Anywhere, go to **Data Sources > Sensors**.
2. Click the **Sensor Apps** tab.
3. In the left navigation list, click **Windows Event Collector**
4. Click the **Status** tab and then the **Download the NXLog Certificate Installer** link.
5. On the Windows Server, execute the script from a PowerShell terminal.
6. At the dialog box prompt, select the certificate file.

   The script automatically asks to remove the previous certificates in the case of an earlier USM Anywhere NXLog installation. AT&T Cybersecurity recommends that you remove the previous certificates to avoid potential conflicts.

When the installation is complete, the terminal window displays a confirmation and provides information about next steps to set up event forwarding.

Windows Event Collector Sensor App Manual Certificate Installation

Although the PowerShell installation is recommended, you can also perform the certificate installation manually. After the initial certificate installation, you will need to use the Microsoft Windows HTTP Services (WinHTTP) Certificate Configuration Tool (WinHttpCertCfg.exe) to complete the configuration of the client certificate.

To manually install the certificate

1. Copy the downloaded certificate file to the Windows Server.
2. Double-click the **USM-NXLog-client.pfx** file to launch the Certificate Import Wizard.
3. For the Store Location, select the **Local Machine**.

> **Note:** Windows Server 2008 does not present the option to import into the Local Machine certificate store. For Windows 2008 installations, use the information in the following Microsoft document to import the certificate into the Local Machine certificate store:
>
> https://technet.microsoft.com/en-us/library/cc754841(v=ws.11).aspx

4. When the wizard prompts you for a password, leave it blank and click **Next**.
5. Select **automatically select the certificate store based on the type of certificate** and click **Next** to finish.

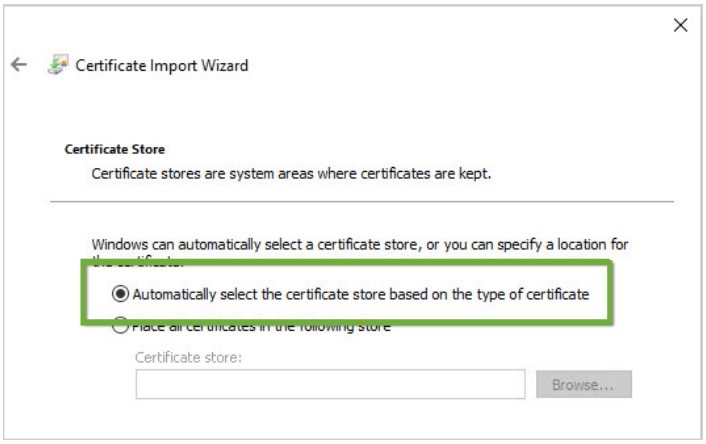

To configure WinHTTP

> **Important:** In order to access the Security event log, the Network Service account must be in the Event Log Readers group.

1. If you do not already have the WinHttpCertCfg.exe tool on your Windows Server, download and install it.
2. Go to the **Administrative Tools** in Windows and open the **Computer Management** utility.
3. Select **Local Users and Groups > Groups > Event Log Readers**.

> **Note:** If your system is deployed as an Active Directory domain controller, Local Users and Groups will not be available. In this case, refer to the Windows documentation to add the network service account to the domain Event Log Readers group.

4. Right-click **Event Log Readers** and select **Add to Group**.

5. Click **Add**.
6. In the Enter the object names to select field, enter **Network Service** as the object name and click **Check Names**.
7. Click **OK** and close the Computer Management utility.
8. Give the Network Service account access to the installed certificate:

   ```
   winhttpcertcfg -g -c LOCAL_MACHINE\my -s USM-NXLog-client -a
   NetworkService
   ```

   If `winhttpcertcfg` is not in the path, you might find it in `C:\Program Files (x86)\Windows Resource Kits\Tools\`.

Windows Event Collector Sensor App Log Forwarding

Microsoft Windows Event Forwarding (WEF) reads any operational or administrative event log on a device and forwards the events you choose to the Windows Event Collector (WEC) sensor app. On the device that you set up as a Windows Event Log collector, you configure subscriptions that pull the desired logs from any number of source computers. No special configuration is required on the source computers, other than that Windows Remote Management (WinRM) should be enabled, the WinRM Windows Firewall exceptions be enabled, and the computer account for the collector must have read permission on the logs that you want to subscribe to.

Set Up Windows Event Forwarding

USM Anywhere provides the log forwarding policy that you use to set up the WEF on your Windows Server.

To get the USM Anywhere log forwarding policy

1. In USM Anywhere, go to **Data Sources > Sensors**.
2. Click the **Sensor Apps** tab.
3. In the left navigation list, select **Windows Event Collector**.
4. Select the USM Anywhere Sensor where you enabled the WEC sensor app.
5. Copy the policy from the field labeled Log Forwarding Policy. You will use this in the next procedure to configure the policy on your Windows Server. The policy follows this pattern:

   ```
   Server=https://<FQDN_of_the_
   collector>:5987/wsman/SubscriptionManager/WEC,Refresh=<REFRESH_INTERVAL_
   IN_SECONDS>,IssuerCA=<CERTIFICATE_THUMBPRINT>
   ```

To configure the policy on your Windows Server

1. On the Windows Server, go to the Control Panel and open the **Local Group Policy Editor**.
2. Select **Computer Configuration > Administrative Templates > Windows Components > Event Forwarding**, and then click **Configure Target Subscription Manager**.
3. Click the **Edit policy setting** link.
4. In the Configure Target Subscription Manager window, make sure that the subscription is marked as **Enabled**.
5. In the Options section of the window, click **Show** to open the subscription managers.
6. In the new Show Contents window, paste the policy that you copied from USM Anywhere in the previous procedure into the new subscription Value field.
7. Click **OK** and close the Local Group Policy Editor.
8. Open the terminal and apply the new configurations by entering this:

```
gpupdate /force
```

Verify the Windows Event Log Collection

You can verify that your Windows Event Log collection configurations are correct by reviewing the event logs.

To review the Windows Event Logs

1. On the Windows Server, open the **Event Viewer**.
2. Go to **Applications and Services Logs > Microsoft > Windows > Eventlog-ForwardingPlugin** and check for any errors.

   You might see warnings if there are any paths that are not configured on your Windows Servers.

If the Windows Event Log collection configuration is without errors or warnings, you can view the events in the USM Anywhere Events List View page.

Windows Event Collector Sysmon Installation

System Monitor (Sysmon) is a Windows system service and device driver that remains resident across system reboots to monitor and log system activity to the Windows Event Log. It provides detailed information about process creations, network connections, and changes to file creation time. Sysmon is a free Windows Sysinternals tool from Microsoft.

Installation of Sysmon is optional, but highly recommended.

To install Sysmon

1. Download the Sysmon ZIP file and unzip it in the target system.
2. Download the Sysmon configuration file to a folder and name the file **sysmon_config.xml**.

   https://cybersecurity.att.com/documentation/resources/downloads/sysmon_config_schema4_0.xml

3. Install Sysmon in the Windows system and execute the following command:

   ```
   sysmon.exe -accepteula -h md5 -n -l -i sysmon_config.xml
   ```

   Sysmon starts logging the information to the Windows Event Log.

4. Open USM Anywhere and verify that you are receiving Sysmon events.

File Integrity Monitoring

File integrity monitoring (FIM) is a mechanism for validating the integrity of operating system and application software files using a verification method between the current file state and a known, good baseline. It is one of the most powerful techniques used to secure IT infrastructures and business data against a wide variety of both known and unknown threats.

The AlienVault Agent

AlienVault offers the lightweight AlienApp Agent as the recommended option for FIM. See the section on The AlienVault Agent to learn more about the agent.

For systems that don't have the AlienVault Agent installed, you can manually enable FIM inside the system.

Manual FIM Configuration Options

If you choose not to use the AlienApp Agent for FIM, you can manually configure FIM on your Linux or Windows system.

Manual FIM Configuration for Linux

For Linux systems that do not have the AlienVault Agent installed, you can enable FIM within USM Anywhere by configuring the osquery agent to monitor and track file changes on those systems. The osquery configuration file (typically named `osquery.conf`) contains the configuration options and queries that osquery uses when it runs. AlienVault provides a default configuration file that you can use to enable FIM for Linux systems in your USM Anywhere environment to identify system and software file changes and forward this information to the USM Anywhere Sensor.

For more information about installing and configuring osquery on your Linux systems, see Linux Log Collection with Osquery.

Manual FIM Configuration for Windows

For Windows systems that do not have the AlienVault Agent installed, you can use FIM to identify changes in system files, folders, and Microsoft Windows registries. To use FIM, you configure Windows systems so that USM Anywhere can view Windows audit object access events. To do so, you need to enable file auditing and update security policy settings. After applying policy changes to include audit object events in Windows security logs, NXLog will forward those events to the USM Anywhere Sensor.

See NXLog CE for Windows Hosts for detailed information about using NXLog to forward these events.

Configuring Policy Settings for Object Access Audit Events

Local Policies determine the security options for a user or service account and are based on the computer and the rights for the account on that computer. These policies can be used to configure an audit policy, which determines which security events will be logged into the Security log on the computer (successful attempts, failed attempts, or both). This Security log is accessible from the Event Viewer.

To define local group policy settings for object access audit events

1. On a selected Windows system, open the **Local Group Policy Editor**.
2. Navigate to **Computer Configuration > Windows Settings > Security Settings > Local Policies > Audit Policy**.

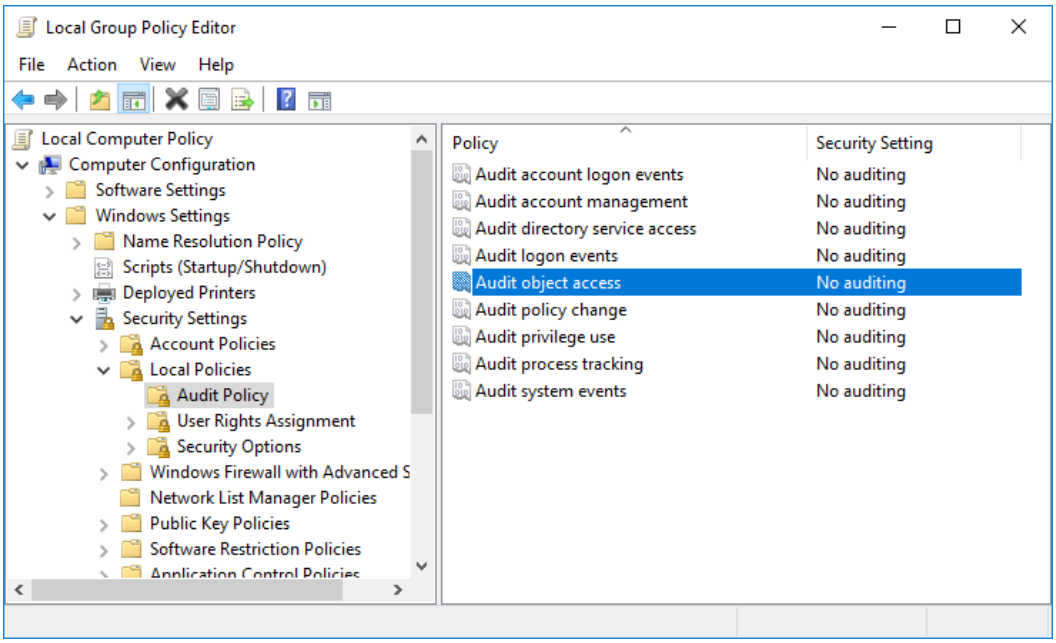

3. Open the **Audit object access** policy.
4. In the dialog, select the **Success** and **Failure** check boxes to enable auditing.

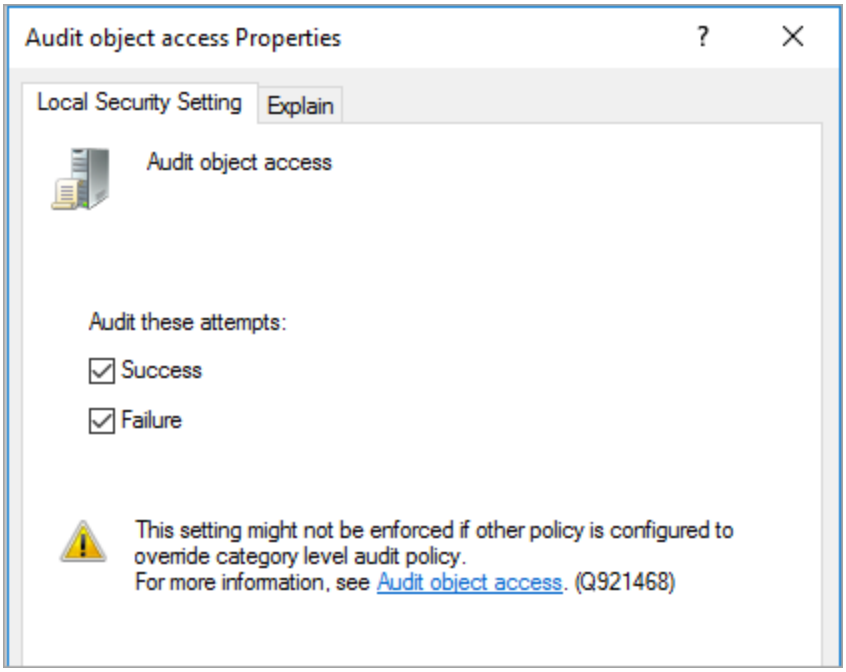

5. Click **Apply** and then **OK**.

Configuring Policy for File Auditing in the Active Directory Domain

In order to track file system changes on a Microsoft Active Directory Domain so that USM Anywhere can view Windows audit object access events, first you must set the Windows group policy to keep track of file system changes.

The following example uses the default domain controller policy in order to track changes on a domain controller. Your actual policy might be different, depending on your particular domain configuration.

To audit changes on a domain controller

1. From the Server Manager, open up Group Policy Management, and expand the domain to select the policy you want to edit.
2. Right-click the policy and choose **Edit**.
3. Select the **Security Options** and change the **Audit: Force audit policy subcategory settings** option to enable it.

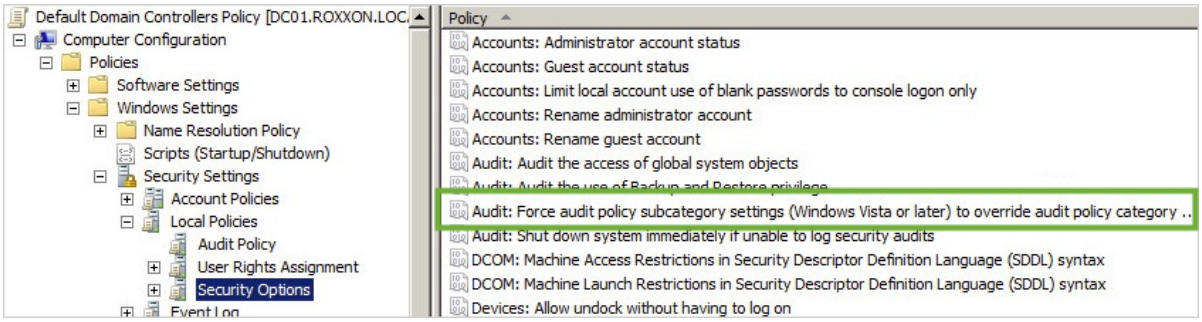

   This allows the more granular advanced audit policy settings instead of the general categories that are enabled by default.

4. Locate and expand the **Advanced Audit Policy Configuration** option.

   > **Note:** This procedure is primarily concerned with object auditing, but you will need to make sure the other policies, such as account lockout, are correct for your organization. Remember, these advanced policies are now taking precedence.

5. For the Audit File System policy, change the configuration to Success and Failure.

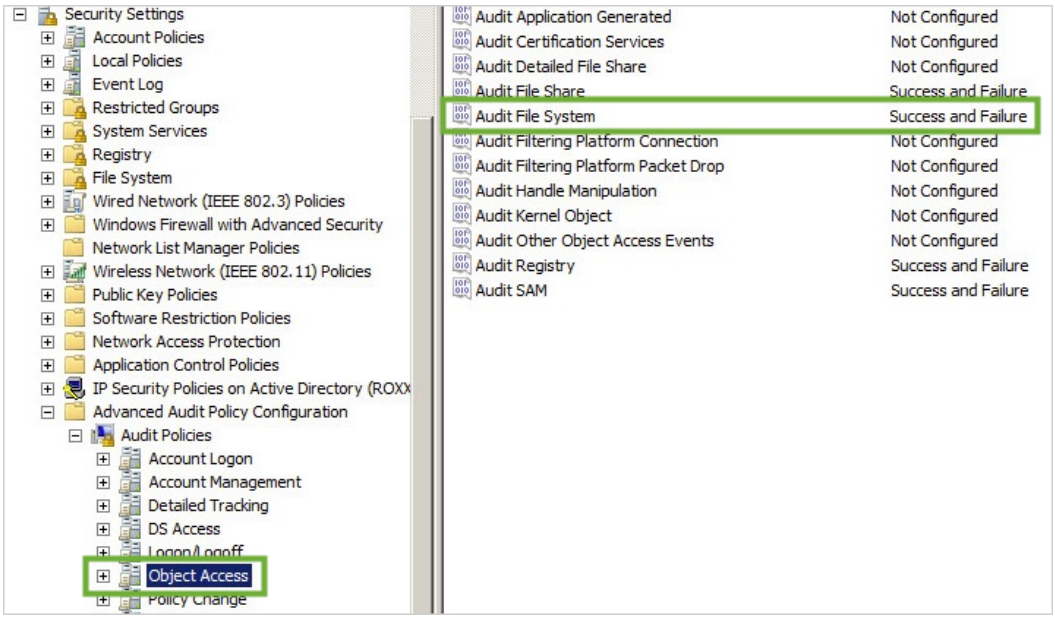

6. Verify that the Group Policy you just edited is enforced, and applied to the domain per your particular configuration.

Configuring a Folder for Auditing

In order for the policy to be effective, you need to enable auditing on the files and directories that you want to monitor. You might be tempted to just enable the entire filesystem, and inherit throughout, and you *could* do that — however, this will be extremely detrimental to the operation of the server, creating hundreds or thousands of events per minute. You should also consider how often you expect changes to the folders you are auditing because this can be quite noisy.

Note: You can only set up file and folder auditing on NTFS drives.

To apply or modify auditing policy settings

1. Open **Windows Explorer** and navigate to the file or folder you want to audit.

   Note: Because the Windows security log is limited in size and new audit events will be stored there, carefully select the files and folders to be audited. Also, consider the amount of disk space that you want to devote to the security log. The maximum size for the security log is set in Event Viewer.

2. Right-click the file or folder and select **Properties**.
3. Select the **Security** tab and click **Advanced**.

4. Select the **Auditing** tab and click **Continue** if prompted.

   This displays the auditing policies for the file or folder.

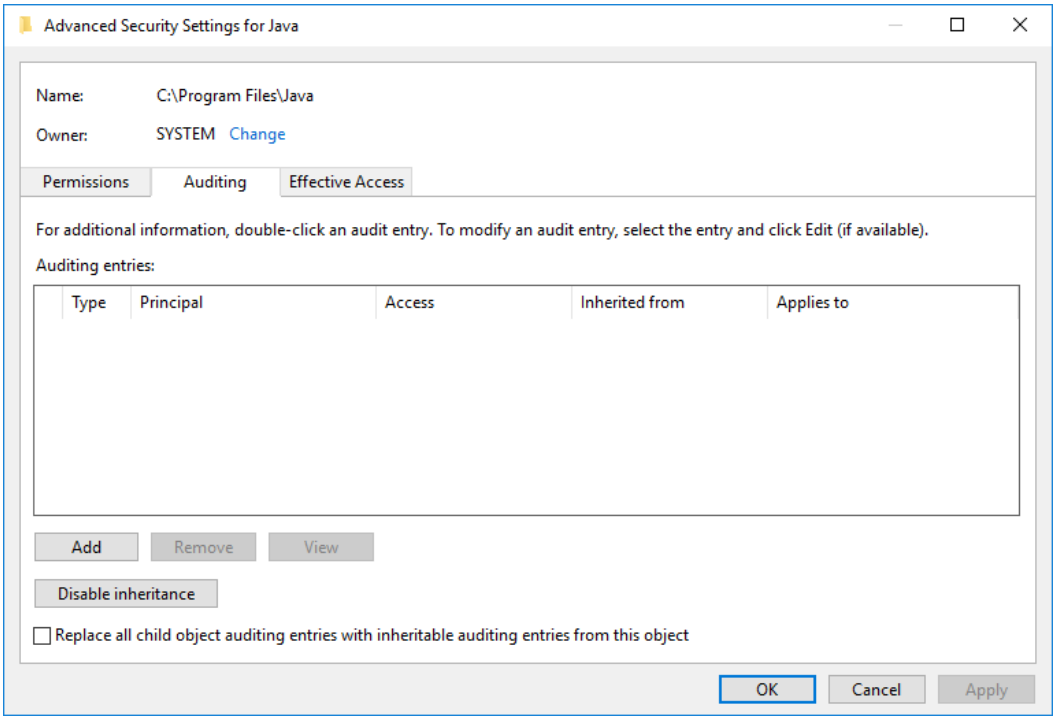

5. Perform one of the following operations:

   - To set up auditing for a new user or group, click **Add**. In the **Enter the object name to select** field, enter the name of the user or group that you want to audit and click **OK**.
   - To remove auditing for an existing group or user, select the group or user name, click **Remove**, and click **OK**. You can skip the remaining steps.
   - To view or change auditing for an existing group or user, select the name and click **Edit**.

6. Set the **Applies to** option to specify the location that you want to audit.

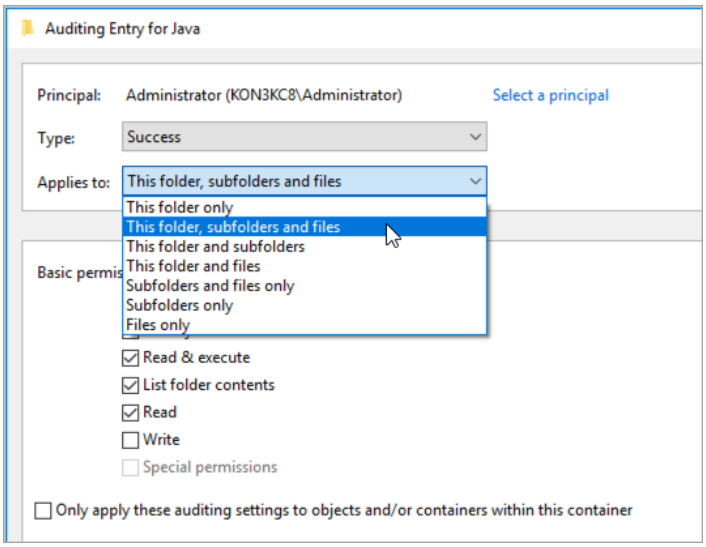

7. In the **permissions** box, select the actions that you want to audit.

   You can click **Show advanced permissions** to display additional permissions for selection. For FIM enablement, Create, Write, Append, and Delete permissions are key.

8. (Optional) If you want to prevent subordinate files and subfolders of the original object from inheriting audit settings, select the **Apply these auditing entries to objects and/or containers within this container only** option.

9. Click **OK**.

Testing and Viewing Events

After enabling object access auditing, view the security log in the Windows Event Viewer to see that the audit events are now collected. You can test to make sure the events are properly generated in Windows Event viewer by creating a file, editing it, moving it out of the folder, and then moving it back. This should generate the events in the Event Viewer of a Windows Server similar to the following example.

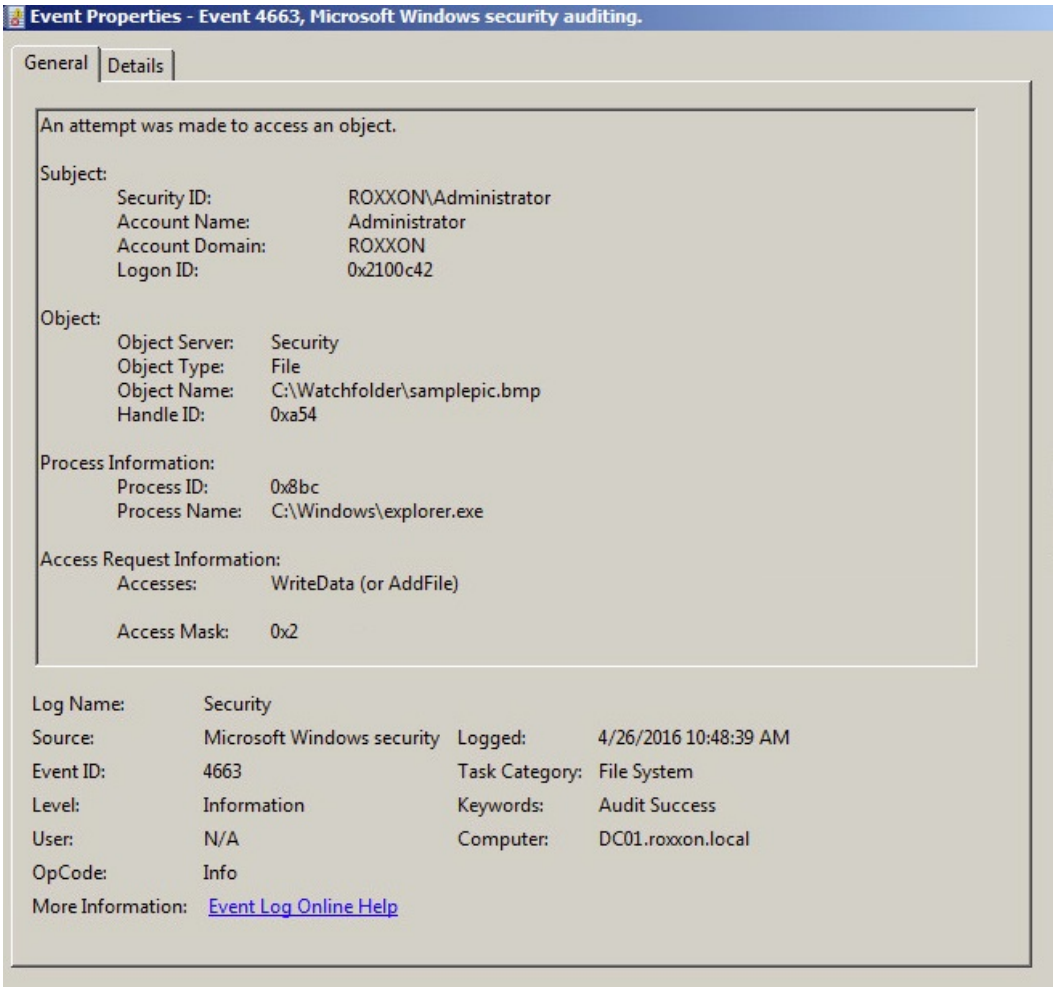

When NXLog is set up to forward these events to the USM Anywhere Sensor, the audit events are available in your USM Anywhere environment.

Scheduling Active Directory Scans from the Job Scheduler Page

Role Availability ✗ Read-Only ✗ Investigator ✓ Analyst ✓ Manager

To effectively manage your Microsoft Windows systems, USM Anywhere can perform scans through an Active Directory (AD) server to collect inventory information. When you configure your VMware Sensor, Microsoft Hyper-V Sensor, or Microsoft Azure Sensor, you can define the credentials that USM Anywhere will use to perform AD scans through the sensor. When you configure these credentials, USM Anywhere performs an initial AD asset scan. You can also schedule a job to perform scans through the Active Directory Scanner and collect updated information about the assets managed by your AD server. The scan returns information for each computer in the AD domain in the following format:

```
Name : WIN2K12-DC
DistinguishedName : CN=WIN2K12-DC,OU=Domain
Controllers,DC=ECORP,DC=local
DNSHostName : WIN2K12-DC.ECORP.local
OperatingSystem : Windows Server 2012 R2 Standard
OperatingSystemServicePack :
OperatingSystemVersion : 6.3 (9600)
IPv4Address : 10.20.30.15
```

The Active Directory Scanner runs a PowerShell (version 5.1 or later) command through Windows Remote Management (WinRM) (version 2.0 or later). See Granting Access to Active Directory for USM Anywhere for information about configuring the AD server to allow access for USM Anywhere,.

To schedule an AD scan job

1. Go to **Settings > Scheduler**.
2. In the left navigation menu, click **Asset Scans**.
3. On the right side of the page, click **Create Scan Job**.

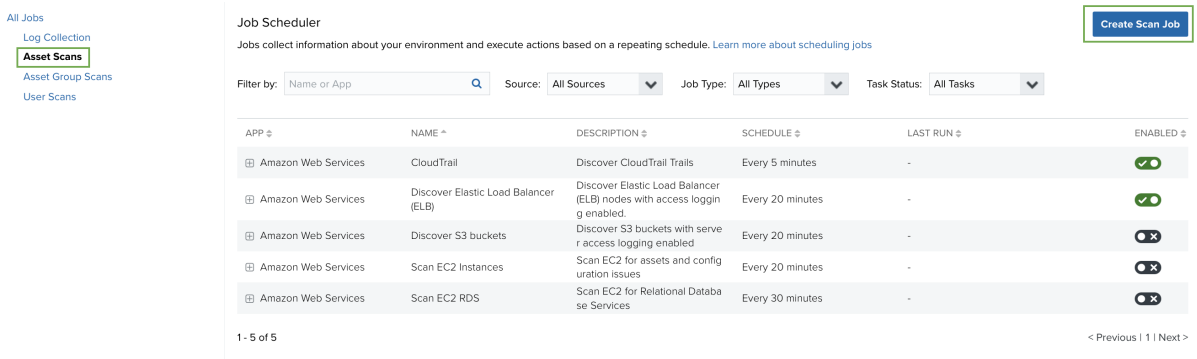

This opens the Schedule New Job dialog box.

4. Enter the name and description for the job.

   The description is optional, but it is a best practice to provide this information so that others can easily understand what it does.

5. Select **Sensor** as the source for your new job.

6. In Action Type, select **Active Directory Scanner**.

7. If you have more than one deployed USM Anywhere Sensor, select the sensor you want to use to run the scan.

   This should be the sensor that is associated with the asset that you want to specify as the target.

8. In App Action, the **Get Active Directory Asset Information** option is already selected.

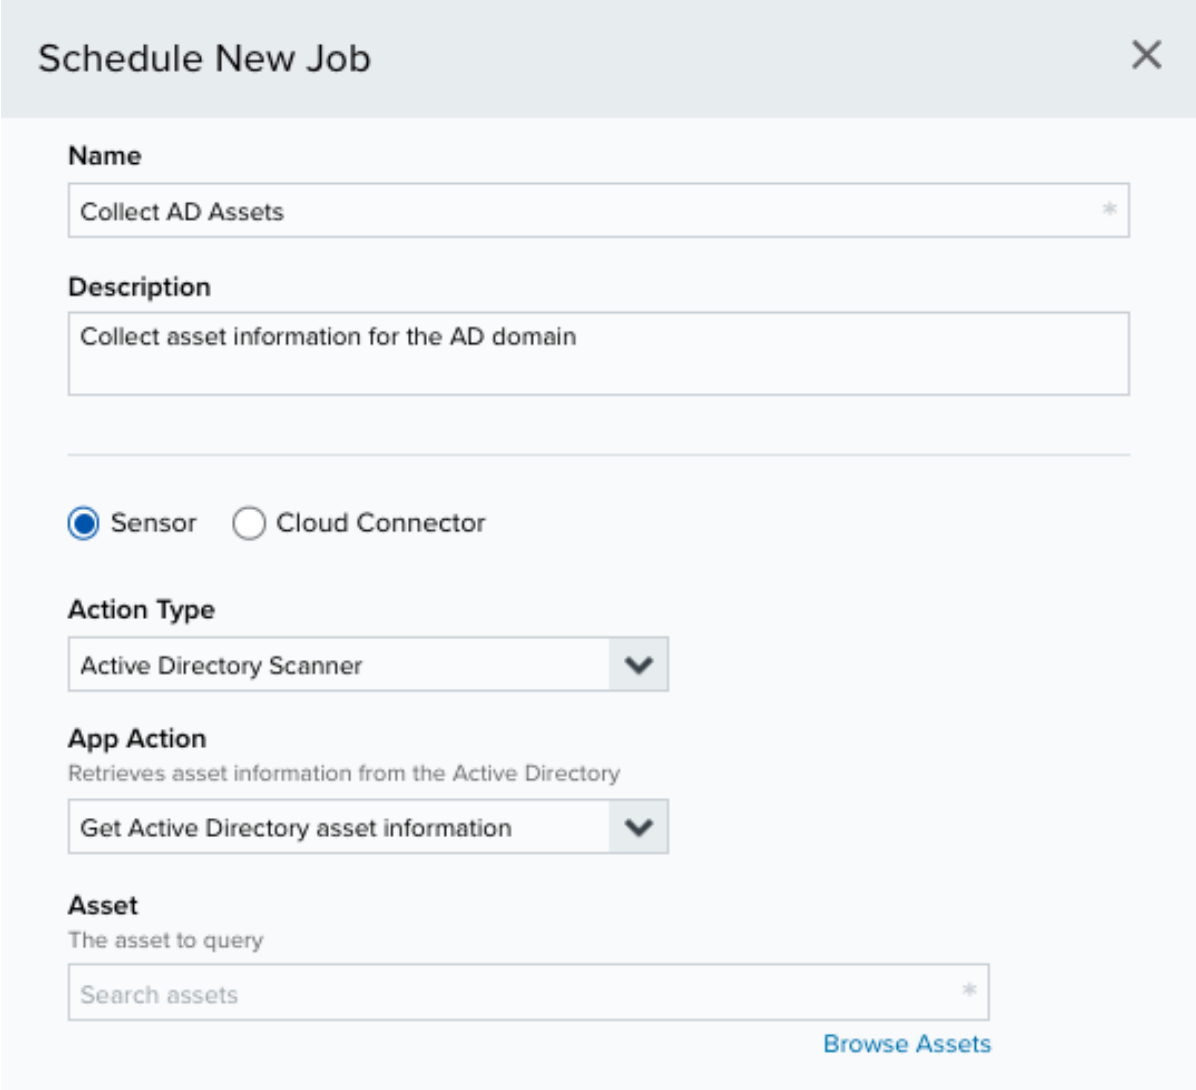

9. Specify the asset that you want to use as a target for the action.

   You can enter the name or IP address of the asset in the field to display matching items that you can select. Or you can click **Browse Assets** to open the Select Asset dialog box and browse the asset list to make your selection.

10. In the Schedule section, specify when USM Anywhere runs the job:

a. Select the increment as **Minute**, **Hour**, **Day**, **Week**, **Month**, or **Year**.

> **Warning:** After a frequency change, monitor the system to check its performance. For example, you can check the system load and CPU. See USM Anywhere System Monitor for more information.

b. Set the interval options for the increment.

The selected increment determines the available options. For example, on a weekly increment, you can select the days of the week to run the job.

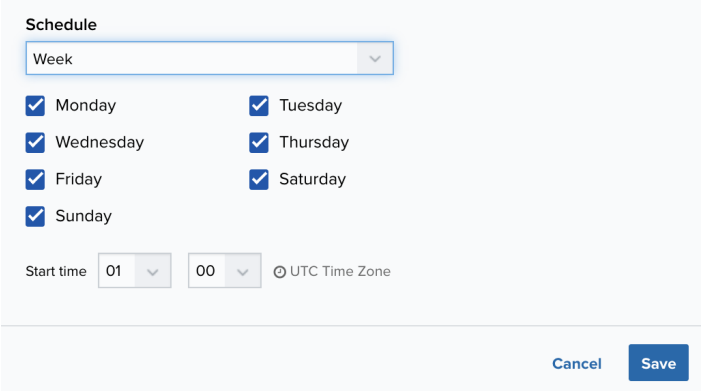

Or on a monthly increment, you can specify a date or a day of the week that occurs within the month.

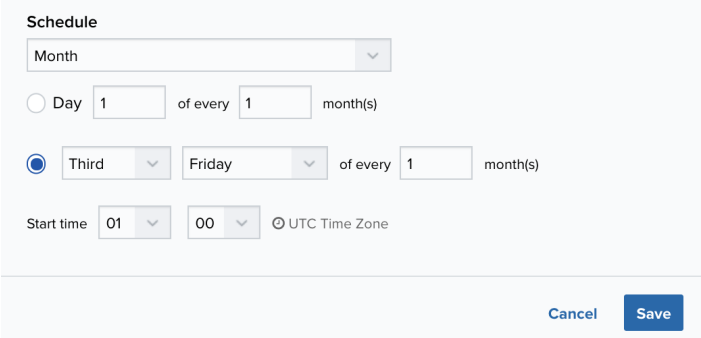

> **Important:** USM Anywhere restarts the schedule on the first day of the month if the option "Every x days" is selected.

c. Set the start time.

This is the time that the job starts at the specified interval. It uses the time zone configured for your USM Anywhere instance (the default is Coordinated Universal Time [UTC]).

11. Click **Save**.

Alarm and Event Notifications

USM Anywhere provides support for direct integration with Slack, Datadog, and PagerDuty as notification methods, as well as integration with Amazon Simple Notification Service (Amazon SNS) to support custom integrations with other messaging services. With direct integration, you can create an orchestration rule that sends notifications to a Slack channel, Datadog event console, or PagerDuty incident management console. With an Amazon SNS integration, you can create an orchestration rule that publishes notification requests to your Amazon SNS for message delivery.

Edition: The Notification integrations are available in the Standard and Premium editions of USM Anywhere.

See the Affordable pricing to fit every budget page for more information about the features and support provided by each of the USM Anywhere editions.

See Orchestration Rules in the *USM Anywhere User Guide* for details about creating orchestration rules.

Before you can create a notification orchestration rule in your environment, you must define one or more of these integrations in USM Anywhere.

This section includes the following topics:

Sending USM Anywhere Notifications to Slack

Role Availability ✗ Read-Only ✗ Investigator ✗ Analyst ✓ **Manager**

From USM Anywhere, you can send an alarm or event notification to a Slack channel to alert team members. This facilitates communication and collaboration within the same messaging tool that your organization uses for incident response. When you have this integration configured in USM Anywhere, you can create orchestration rules to automatically send these notifications when an event or alarm matches the rule criteria.

Edition: The Notification integrations are available in the Standard and Premium editions of USM Anywhere.

See the Affordable pricing to fit every budget page for more information about the features and support provided by each of the USM Anywhere editions.

Note: While the direct integration with USM Anywhere is the easiest and most straightforward way to send messages to your Slack team from USM Anywhere, you can use the Amazon Simple Notification Service (SNS) messaging service as an alternative.

In this case, you create the webhook in Slack and then set up the integration in the Lambda function that you created in Amazon Web Services (AWS) to support USM Anywhere messaging. See Sending Notifications Through Amazon SNS and Set Up a Slack Integration through Amazon SNS for more information.

Create the Slack Webhook

Slack provides a mechanism to create incoming webhooks to post messages from external sources into Slack. They use normal HTTP requests with a JSON payload, which includes the message and some additional options. You must first create this webhook for your Slack team to configure the integration with USM Anywhere.

Important: To add an incoming webhook for the Slack team, you must be the team owner or be a team member where the owner has granted the permission to install apps and custom integrations to all team members. See Sending messages using Incoming Webhooks for more information.

To create the incoming webhook for Slack

1. Log in to your Slack team and go to https://api.slack.com/incoming-webhooks.
2. Review the information and click the **Getting started with Incoming Webhooks** link to open the page for a new configuration.

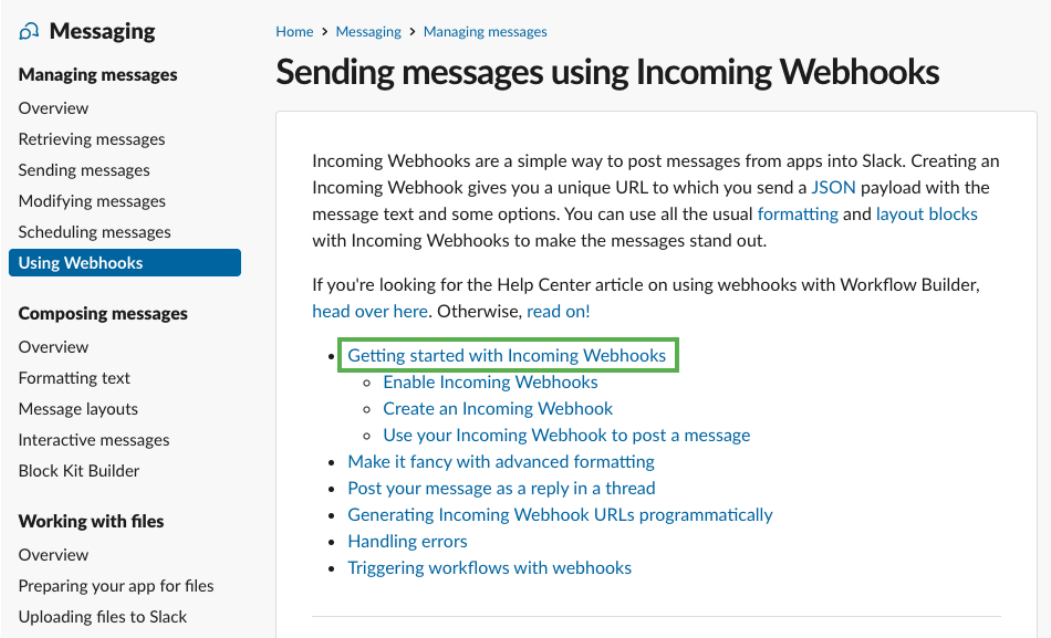

3. Click the **Create your Slack app** button to create a Slack app if you don't have one already.

   The Create an app dialog box opens.

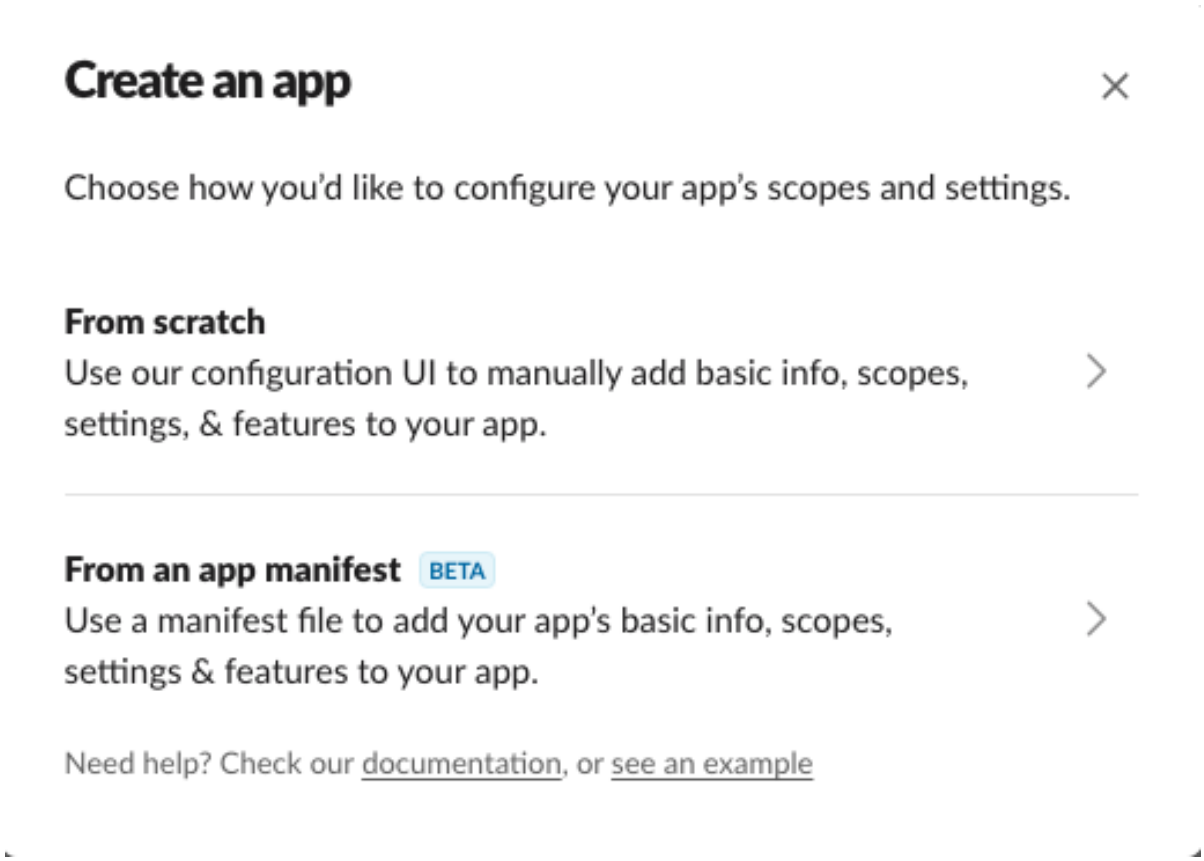

4. Click **From scratch** to create a Slack app if you don't have one already.

   The Create a Slack App dialog box opens.

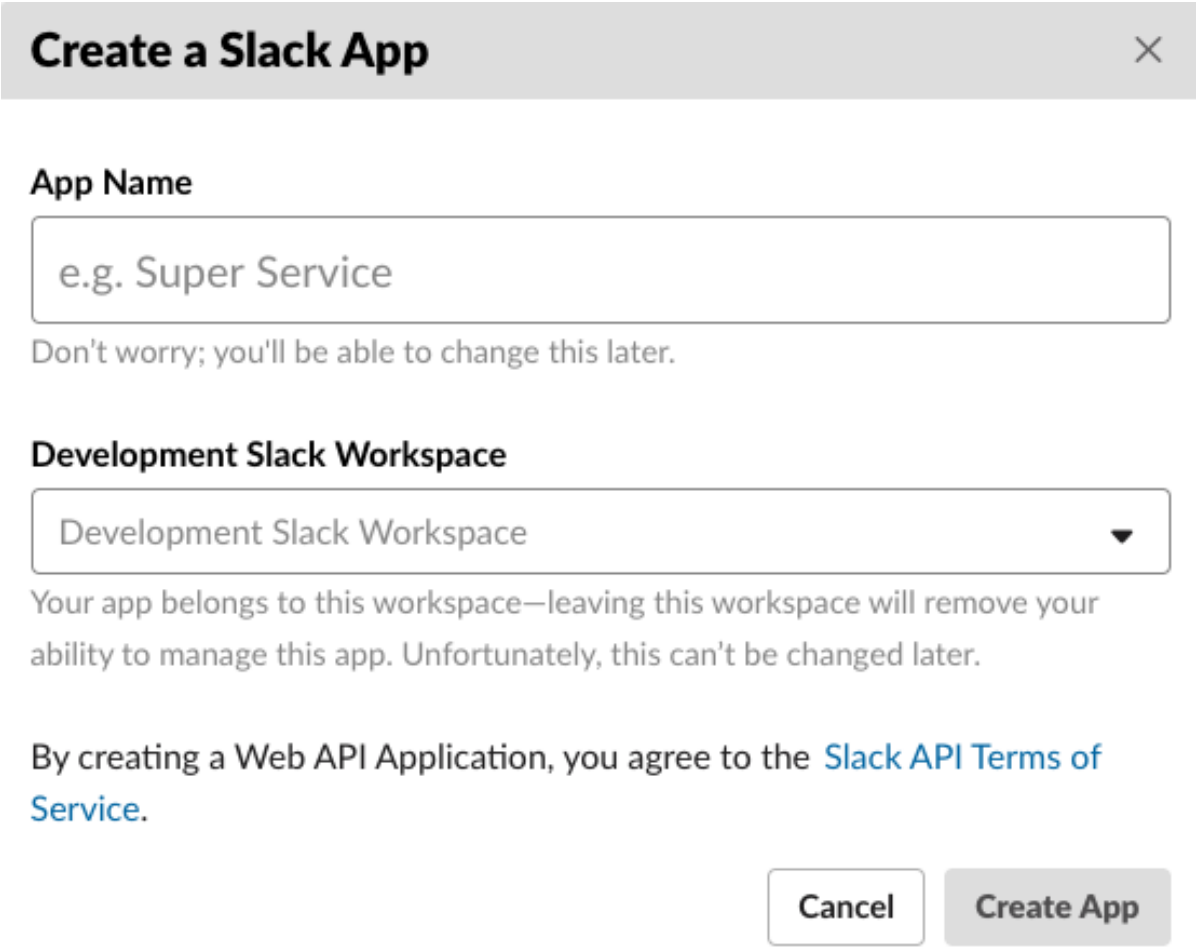

5. Enter a name for the app.
6. Choose a development Slack workspace you want to use for USM Anywhere notifications.

   If you do not already have a channel for this purpose, select the **Sign into a different workspace** option. You can create a new channel, for example, as either a public or private channel and invite the appropriate team members.

7. Click **Create App**.

   The Basic Information page opens.

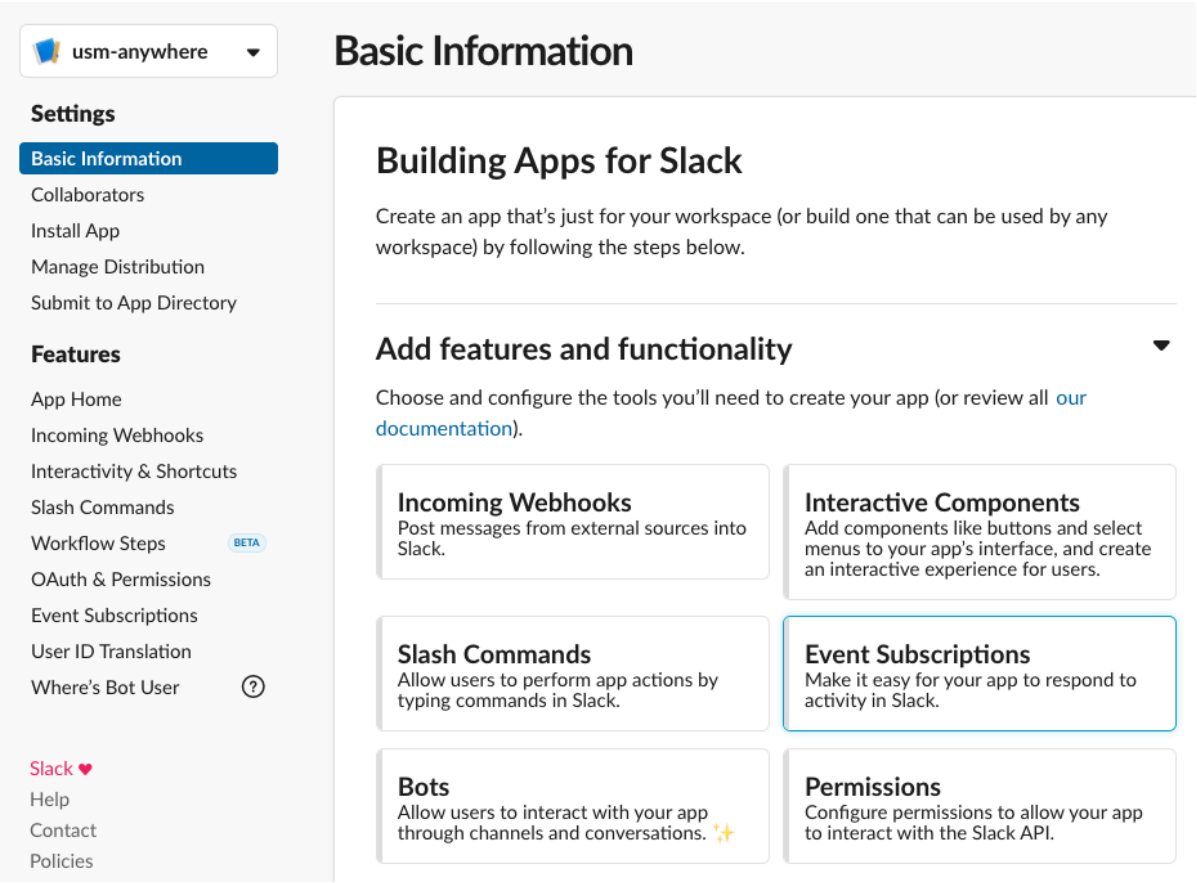

8. Click **Incoming Webhooks**.

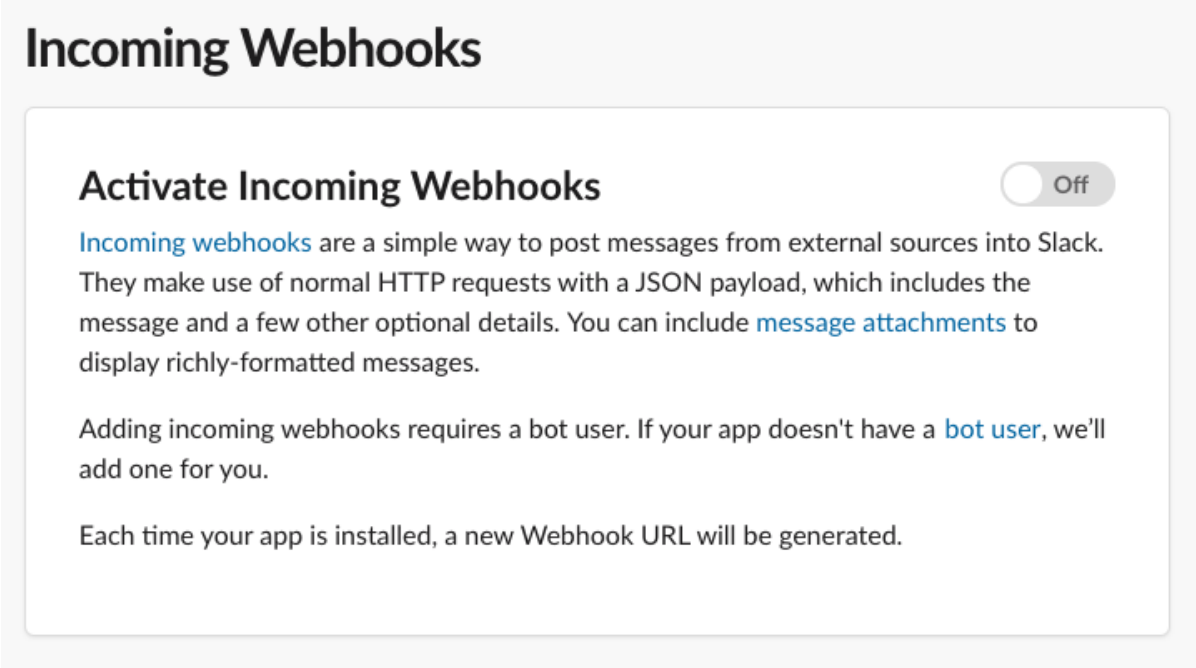

9. Click the off icon to activate the incoming webhooks.

   This turns the icon on and displays it as green.

10. Choose the channel you want to use for USM Anywhere notifications.
11. Click **Request to Add New Webhook** to send the request.
12. When the request is approved, you can see the webhook URL.

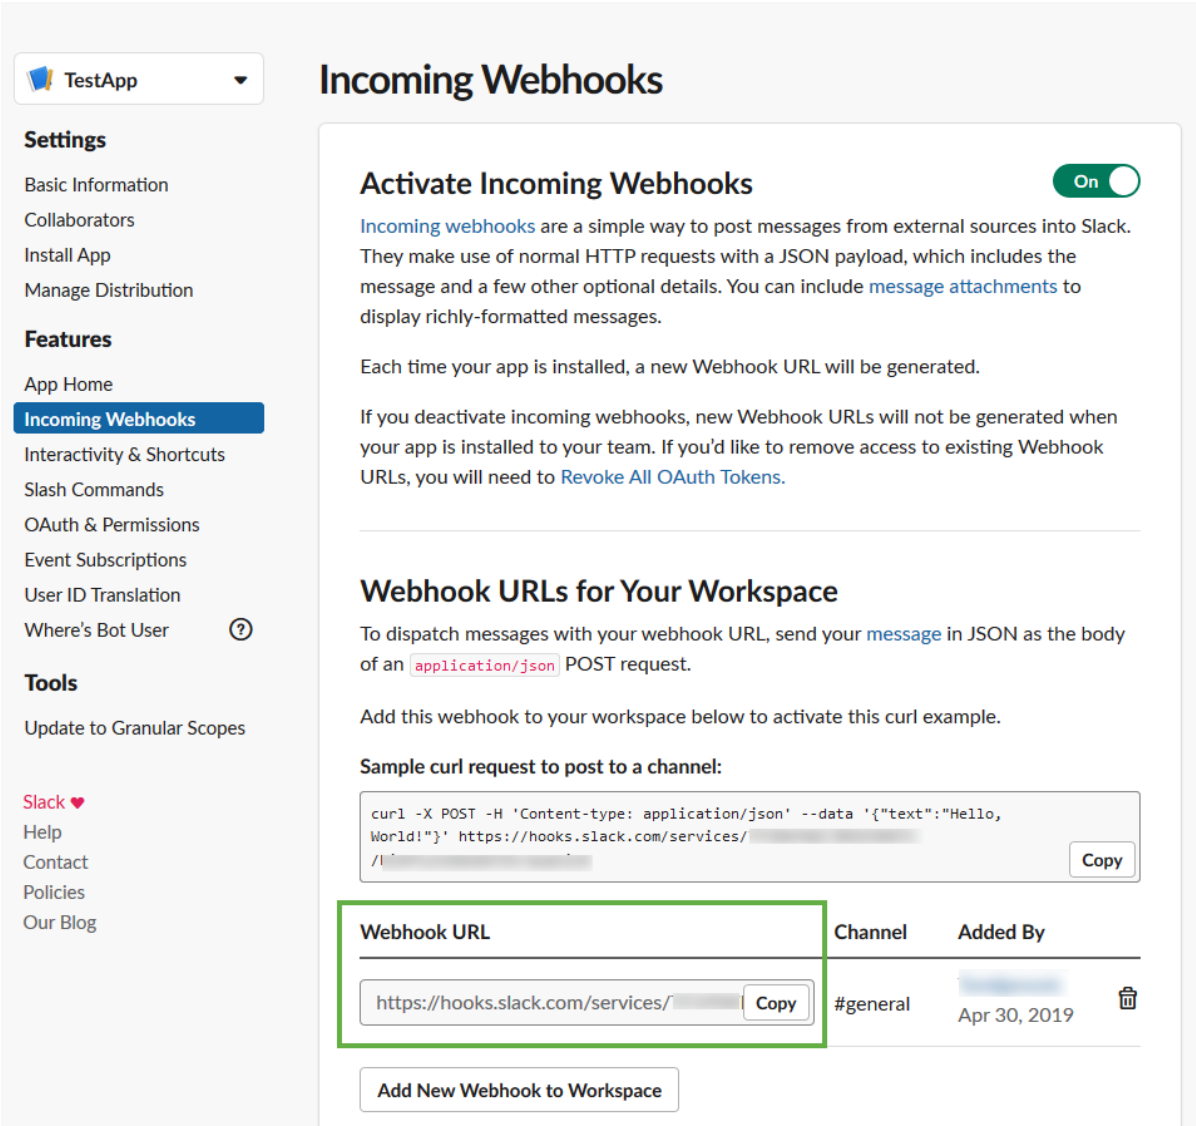

13. Copy the displayed Webhook URL.

Configure the Slack Webhook in USM Anywhere

After you have generated and copied the incoming webhook for your Slack team, you can configure the Slack connection in USM Anywhere. After this configuration is in place, any orchestration rules set up for Slack notification will send the triggered notification to the Slack team channel.

To configure the connection between USM Anywhere and the Slack channel

1. In the USM Anywhere web user interface (UI), go to **Settings > Notifications**.
2. In the left navigation panel, click **Slack**.
3. In the Slack Webhook URL field, paste the webhook URL that you copied in the Slack API tool.

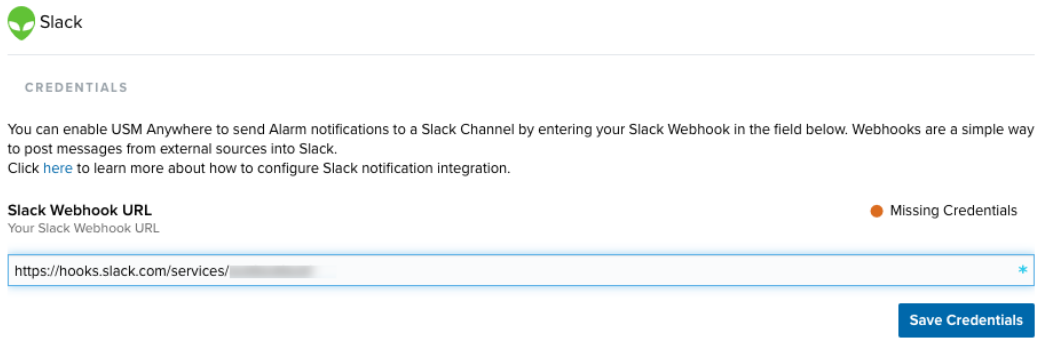

4. Click **Save Credentials**.

Add an Orchestration Rule for Slack Notifications

Create an orchestration rule to match new alarms or events and trigger a notification to the Slack channel. You can use an existing alarm or event with the desired characteristics to easily set the matching conditions for the rule.

To create an orchestration rule to trigger a Slack notification

1. Go to **Activity > Alarms** or **Activity > Events**.
2. Click the alarm or event to open the details.
3. Click **Create Rule** and select **Create Notification Rule**.

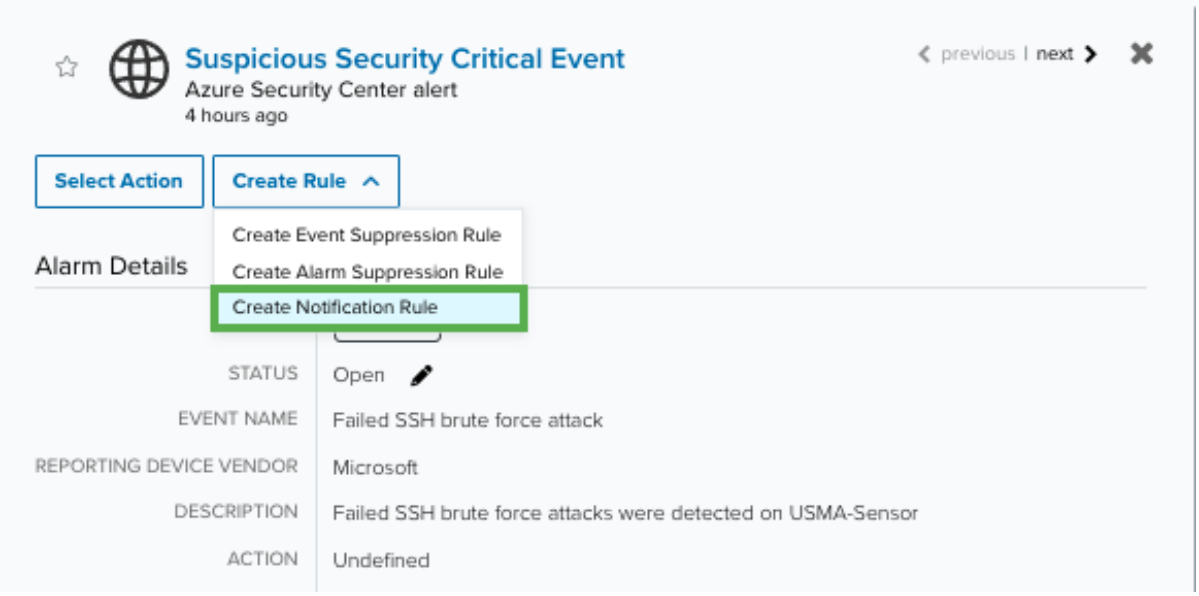

4. You have already suggested property values to create a matching condition, but if you want to add new property values, click **Add Condition**.

> **Note:** If the field is related to the name of a country, you should use the country code defined by the ISO 3166.

> **Note:** The Sources or Destinations field needs to match the universally unique identifier (UUID) of the event or alarm. You can use the Source Name or Destination Name field instead.

> **Important:** Instead of using the `equals` and `equals, case insensitive` operators for array fields, AT&T Cybersecurity recommends the use of the `in` or `contains` operators.

> **Note:** If you need to add a property value that maps with a property key, you need to know the mapping of the field. See Determining the Mapping of a Field for more information.

1. (Optional.) Click **Add Group** to group your conditions.

> **Note:** See Operators in the Orchestration Rules for more information.

Note: The current rule box shows you the syntax of your rule, and the rule verification box reviews that syntax before saving the rule.

2. Click **Next**.

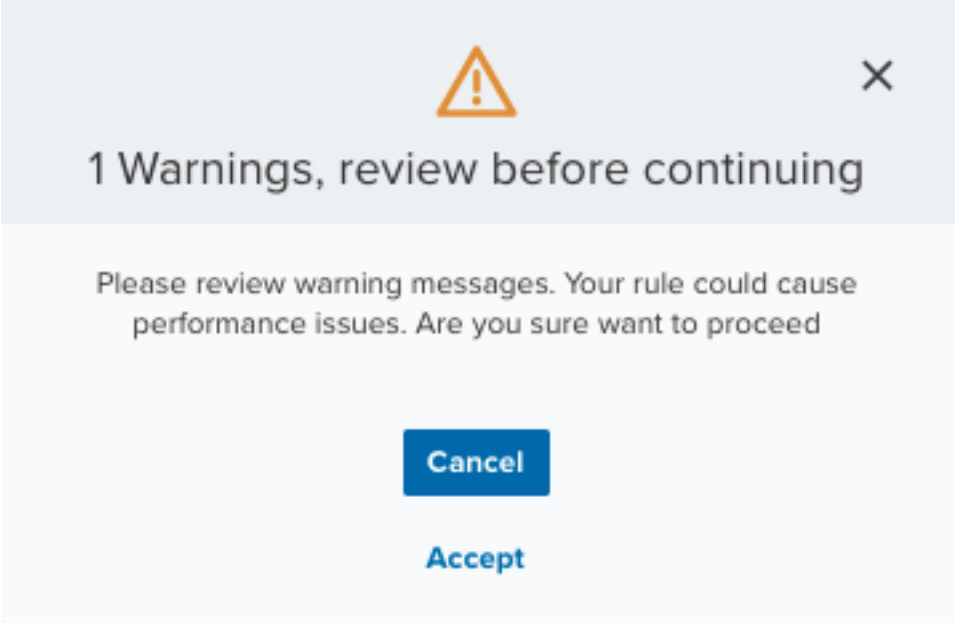

Important: A dialog box opens if there are warning messages. Click **Cancel** to review the warning messages, or click **Accept** to continue creating the rule.

3. Enter a name for the rule.
4. (Optional.) Enter a description for identifying this rule.
5. For Notification Method, select the **Slack** option.
6. Enter the Slack Alert Username.

   The username must be a valid team member for the Slack channel.

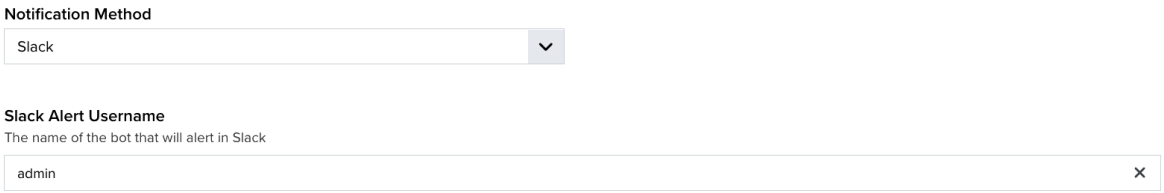

7. Modify these two options:

- **Occurrences:** Specify the number of event occurrences that produce a match on the conditional expression to trigger the rule. You can enter the number of occurrences or use the arrow to scroll the value up or down. You need to enter a number between 1 and 100.
- **Length:** Specify the length of the timespan used to identify a match for multiple occurrences. Enter the number and choose a value of seconds, minutes, or hours.

  This duration identifies the amount of time that transpires from the beginning to the end of the occurrence. If the number of occurrences is not met within this period, the rule is not a match.

  In this example, the rule applies when the configured conditions happen five times every three hours.

These two options function together to specify the number of occurrences within a time period that will produce a match for the rule. For example, you can define a rule to trigger an alarm for an unauthorized access attempt when a failed SSH login occurs three times within a five-minute window.

8. Click **Save**.

   The created rule displays in the list of rules. You can see it from **Settings > Rules > Orchestration Rules**. See Orchestration Rules for more information.

   Important: It takes a few minutes for an orchestration rule to become active.

Set Up a Slack Integration through Amazon SNS

If you prefer to use Amazon SNS to forward notifications to your Slack channel, you can add the webhook that you created to the Lambda function in your AWS account.

Important: For this integration type, you do not add the Slack webhook in USM Anywhere. When you create the orchestration rule, you select the Amazon SNS notification method.

Before you can complete this integration, you must have an SNS topic and a Lambda function for USM Anywhere notifications set up in your AWS account (see Set Up an Amazon SNS Topic) and a Slack incoming webhook (see Create the Slack Webhook).

To integrate the Slack webhook with the USM Anywhere through Amazon SNS

1. In the Lambda function code, paste this code and replace [INSERT_WEBHOOK_URL] with the Slack webhook URL.
2. Use the default Role setting (*Create a new role from templates*) and specify the Role name as `lambda_basic_execution`.

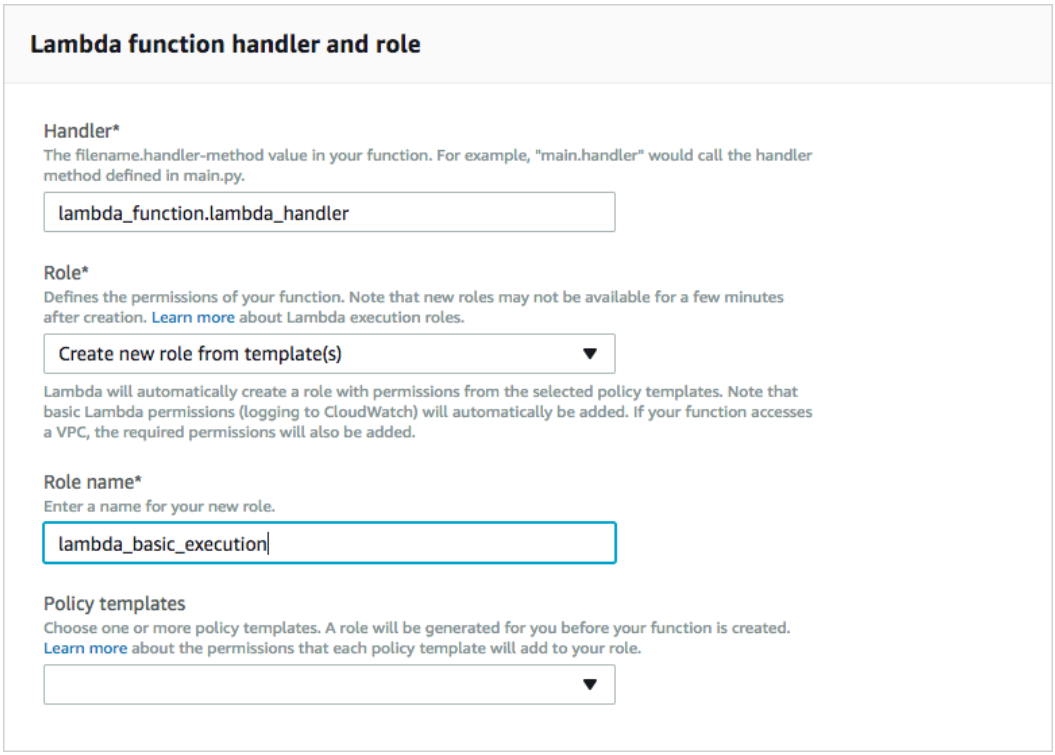

3. Expand the Advanced settings and set the Timeout to **10 seconds**.
4. Click **Next**.
5. Click **Create function**.

To check the integration with Slack

1. Go to your Lambda function, click **Monitoring**, and verify the Invocation Count graph data.
2. Check your Slack channel for notifications.

Sending USM Anywhere Notifications to Datadog

Role Availability ✗ Read-Only ✗ Investigator ✗ Analyst ✓ **Manager**

From USM Anywhere, you can send an alarm or event notification to your Datadog event console so that team members are alerted. This facilitates communication and collaboration within the same messaging tool that your organization uses for infrastructure monitoring. When you have this integration configured in USM Anywhere, you can create orchestration rules to automatically send these notifications when an event or alarm matches the rule criteria.

Edition: The Notification integrations are available in the Standard and Premium editions of USM Anywhere.

See the Affordable pricing to fit every budget page for more information about the features and support provided by each of the USM Anywhere editions.

Note: While direct integration with USM Anywhere is the easiest and most straightforward way to send messages to your Datadog environment from USM Anywhere, you can use the Amazon SNS messaging service as an alternative. In this case, you create the API key in Datadog and then set up the integration in the Lambda function that you created in AWS to support USM Anywhere messaging. S ee Sending Notifications Through Amazon SNS and Set Up a Datadog Events Integration Through Amazon SNS for more information.

Create a Datadog API Key

Datadog provides a mechanism to create API keys as a way to post data from external sources into Datadog events. All requests to the Datadog API must be authenticated. Requests that write data require reporting access and require an API key. You must first create this API key to configure the integration with USM Anywhere.

To create the API key for Datadog

1. Log in to your Datadog account and go to https://app.datadoghq.com/account/settings#api.
2. For the **New API key**, enter a name for the key and click **Create API key**.

   Make sure to copy the generated key value and store it in a secured location.

3. (Amazon SNS Only.) For the **New application key**, click **Create Application key** and copy the generated value.

> **Note:** This key is not used for a direct integration with USM Anywhere. However, if you plan to use the Amazon SNS messaging service for a custom integration, any requests that read data require full access and an application key.

Configure the Datadog API Key in USM Anywhere

After you have generated and copied the API key for your Datadog environment, you can configure USM Anywhere for Datadog notifications. After this configuration is in place, any orchestration rules set up for Datadog notification will send the triggered notification to your Datadog events.

To configure the connection between Datadog events and USM Anywhere

1. In the USM Anywhere web UI, go to **Settings > Notifications**.
2. In the left navigation panel, click **Datadog**.
3. In the Datadog API key field, paste the key value that you generated in the Datadog API tool.

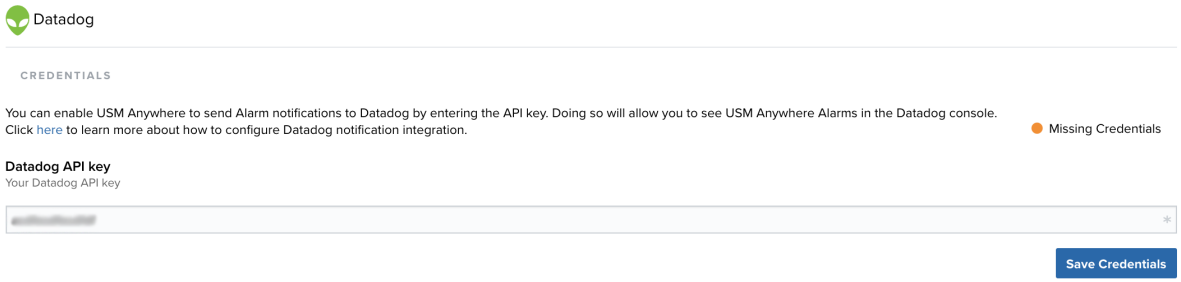

4. Click **Save Credentials**.

Add an Orchestration Rule for Datadog Notifications

Create an orchestration rule to match new alarms or events and trigger a notification to Datadog events. You can use an existing alarm or event with the desired characteristics to easily set the matching conditions for the rule.

To create an orchestration rule to trigger a Datadog notification

1. Go to **Activity > Alarms** or **Activity > Events**.
2. Click the alarm or event to open the details.
3. Click **Create Rule** and select **Create Notification Rule**.

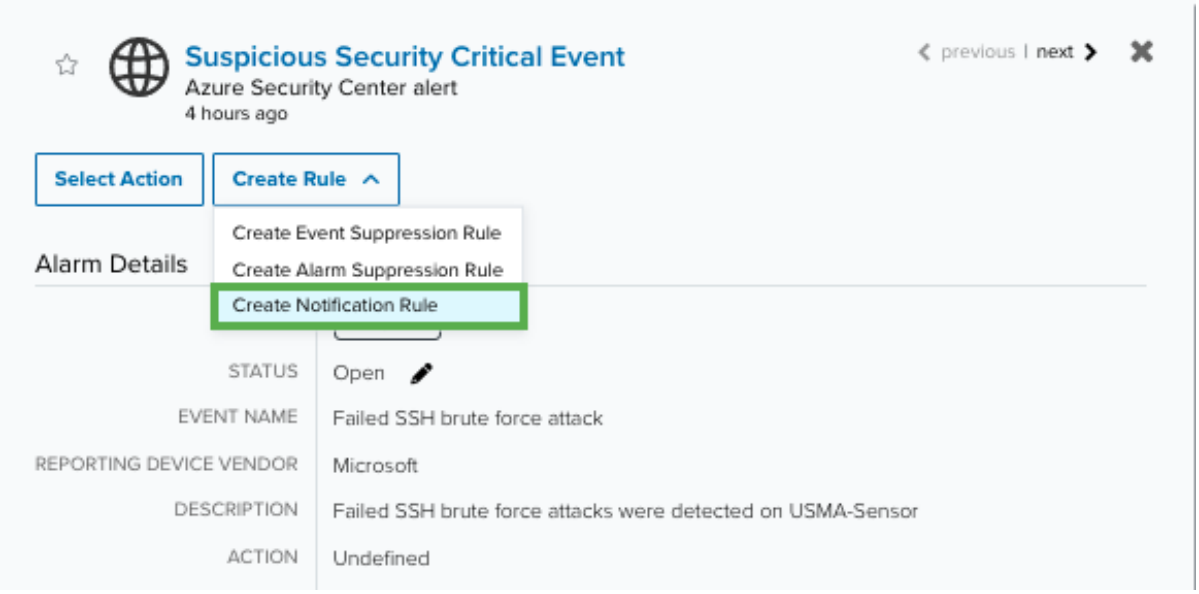

4. You have already suggested property values to create a matching condition, but if you want to add new property values, click **Add Condition**.

> **Note:** If the field is related to the name of a country, you should use the country code defined by the ISO 3166.

> **Note:** The Sources or Destinations field needs to match the universally unique identifier (UUID) of the event or alarm. You can use the Source Name or Destination Name field instead.

> **Important:** Instead of using the `equals` and `equals, case insensitive` operators for array fields, AT&T Cybersecurity recommends the use of the `in` or `contains` operators.

> **Note:** If you need to add a property value that maps with a property key, you need to know the mapping of the field. See Determining the Mapping of a Field for more information.

1. (Optional.) Click **Add Group** to group your conditions.

> **Note:** See Operators in the Orchestration Rules for more information.

Note: The current rule box shows you the syntax of your rule, and the rule verification box reviews that syntax before saving the rule.

2. Click **Next**.

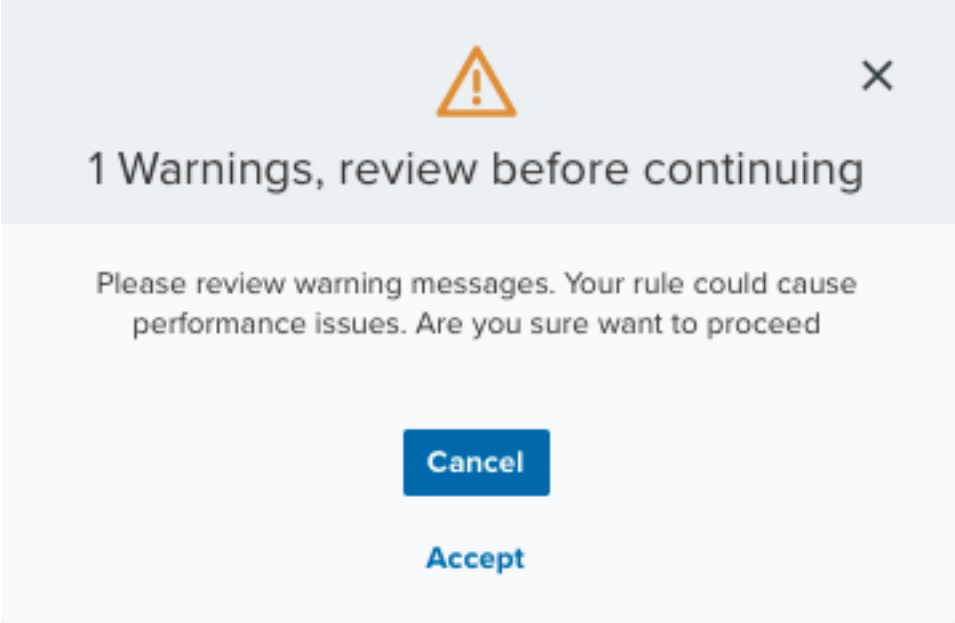

Important: A dialog box opens if there are warning messages. Click **Cancel** to review the warning messages, or click **Accept** to continue creating the rule.

3. Enter a name for the rule.
4. (Optional.) Enter a description for identifying this rule.
5. For Notification Method, select the **Slack** option.
6. Enter the Slack Alert Username.

   The username must be a valid team member for the Slack channel.

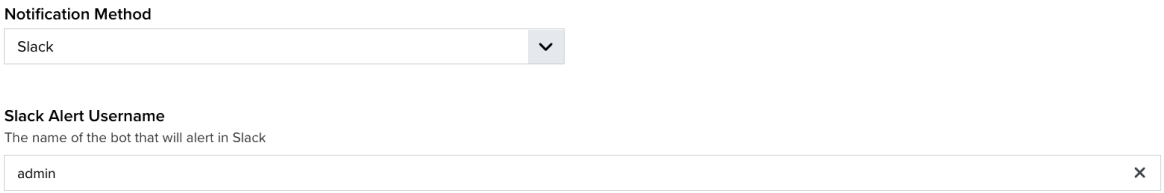

7. Modify these two options:

- **Occurrences:** Specify the number of event occurrences that produce a match on the conditional expression to trigger the rule. You can enter the number of occurrences or use the arrow to scroll the value up or down. You need to enter a number between 1 and 100.
- **Length:** Specify the length of the timespan used to identify a match for multiple occurrences. Enter the number and choose a value of seconds, minutes, or hours.

  This duration identifies the amount of time that transpires from the beginning to the end of the occurrence. If the number of occurrences is not met within this period, the rule is not a match.

  In this example, the rule applies when the configured conditions happen five times every three hours.

These two options function together to specify the number of occurrences within a time period that will produce a match for the rule. For example, you can define a rule to trigger an alarm for an unauthorized access attempt when a failed SSH login occurs three times within a five-minute window.

8. Click **Save**.

   The created rule displays in the list of rules. You can see it from **Settings > Rules > Orchestration Rules**. See Orchestration Rules for more information.

   Important: It takes a few minutes for an orchestration rule to become active.

5. For **Notification Method**, select the **Datadog** option.

6. Set the **Datadog Priority**.

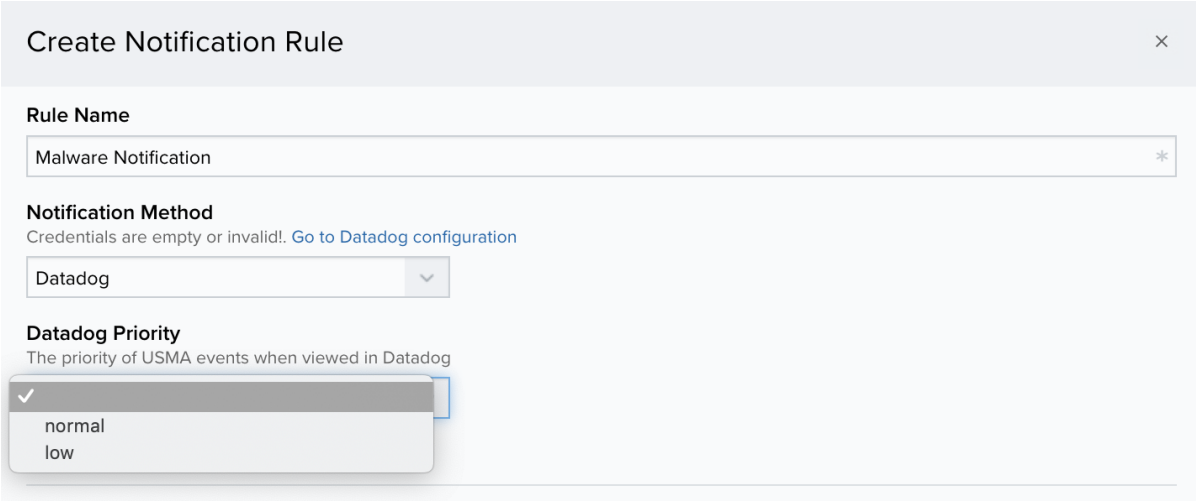

7. At the bottom of the dialog box, set the rule condition parameters to specify the criteria for a matching alarm or event to trigger the rule.

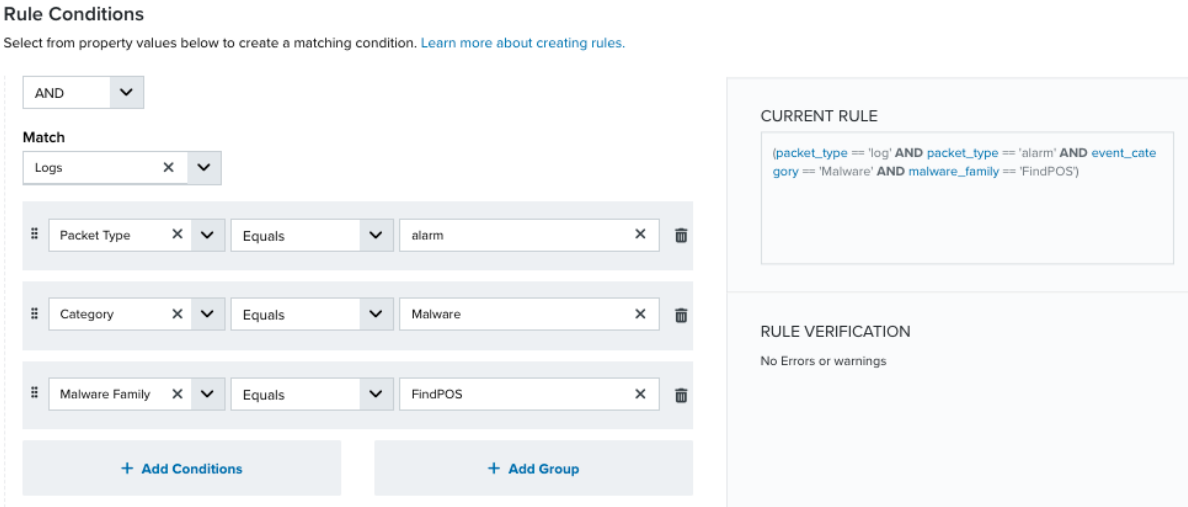

- This section provides suggested property/value pairs from the selected alarm or event that you can use as conditions for the rule. Click the icon to delete the items that you do not want to include in the matching conditions. You can also add other conditions that are not suggested.
- If you create the rule from the Rules page, you must use the Add Condition and Add Group functions to define the property/value pairs that you want to use as conditions

for the rule.

- At the bottom of the dialog box, click **More** to display the optional multiple occurrence and window-length parameters.

Conditional Expression

Select an operator and add one or more conditions to form the conditional expression. You can include a condition group to evaluate a subset of conditions. The Current Rule pane displays the constructed expression in standard syntax. The box displays a red border if the expression is syntactically invalid as currently specified. A valid expression is required to save the rule definition.

Select the operator used to determine the match for multiple conditions:

- **AND**: Match all conditions.
- **OR**: Match any one condition.
- **AND NOT**: Exclude items matching all conditions after the first.
- **OR NOT**: Include all items that do not match any conditions after the first.

Click **Add Condition** to add a condition. For each condition, specify the field name, evaluator, and value. If the evaluation returns *true* for the condition, it is a match.

Click **Add Group** to add a condition group. A new group includes a condition and its own operator used to match the conditions within the group. You can nest condition groups.

Occurrences

Specify the number of event or alarm occurrences that produce a match on the conditional expression to trigger the rule. The default value is 1. You can enter the number of occurrences or use the arrow to scroll the value up or down.

USM Anywhere uses this in conjunction with the Length option to specify the number of occurrences within a time period that will trigger the rule. For example, you can define a rule to trigger for an unauthorized access attempt when a failed SSH login occurs three times within a five-minute window.

Length

Specify the length of the window to identify a match for multiple occurrences. Enter the number and choose a time unit value of seconds, minutes, or hours. This time period identifies the amount of time that transpires from the first occurrence to the last occurrence. If the number of occurrences is not met within this period, the rule does not

trigger.

8. Click **Save Rule**.

Set Up a Datadog Events Integration Through Amazon SNS

If you prefer to use Amazon SNS to forward notifications to your Datadog Events, you can add the API key to the Lambda function in your AWS account.

> **Important:** For this integration type, you do not add the Datadog API key in USM Anywhere. When you create the orchestration rule, you select the Amazon SNS notification method.

Before you can complete this integration, you must have an SNS topic and a Lambda Function for USM Anywhere notifications set up in your AWS account (see Set Up an Amazon SNS Topic) and a Datadog API key (see Create a Datadog API Key).

To integrate USM Anywhere notifications with Datadog Events through Amazon SNS

1. In the Lambda function code, paste this code and replace [INSERT_DATADOG_API_KEY] and [INSERT_DATADOG_APPLICATION_KEY] with your Datadog keys.

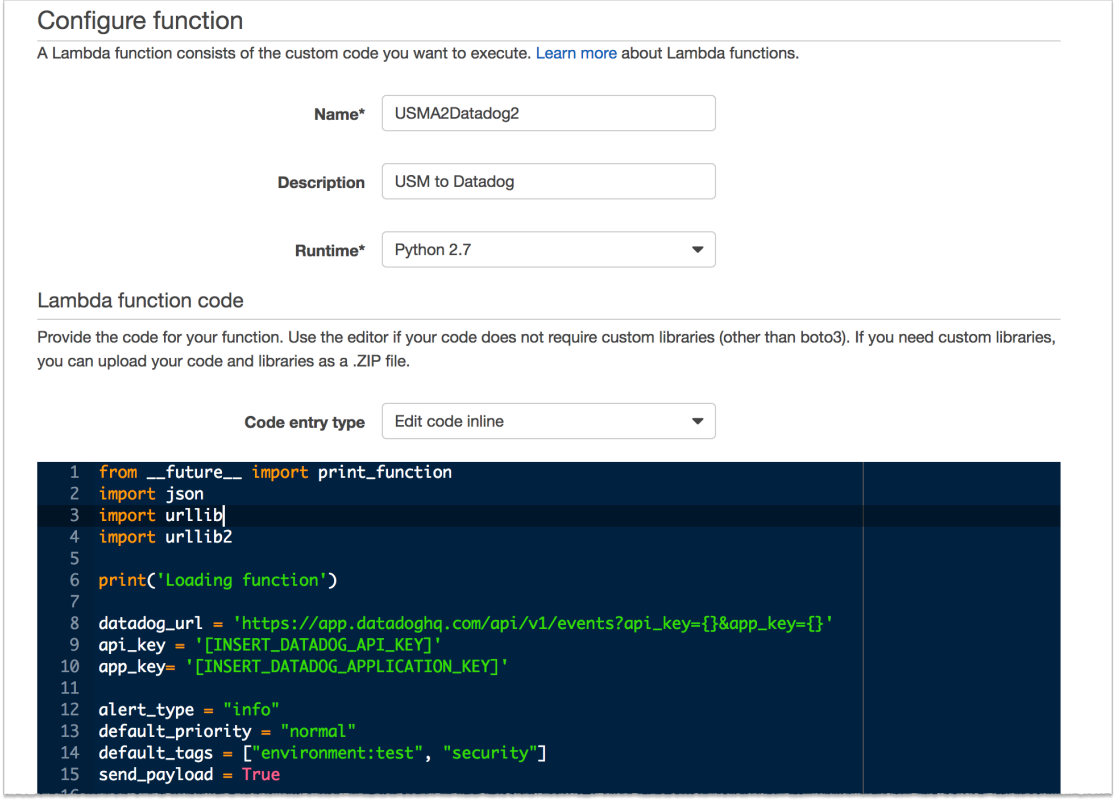

You can also modify the Datadog fields and adapt them to your environment, similar to the following:

```
alert_type = "info"
default_priority = "normal"
default_tags = ["environment:test", "security"]
send_payload = True
```

2. Use the default **Role** setting (Create a new role from templates) and specify the **Role name** as `lambda_basic_execution`.

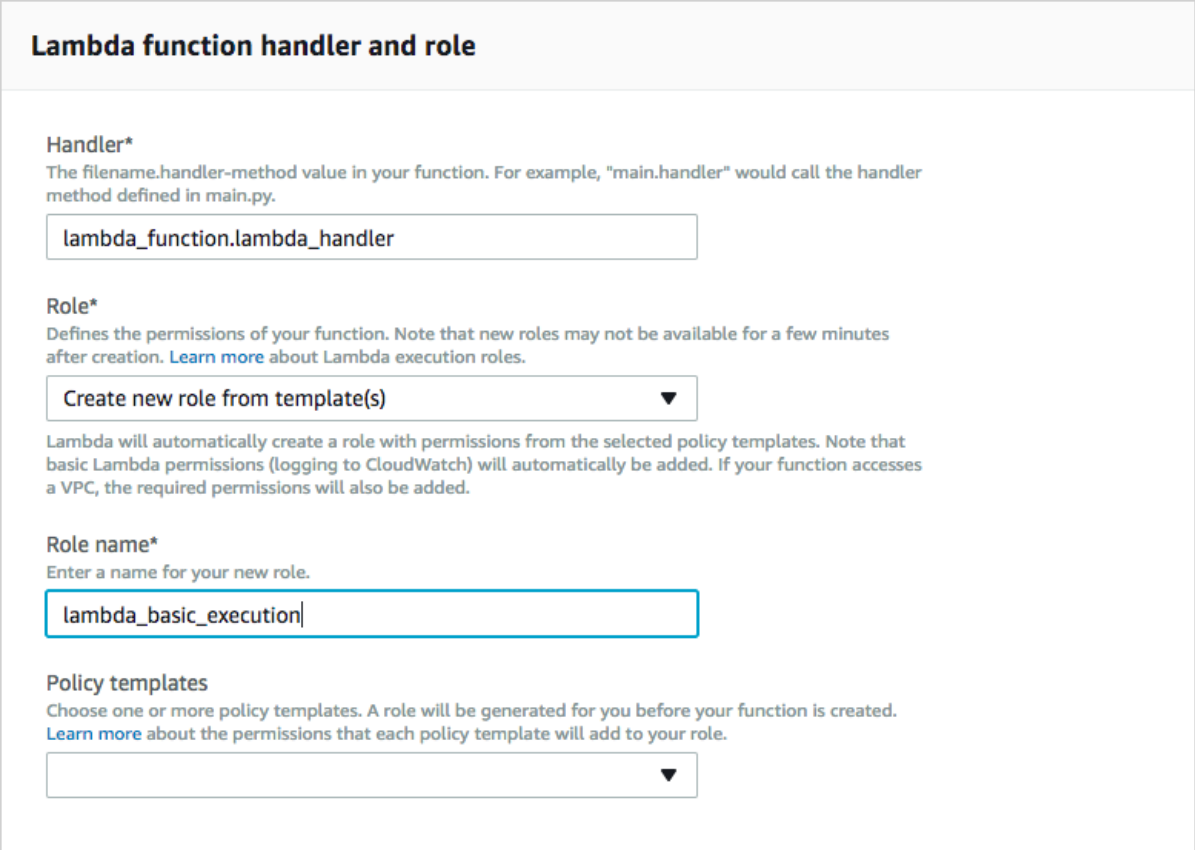

3. Expand the **Advanced** settings and set the **Timeout** to 10 seconds.
4. Click **Next**.
5. Click **Create function**.

To check the integration with Datadog

1. Go to your Lambda function, click **Monitoring**, and verify that the Invocation Count graph shows some data.
2. Click **View logs in CloudWatch** and open the last entry.

You should see entries similar to the following:

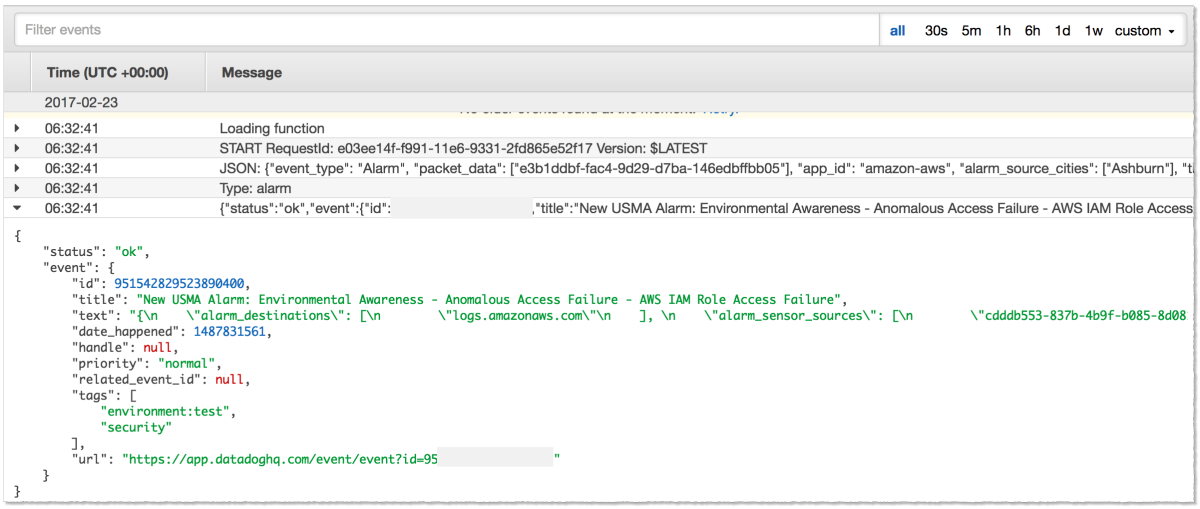

3. Go to the Datadog event URL and check that you see the USM Anywhere alarm in the Datadog console.

Sending USM Anywhere Notifications to PagerDuty

Role Availability ✗ Read-Only ✗ Investigator ✗ Analyst ✓ **Manager**

From USM Anywhere, you can send an alarm or event notification to your PagerDuty incident management console so that team members receive alerts. This facilitates communication and collaboration within the same messaging tool that your organization uses for incident response. When you have this integration configured in USM Anywhere, you can create orchestration rules to automatically send these notifications when an event or alarm matches the rule criteria.

Edition: The Notification integrations are available in the Standard and Premium editions of USM Anywhere.

See the Affordable pricing to fit every budget page for more information about the features and support provided by each of the USM Anywhere editions.

Create the PagerDuty Integration

PagerDuty provides a mechanism to create services that include integrations to its Events API as a way to post data from external sources into PagerDuty incidents. The service configuration determines how PagerDuty handles the incoming incident. You must first create the integration key for a PagerDuty service before you set up the configuration in USM Anywhere to send these notifications.

Note: A PagerDuty service typically represents an application, component, or team for opening incidents. If you already have a defined service and you want to incorporate USM Anywhere notifications with it, you can simply add a new integration to that service and use the parameters outlined in the following procedure.

To create a PagerDuty service and integration for USM Anywhere

1. Log in to your PagerDuty account.
2. In the top menu, select **Configuration > Services**.
3. At the top of the page, click **Add New Service**.
4. In the General Settings, enter a name for the new service (such as *AT&T Cybersecurity*).
5. In Integration Settings, set the type and name for the integration.
   - Choose **Use our API Directly** and select **Events API v2**.
   - Enter an Integration Name (such as *USM Anywhere*).

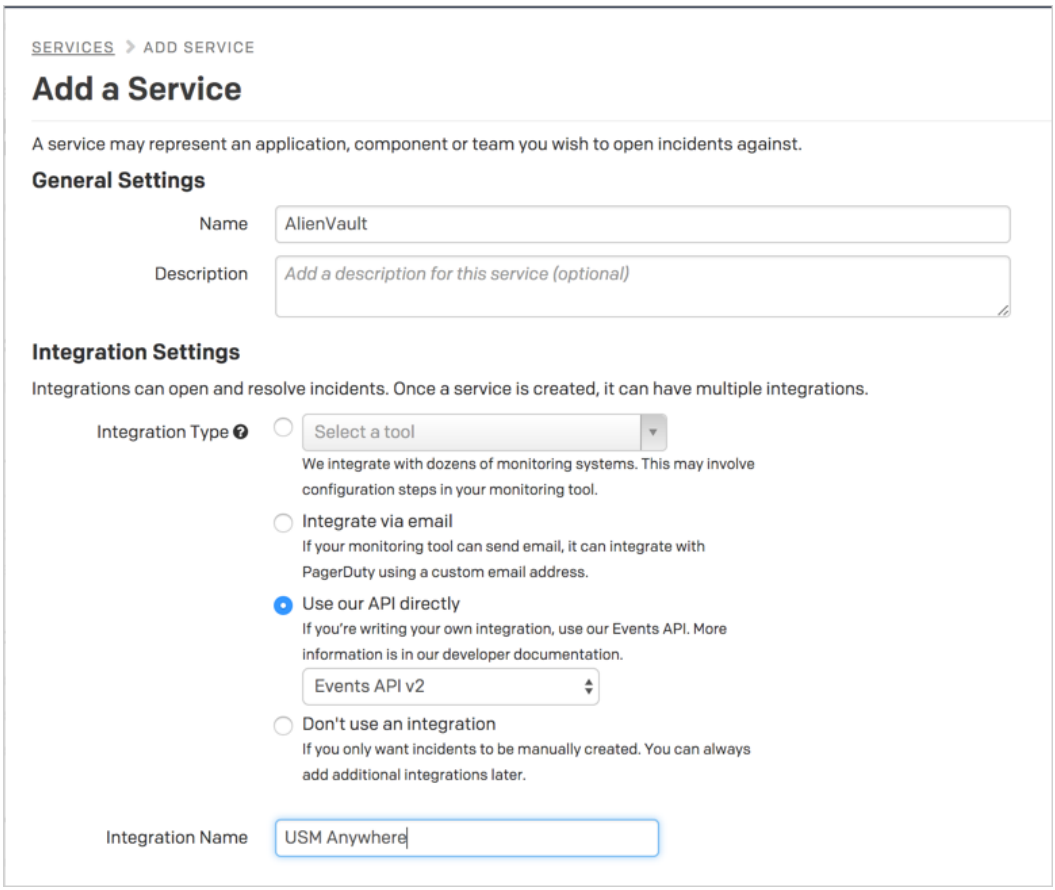

6. Set the Incident Settings and Incident Behavior according to how you want PagerDuty to handle the incidents (notifications) from USM Anywhere.
7. Click **Add Service**.
8. In the Integrations tab, copy the Integration Key for the new integration.

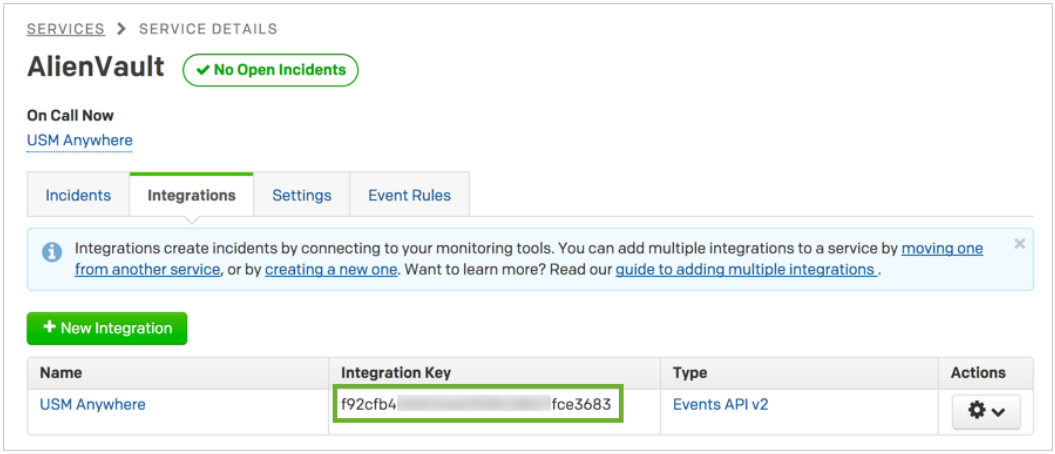

Make sure to copy the key value to a secured location.

Configure the PagerDuty Integration in USM Anywhere

After you have created the PagerDuty integration and copied the key, you can configure USM Anywhere for PagerDuty notifications. After this configuration is in place, any orchestration rules set up for PagerDuty notification will send the triggered notification to the PagerDuty service for incident handling.

To configure the connection between PagerDuty and USM Anywhere

1. In the USM Anywhere web UI, go to **Settings > Notifications**.
2. In the left navigation panel, click **PagerDuty**.
3. In the PagerDuty integration key field, paste the key value that you copied from your PagerDuty service integration.

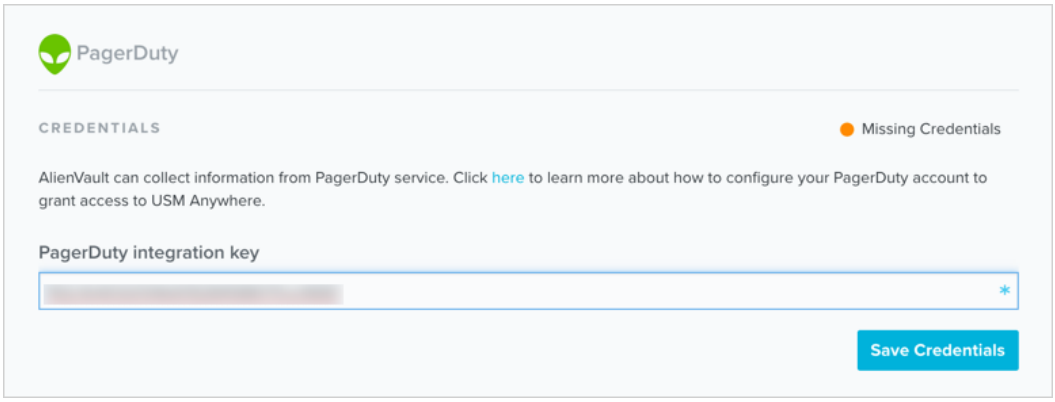

4. Click **Save Credentials**.

Add an Orchestration Rule for PagerDuty Notifications

Create an orchestration rule to match new alarms or events and trigger a notification to your PagerDuty service. You can use an existing alarm or event with the desired characteristics to easily set the matching conditions for the rule.

To create an orchestration rule to trigger a PagerDuty notification

1. Go to **Activity > Alarms** or **Activity > Events**.
2. Click the alarm or event to open the details.
3. Click **Create Rule** and select **Create Notification Rule**.

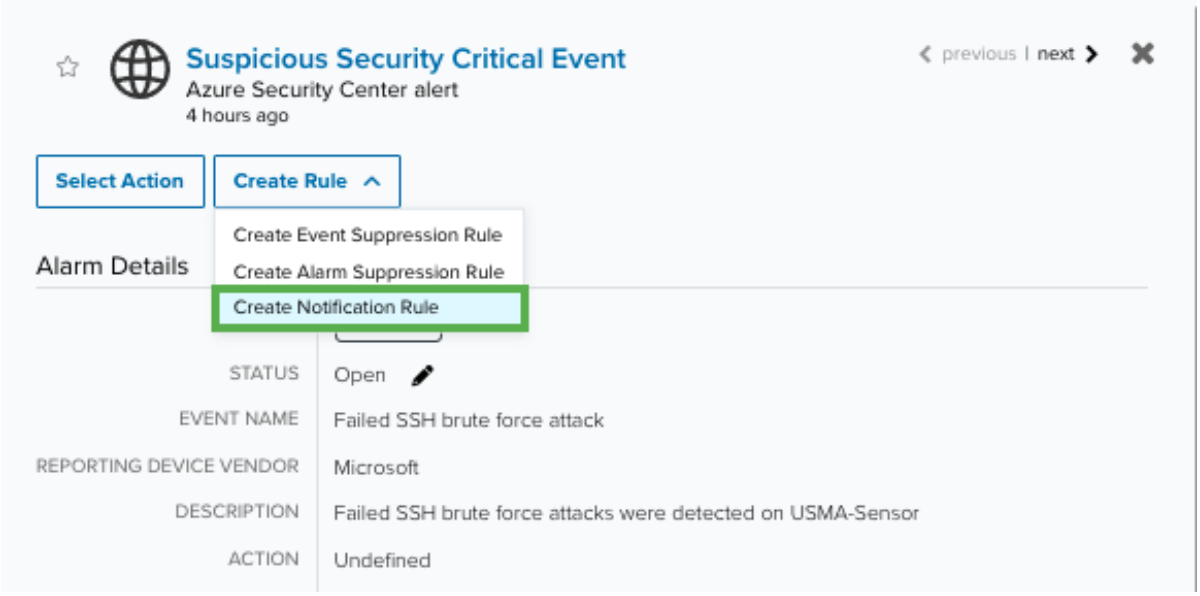

4. You have already suggested property values to create a matching condition, but if you want to add new property values, click **Add Condition**.

> **Note:** If the field is related to the name of a country, you should use the country code defined by the ISO 3166.

> **Note:** The Sources or Destinations field needs to match the universally unique identifier (UUID) of the event or alarm. You can use the Source Name or Destination Name field instead.

> **Important:** Instead of using the `equals` and `equals, case insensitive` operators for array fields, AT&T Cybersecurity recommends the use of the `in` or `contains` operators.

> **Note:** If you need to add a property value that maps with a property key, you need to know the mapping of the field. See Determining the Mapping of a Field for more information.

1. (Optional.) Click **Add Group** to group your conditions.

> **Note:** See Operators in the Orchestration Rules for more information.

Note: The current rule box shows you the syntax of your rule, and the rule verification box reviews that syntax before saving the rule.

2. Click **Next**.

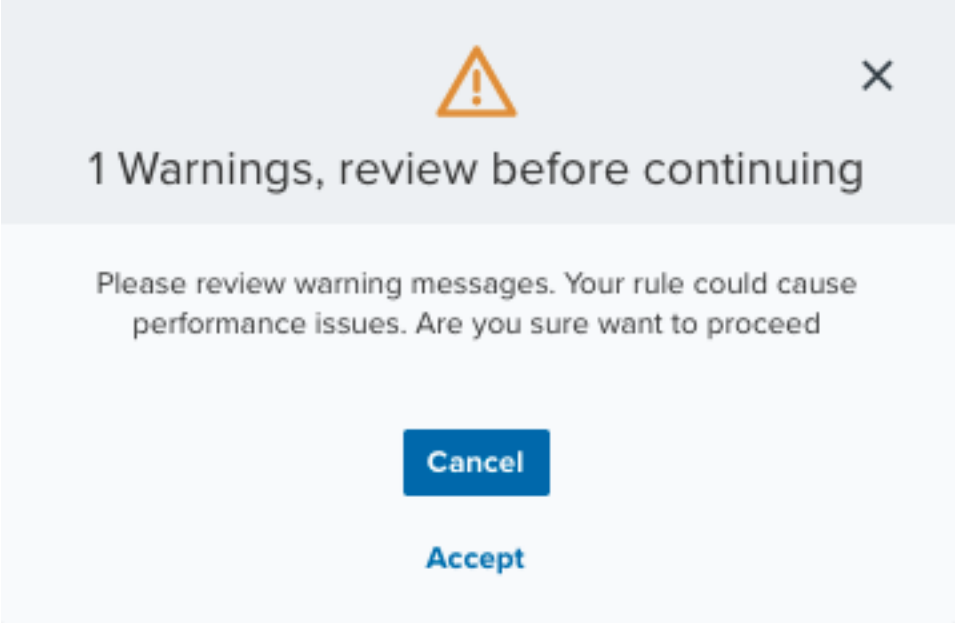

Important: A dialog box opens if there are warning messages. Click **Cancel** to review the warning messages, or click **Accept** to continue creating the rule.

3. Enter a name for the rule.
4. (Optional.) Enter a description for identifying this rule.
5. For Notification Method, select the **Slack** option.
6. Enter the Slack Alert Username.

   The username must be a valid team member for the Slack channel.

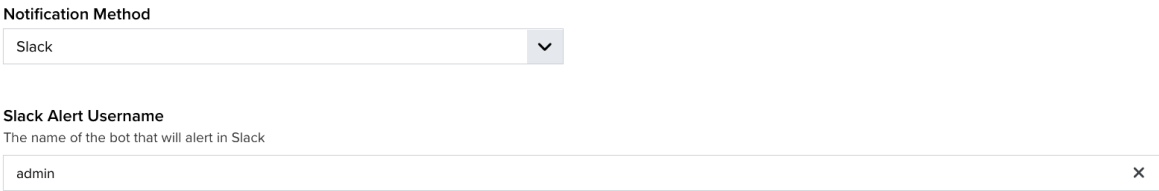

7. Modify these two options:

- **Occurrences:** Specify the number of event occurrences that produce a match on the conditional expression to trigger the rule. You can enter the number of occurrences or use the arrow to scroll the value up or down. You need to enter a number between 1 and 100.
- **Length:** Specify the length of the timespan used to identify a match for multiple occurrences. Enter the number and choose a value of seconds, minutes, or hours.

  This duration identifies the amount of time that transpires from the beginning to the end of the occurrence. If the number of occurrences is not met within this period, the rule is not a match.

  In this example, the rule applies when the configured conditions happen five times every three hours.

These two options function together to specify the number of occurrences within a time period that will produce a match for the rule. For example, you can define a rule to trigger an alarm for an unauthorized access attempt when a failed SSH login occurs three times within a five-minute window.

8. Click **Save**.

   The created rule displays in the list of rules. You can see it from **Settings > Rules > Orchestration Rules**. See Orchestration Rules for more information.

   Important: It takes a few minutes for an orchestration rule to become active.

5. For **Notification Method**, select the **PagerDuty** option.

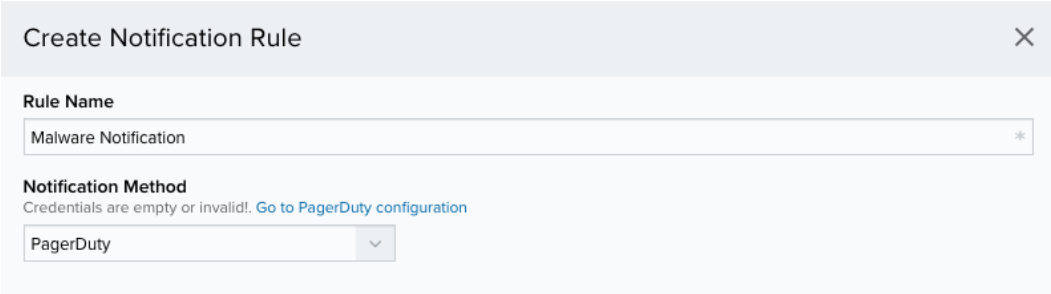

6. At the bottom of the dialog box, set the rule condition parameters to specify the criteria for a matching alarm or event to trigger the rule.

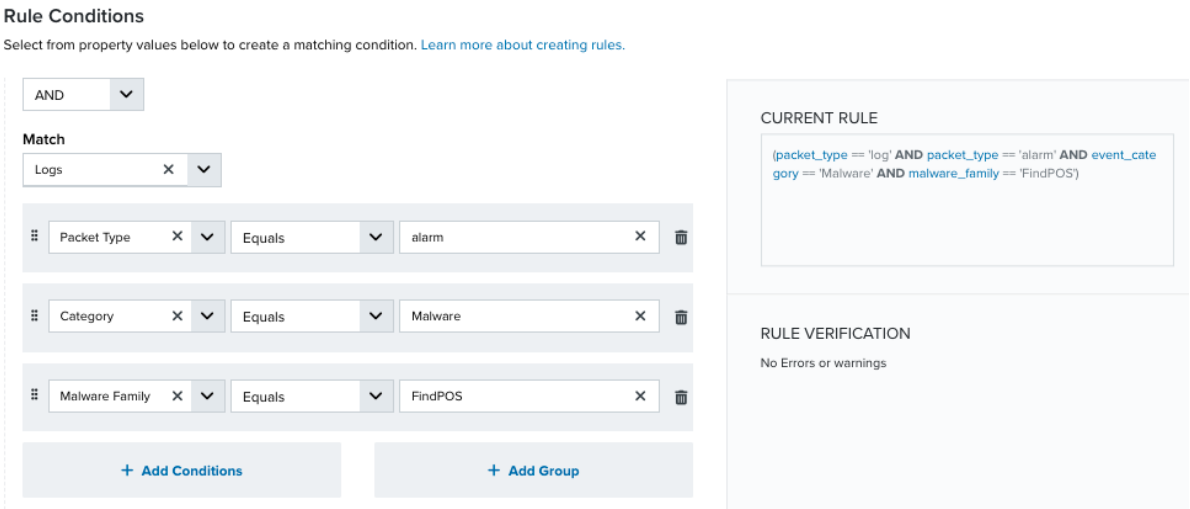

- This section provides suggested property/value pairs from the selected alarm or event that you can use as conditions for the rule. Click the icon to delete the items that you do not want to include in the matching conditions. You can also add other conditions that are not suggested.
- If you create the rule from the Rules page, you must use the Add Condition and Add Group functions to define the property/value pairs that you want to use as conditions for the rule.
- At the bottom of the dialog box, click **More** to display the optional multiple occurrence and window-length parameters.

Conditional Expression

Select an operator and add one or more conditions to form the conditional expression. You can include a condition group to evaluate a subset of conditions. The Current Rule pane displays the constructed expression in standard syntax. The box displays a red border if the expression is syntactically invalid as currently specified. A valid expression is required to save the rule definition.

Select the operator used to determine the match for multiple conditions:

- **AND**: Match all conditions.
- **OR**: Match any one condition.
- **AND NOT**: Exclude items matching all conditions after the first.
- **OR NOT**: Include all items that do not match any conditions after the first.

Click **Add Condition** to add a condition. For each condition, specify the field name, evaluator, and value. If the evaluation returns *true* for the condition, it is a match.

Click **Add Group** to add a condition group. A new group includes a condition and its own operator used to match the conditions within the group. You can nest condition groups.

Occurrences

Specify the number of event or alarm occurrences that produce a match on the conditional expression to trigger the rule. The default value is 1. You can enter the number of occurrences or use the arrow to scroll the value up or down.

USM Anywhere uses this in conjunction with the Length option to specify the number of occurrences within a time period that will trigger the rule. For example, you can define a rule to trigger for an unauthorized access attempt when a failed SSH login occurs three times within a five-minute window.

Length

Specify the length of the window to identify a match for multiple occurrences. Enter the number and choose a time unit value of seconds, minutes, or hours. This time period identifies the amount of time that transpires from the first occurrence to the last occurrence. If the number of occurrences is not met within this period, the rule does not trigger.

7. Click **Save Rule**.

Review USM Anywhere Notifications in PagerDuty

PagerDuty creates incidents for the service from the notifications that USM Anywhere sends. You can review and respond to these incidents from your PagerDuty incidents dashboard. When you expand the details for the incident, in the Client area, click the link to open the event or alarm in USM Anywhere.

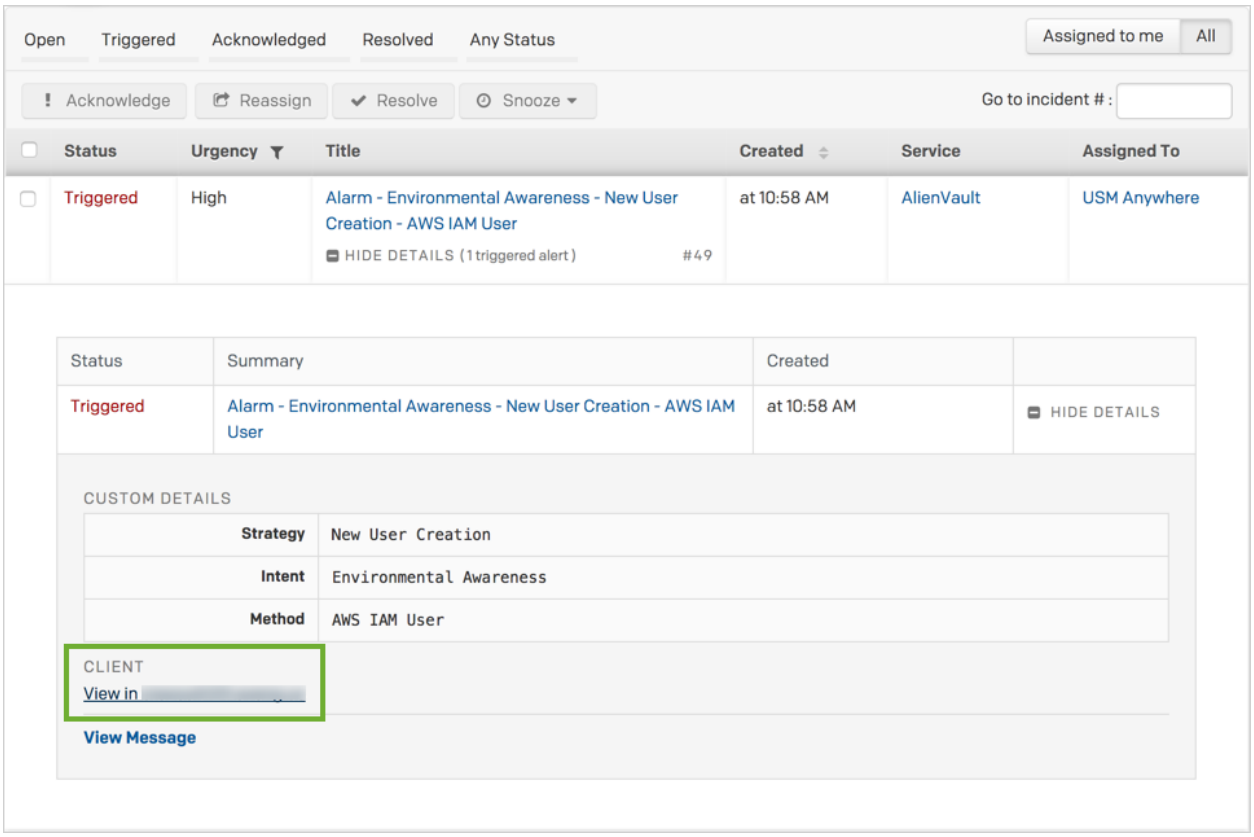

Sending Notifications Through Amazon SNS

Amazon Simple Notification Service (SNS) is a flexible messaging and mobile notifications service for coordinating the delivery of messages to subscribing endpoints and clients. You can configure SNS using the AWS Management Console, AWS Command Line Interface, or using the AWS SDK. By subscribing AWS Lambda functions to Amazon SNS topics, you can perform custom message handling.

USM Anywhere provides an integration point for Amazon SNS to connect to your SNS topic through the Amazon SNS APIs. When you have this integration configured in AWS and USM Anywhere, you can create orchestration rules to automatically send these notifications when an event or alarm matches the rule criteria.

Note: The default AWS SMS spending quota is set at $1.00 (USD) per month. SMS notifications will cease once the spending quota is exceeded; however, your spending quota can be increased upon request. For instructions, see Requesting increases to your monthly SMS spending quota for Amazon SNS.

Edition: The Notification integrations are available in the Standard and Premium editions of USM Anywhere.

See the Affordable pricing to fit every budget page for more information about the features and support provided by each of the USM Anywhere editions.

Completing the Amazon SNS integration for USM Anywhere notifications includes the following tasks:

- Set Up an Amazon SNS Topic
- Create an AWS Access Key
- Configure Amazon SNS Notifications in USM Anywhere

Set Up an Amazon SNS Topic

When using Amazon Simple Notification Service (SNS), you create a topic and control access to it by defining policies that determine which publishers and subscribers can communicate with the topic. As a publisher, USM Anywhere can then send messages (notifications) to topics for which it has the needed credentials (access key).

According to the Amazon Web Services (AWS) Documentation, when an Amazon SNS topic has an AWS Lambda function subscribed to it, it invokes the AWS Lambda function with the payload of a published message. The AWS Lambda function receives the message as an input parameter. It can manipulate the information in the message, publish the message to other Amazon SNS topics, or send the message to other AWS services.

To set up an Amazon SNS topic for USM Anywhere notifications

1. Log in to your AWS account and go to the Amazon SNS console.
2. Create a new Amazon SNS topic in the AWS SNS dashboard page:
   - Click **Topics**.
   - Click **Create topic**.

- Enter a topic name and a display name.
- Click **Create topic**.

To create an AWS Lambda function for USM Anywhere notifications

1. Open the AWS Lambda page and click **Create a function**.
2. Click **Author from scratch**.
3. Create a new AWS Lambda function:
   - Enter a name.
   - In the Runtime menu, select the current version of Python.
   - In the Execution Role section, create a new role with basic AWS Lambda permissions, use an existing role, or create a new role from AWS policy templates.
   - Click **Create Function**.

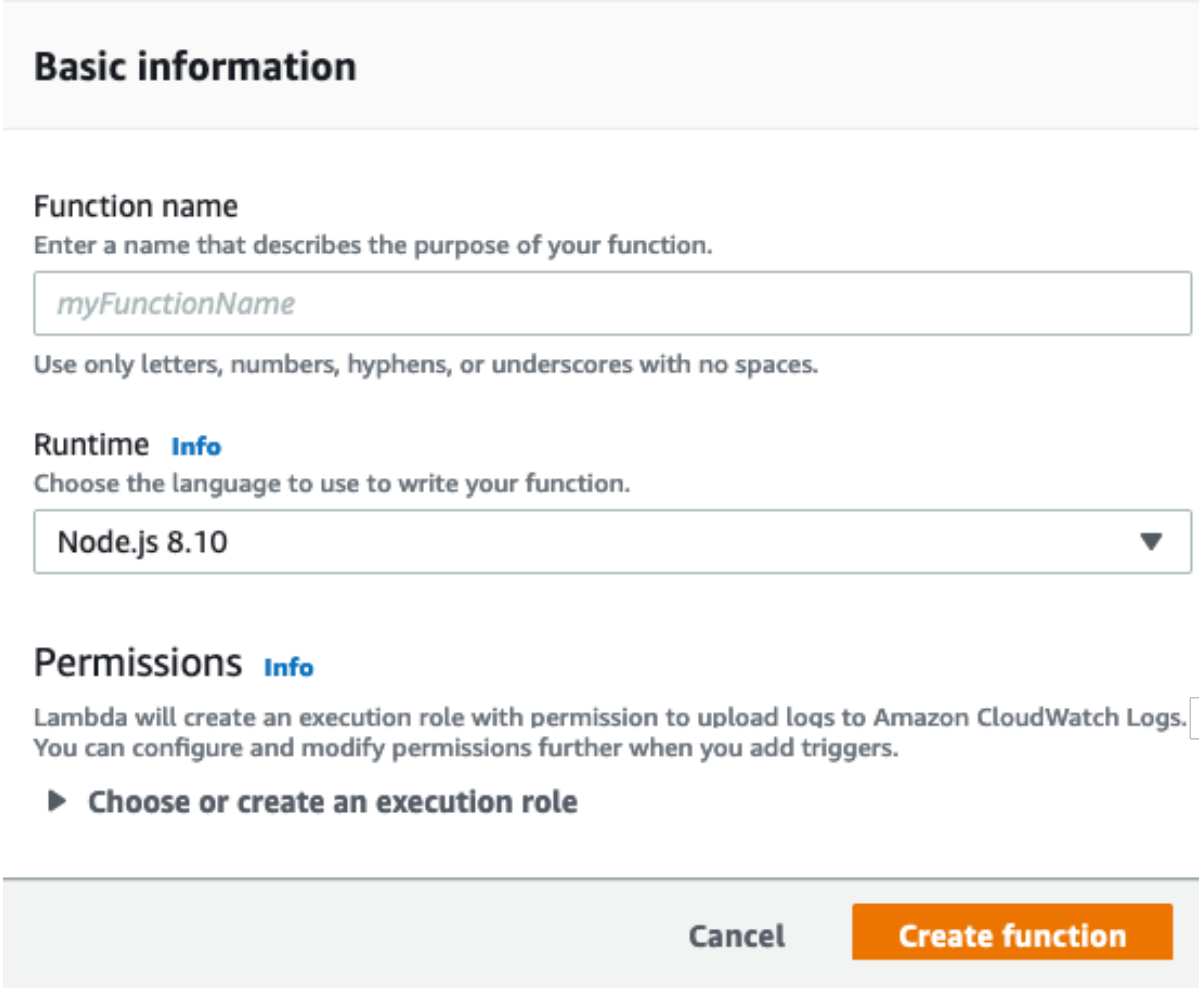

4. Assign it to the Amazon SNS topic:
   - Click **Add Trigger**.
   - Select **SNS** from the list.
   - Enter the Amazon Resource Name (ARN) of the Amazon SNS topic your created.
   - Select **Enable Trigger**.
   - Click **Add**.
5. Select the AWS Lambda function:
   - Enter the following code to send the populated fields from an alarm or event in USM Anywhere:

```
import json
def lambda_handler(event, context):
    message = json.loads(event['Records'][0]['Sns']['Message'])
    print("JSON: " + json.dumps(message))
    return message
```

> **Note:** AT&T Cybersecurity cannot provide a list of fields specific to Amazon SNS because the list varies depending on each data source.

   - Under Basic settings, set Timeout to 10 seconds.
   - Click **Save**.

Next...

If not done already, you need to create an access key in AWS for USM Anywhere to communicate with the AWS APIs.

Create an AWS Access Key

USM Anywhere requires an access key to make programmatic calls to AWS API operations. These access keys consist of an access key ID and a secret access key.

To create an AWS Access Key ID

1. Log in to your AWS Account and go to the Amazon SNS console.
2. Create a new user (see the Add User page).
3. Select **Programmatic access**.
4. Click **Next: Permissions**.

5. Click **Attach existing policies directly**.
6. Click **Create policy**.
7. Create a policy with the following code:

```
{
     "Version":"2012-10-17",
     "Statement":[{
       "Effect":"Allow",
       "Action":"sns:Publish",
       "Resource":"arn:aws:sns:us-east-1:<ACCOUNT_ID>:<USMA>"
       }
     ]
}
```

8. Replace <ACCOUNT_ID> and <USMA> with the ID of your AWS Account and the name of the SNS Topic you created. See Set Up an Amazon SNS Topic for details.
9. Attach the new policy you created.
10. Attach also the **AmazonSNSReadOnlyAccess** policy or manually add permissions to list topics ("Resource": "*").
11. Click **Next** and **Create User**.

Note: Copy the access key ID and secret access key, which you will need to configure Amazon SNS in USM Anywhere.

Configure Amazon SNS Notifications in USM Anywhere

Role Availability	Read-Only	Investigator	Analyst	Manager
	✗	✗	✗	✓

After you set up the Amazon SNS topic and create the access key for Amazon Web Services (AWS), you can configure Amazon SNS notifications in USM Anywhere.

To configure Amazon SNS Credentials for notifications

1. Go to **Settings > Notifications**.
2. In the left navigation panel, click **Amazon SNS**.
3. Select the AWS Region name.

4. Enter the Access key and Secret key. See Create an AWS Access Key for more information.

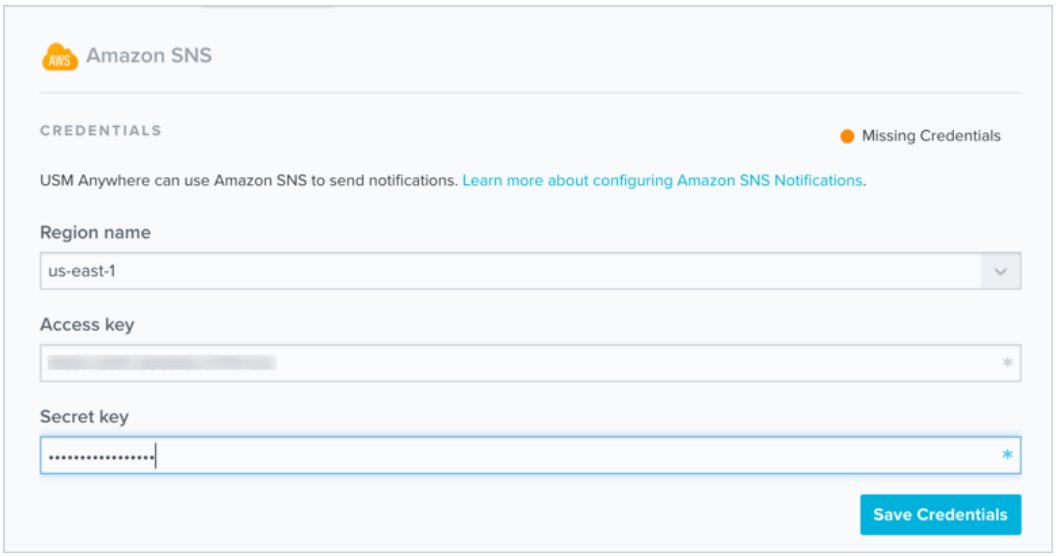

5. Click **Save Credentials**.

To create an orchestration rule for sending a notification request to Amazon SNS

1. Go to **Activity > Alarms** or **Activity > Events**.
2. Click the alarm or event to open the details.
3. Click **Create Rule** and select **Create Notification Rule**.

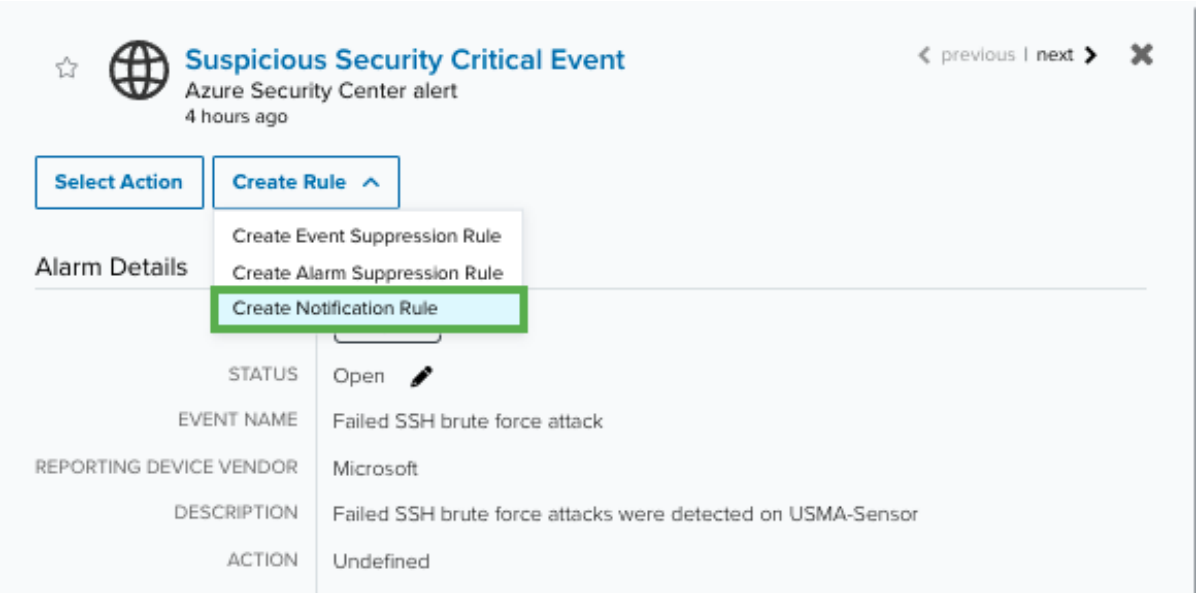

4. You have already suggested property values to create a matching condition, but if you want to add new property values, click **Add Condition**.

> **Note:** If the field is related to the name of a country, you should use the country code defined by the ISO 3166.

> **Note:** The Sources or Destinations field needs to match the universally unique identifier (UUID) of the event or alarm. You can use the Source Name or Destination Name field instead.

> **Important:** Instead of using the `equals` and `equals, case insensitive` operators for array fields, AT&T Cybersecurity recommends the use of the `in` or `contains` operators.

> **Note:** If you need to add a property value that maps with a property key, you need to know the mapping of the field. See Determining the Mapping of a Field for more information.

1. (Optional.) Click **Add Group** to group your conditions.

> **Note:** See Operators in the Orchestration Rules for more information.

> **Note:** The current rule box shows you the syntax of your rule, and the rule verification box reviews that syntax before saving the rule.

2. Click **Next**.

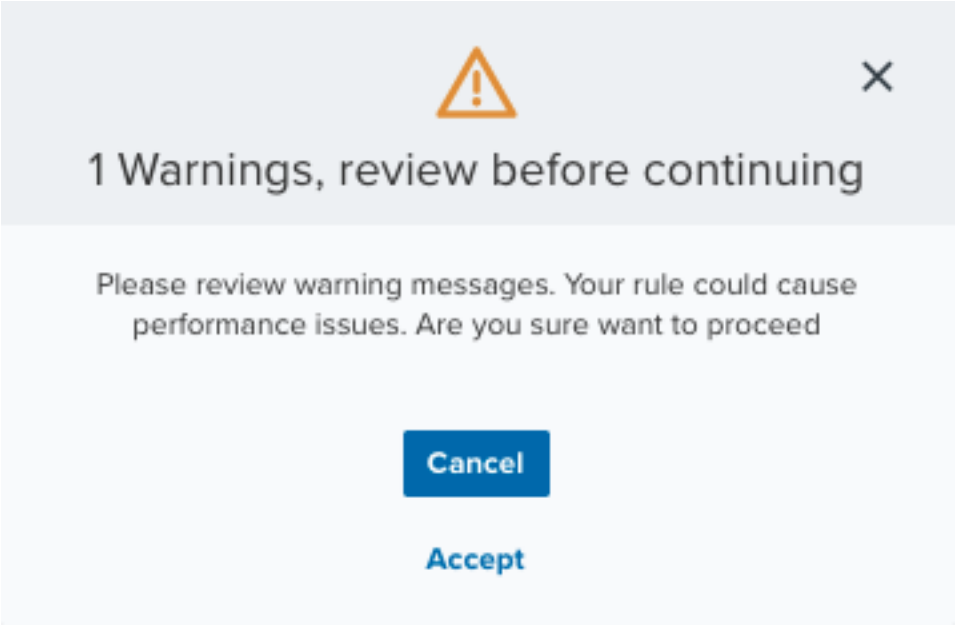

Important: A dialog box opens if there are warning messages. Click **Cancel** to review the warning messages, or click **Accept** to continue creating the rule.

3. Enter a name for the rule.
4. (Optional.) Enter a description for identifying this rule.
5. For Notification Method, select the **Slack** option.
6. Enter the Slack Alert Username.

   The username must be a valid team member for the Slack channel.

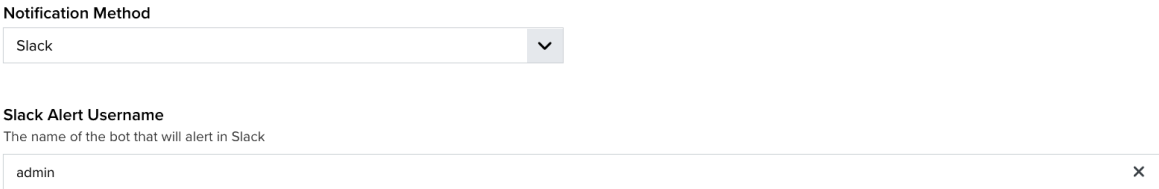

7. Modify these two options:

- **Occurrences:** Specify the number of event occurrences that produce a match on the conditional expression to trigger the rule. You can enter the number of occurrences or use the arrow to scroll the value up or down. You need to enter a number between 1 and 100.
- **Length:** Specify the length of the timespan used to identify a match for multiple occurrences. Enter the number and choose a value of seconds, minutes, or hours.

  This duration identifies the amount of time that transpires from the beginning to the end of the occurrence. If the number of occurrences is not met within this period, the rule is not a match.

  In this example, the rule applies when the configured conditions happen five times every three hours.

These two options function together to specify the number of occurrences within a time period that will produce a match for the rule. For example, you can define a rule to trigger an alarm for an unauthorized access attempt when a failed SSH login occurs three times within a five-minute window.

8. Click **Save**.

   The created rule displays in the list of rules. You can see it from **Settings > Rules > Orchestration Rules**. See Orchestration Rules for more information.

   Important: It takes a few minutes for an orchestration rule to become active.

5. For **Notification Method**, select the **Amazon SNS** option.
6. Enter the **SNS Topic Name** you created in the AWS console. See Set Up an Amazon SNS Topic for more information.

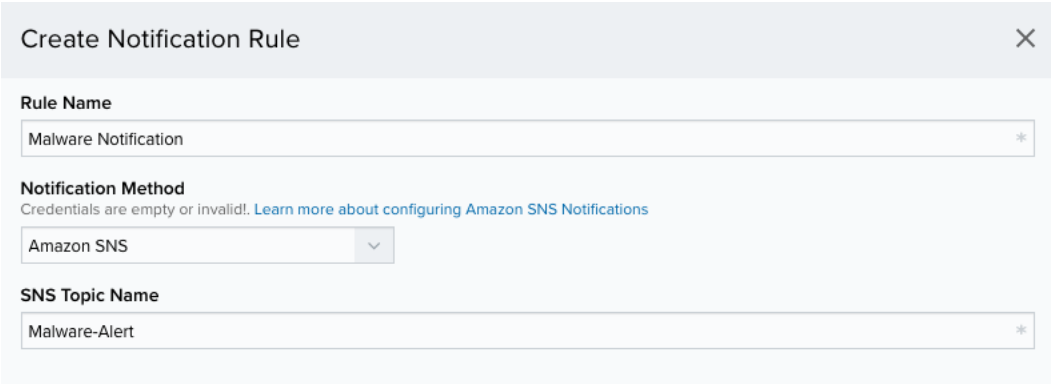

7. At the bottom of the dialog box, set the rule condition parameters to specify the criteria for a matching alarm or event to trigger the rule.

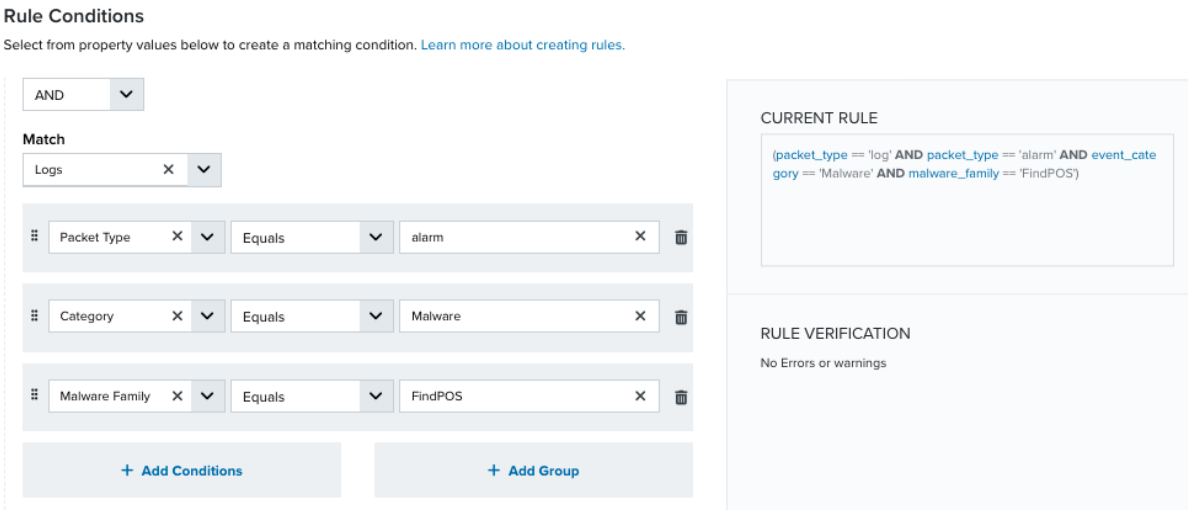

- This section provides suggested property/value pairs from the selected alarm or event that you can use as conditions for the rule. Click the icon to delete the items that you do not want to include in the matching conditions. You can also add other conditions that are not suggested.
- If you create the rule from the Rules page, you must use the Add Condition and Add Group functions to define the property/value pairs that you want to use as conditions for the rule.
- At the bottom of the dialog box, click **More** to display the optional multiple occurrence and window-length parameters.

Conditional Expression

Select an operator and add one or more conditions to form the conditional expression. You can include a condition group to evaluate a subset of conditions. The Current Rule pane displays the constructed expression in standard syntax. The box displays a red border if the expression is syntactically invalid as currently specified. A valid expression is required to save the rule definition.

Select the operator used to determine the match for multiple conditions:

- **AND**: Match all conditions.
- **OR**: Match any one condition.
- **AND NOT**: Exclude items matching all conditions after the first.
- **OR NOT**: Include all items that do not match any conditions after the first.

Click **Add Condition** to add a condition. For each condition, specify the field name, evaluator, and value. If the evaluation returns *true* for the condition, it is a match.

Click **Add Group** to add a condition group. A new group includes a condition and its own operator used to match the conditions within the group. You can nest condition groups.

Occurrences

Specify the number of event or alarm occurrences that produce a match on the conditional expression to trigger the rule. The default value is 1. You can enter the number of occurrences or use the arrow to scroll the value up or down.

USM Anywhere uses this in conjunction with the Length option to specify the number of occurrences within a time period that will trigger the rule. For example, you can define a rule to trigger for an unauthorized access attempt when a failed SSH login occurs three times within a five-minute window.

Length

Specify the length of the window to identify a match for multiple occurrences. Enter the number and choose a time unit value of seconds, minutes, or hours. This time period identifies the amount of time that transpires from the first occurrence to the last occurrence. If the number of occurrences is not met within this period, the rule does not trigger.

8. Click **Save Rule**.

When a matching alarm or event is generated in USM Anywhere, you can go to your AWS console and select the Lambda function you created to verify that the function is being called. You can also open the Amazon CloudWatch logs to see the message in JavaScript Object Notation (JSON) format.

Troubleshooting and Remote Sensor Support

Use Remote Support to allow the AT&T Cybersecurity Technical Support to access and diagnose the components identified in a support ticket. USM Anywhere offers remote technical support through the USM Anywhere Sensor console. All data exchanged with AT&T Cybersecurity Support is encrypted for security. The information exchanged is only available to AT&T Cybersecurity Support and Engineering teams.

Typically, you open a ticket with AT&T Cybersecurity Support first and only establish a remote support connection upon their request. You can establish multiple sessions using the same ticket number for different sensors. But a support engineer could ask you to open a new ticket if it is an unrelated issue. During the remote support session you can communicate with the AT&T Cybersecurity Technical Support by phone or email at any time.

This section includes the following topics:

Checking Connectivity to the Remote Server 451

Creating a Remote Support Session 453

Sensor System Menu 456

Collecting Debug Information 466

View Network Testing Information 468

Retrieve Unique Identifier Information 476

Checking Connectivity to the Remote Server

Role Availability ✗ Read-Only ✗ Investigator ✓ Analyst ✓ Manager

Before you can remotely connect with AT&T Cybersecurity Technical Support, you need to verify your connection to the Remote Support server from the sensor. The USM Anywhere Sensor uses port 22 and 443 for SSH communications with the USM Anywhere Remote Support server. If there is an issue with your connectivity, make sure that port 22 and 443 are open to `prod-usm-saas-tractorbeam.alienvault.cloud` or `prod-gov-usm-saas-tractorbeam.gov.alienvault.us` (for AT&T TDR for Gov). If the ports *are open* and you still *have no connectivity*, check for any other physical problem on your side. If none are found, contact AT&T Cybersecurity Technical Support to find out if their server is temporarily down.

To check the network connectivity for the sensor

1. While logged in to USM Anywhere, check your networking status by going to **Settings > System**.
2. If you have more than one deployed sensor, select the one that you want to verify.

   If the page reports that this endpoint is unreachable, you may have a problem.

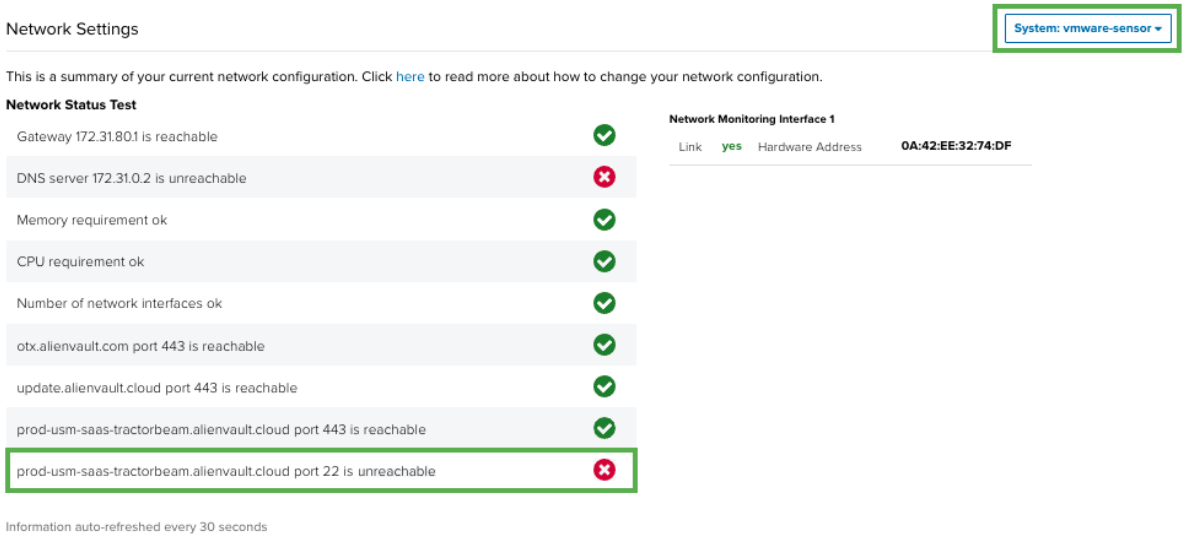

To verify remote support connectivity directly on the USM Anywhere Sensor

1. Open your virtualization management console and connect to the USM Anywhere Sensor virtual machine (VM).

   Alternatively, you can open an SSH session to the sensor VM. When using an SSH session, the default username is *sysadmin*.

   > **Important:** If you are accessing a Microsoft Azure Sensor through SSH and you specified a username other than the default (*sysadmin*) for your SSH access, you must use the following commands at the command line to "sudo up" and access the sensor console:
   >
   > ```
   > # sudo su - sysadmin
   > ```

2. From the USM Anywhere Sensor console System Menu, select **Maintenance** and press **Enter**.

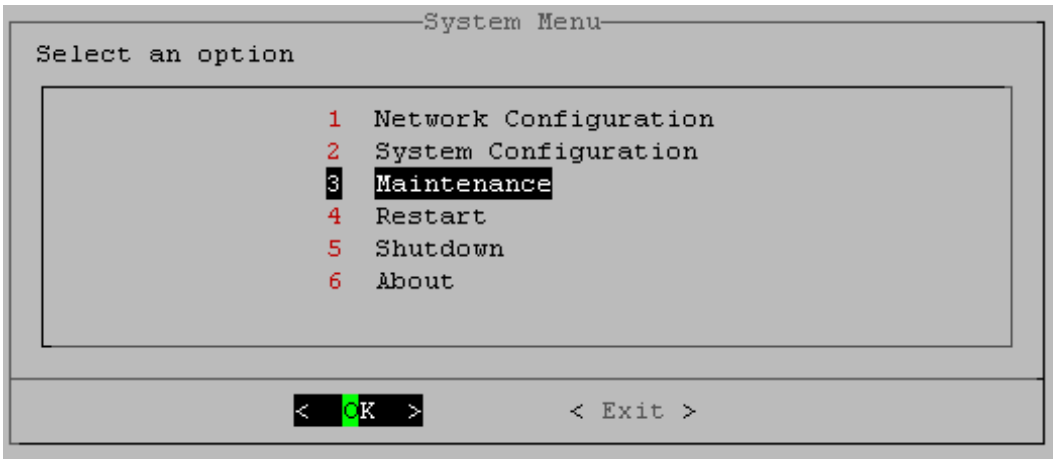

3. From the Maintenance menu, select **Remote Support** and press **Enter**.

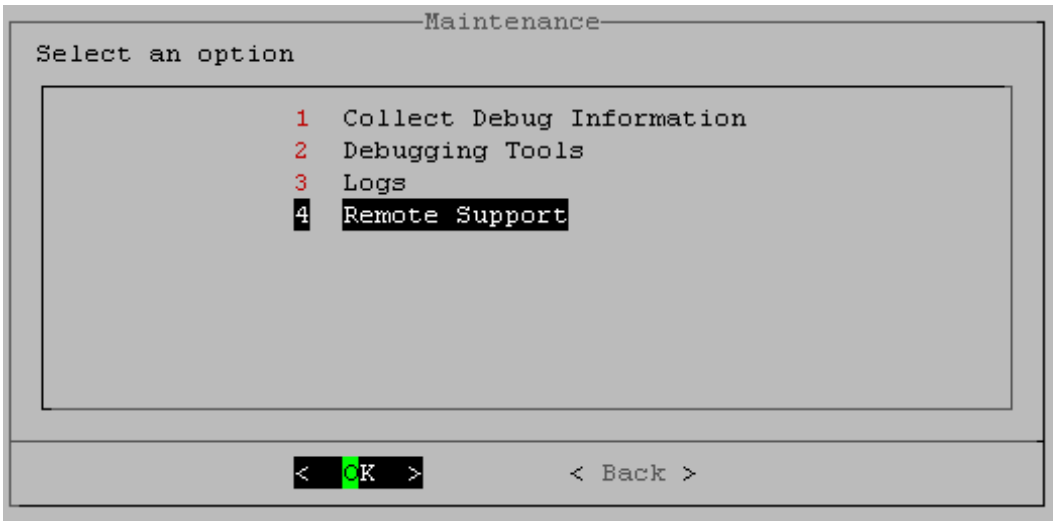

4. From the Remote Support menu, select **Show Remote Support Status**, press the arrow-down (↓) key, and then **Enter**.

   The system displays an alert message that the check is in progress.

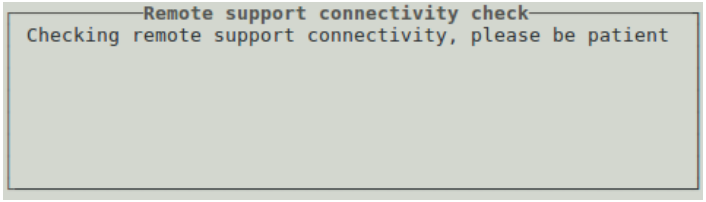

   When the check is complete with a connection, you see a success alert.

   If there's no connection, the system displays an alert message that the remote server is unreachable.

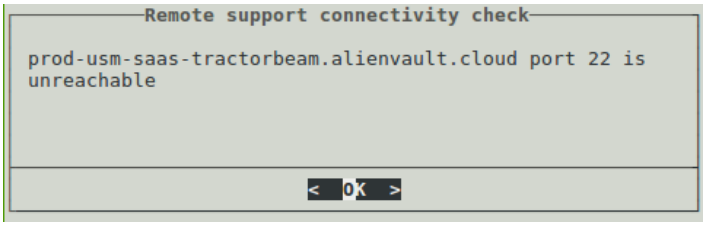

   Note: If the system does make a connection, you see a success prompt.

5. Press **Enter**.

Creating a Remote Support Session

After you confirm that you have a connection to AT&T Cybersecurity Technical Support, you are ready to start a session. When AT&T Cybersecurity Technical Support Engineers complete their work on an issue, they communicate the results to you by email and update your ticket.

Note: USM Anywhere provides an audit trail of the AT&T Cybersecurity Technical Support Engineers accessing your instance by creating a temporary user with the username of `<user>@alienvault.com`. These users are disabled after the session ends, and you can view them under Settings > Users.

To enable remote sensor support

1. Open your virtualization management console and connect to the USM Anywhere Sensor virtual machine (VM).

   Alternatively, you can open an SSH session to the sensor VM. When using an SSH session, the default username is *sysadmin*.

   Important: If you are accessing a Microsoft Azure Sensor through SSH and you specified a username other than the default (*sysadmin*) for your SSH access, you must use the following commands at the command line to "sudo up" and access the sensor console:

   ```
   # sudo su - sysadmin
   ```

2. From the USM Anywhere Sensor console System Menu, select **Maintenance** and press **Enter**.

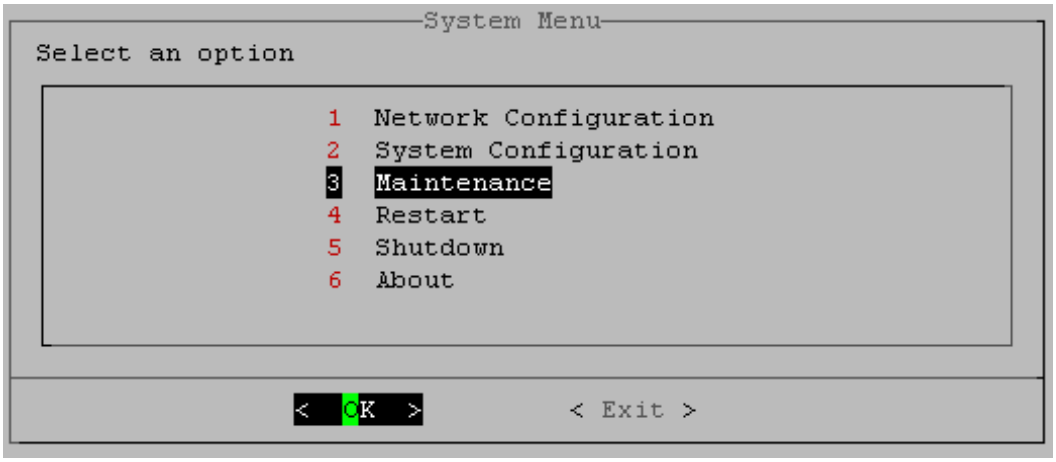

3. From the Maintenance menu, select **Remote Support** and press **Enter**.

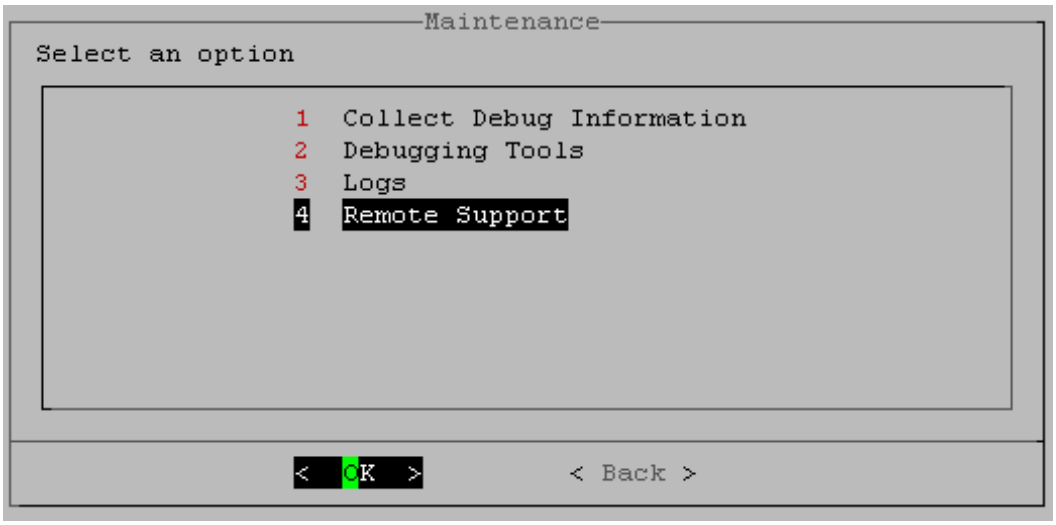

4. From the **Remote Support** menu, select **Enable Remote Support** and press Enter (**< OK >**).

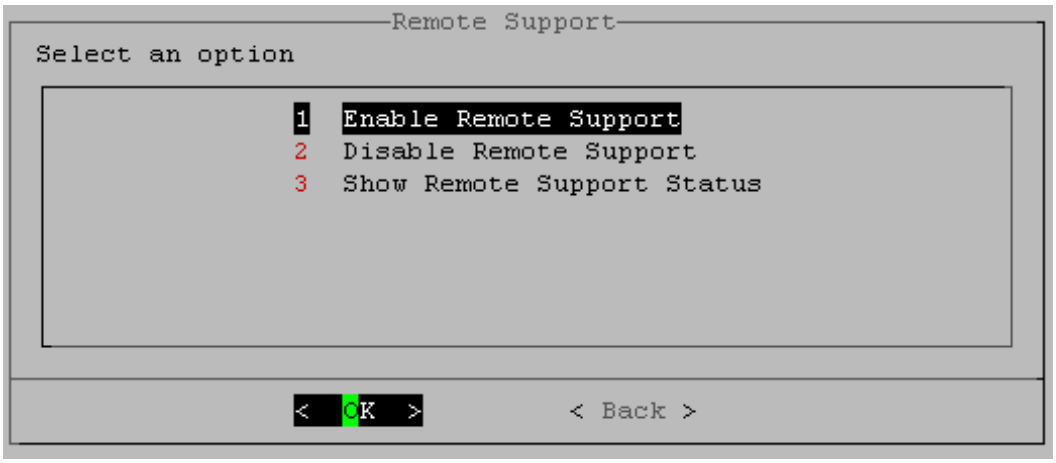

5. On the **Enable Remote Support** screen, enter the eight-digit ticket number and press Enter (**< OK >**).

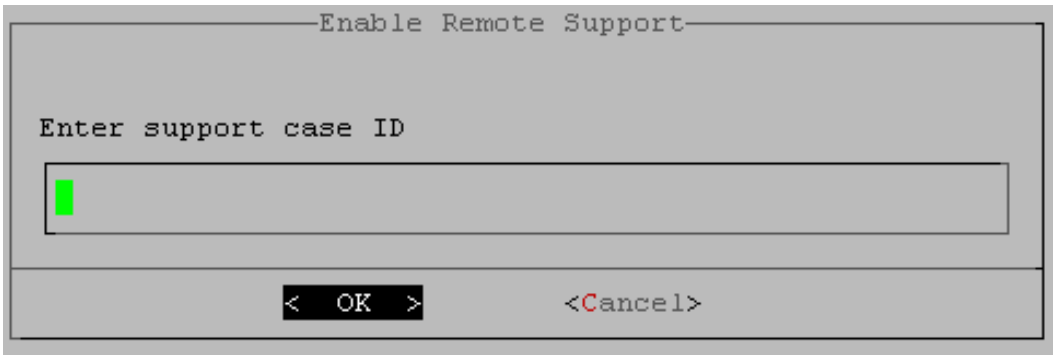

> **Important:** Be careful not to enter any spaces before or after the number or the operation will fail.

A progress bar appears and your request begins processing, which may take several seconds.

When the connection is established with the Support server, the system displays a connection message.

```
Connected to AlienVault Support. Press Enter to continue.
```

6. Press Enter.

   The USM Anywhere Sensor console returns you to the to the Remote Support screen.

7. To disconnect after your session is done, select **Disable Remote Support** and press **Enter**.

   The **Manage Connectivity** information screen appears and prompts you to confirm.

```
Are you sure you want to disconnect from AlienVault Remote Support?
```

8. Select **Yes**.

   The screen goes black and, after several seconds, you receive a notification that the secure connection is now disconnected. You can then back out of the previous menus and close the sensor console.

Sensor System Menu

Network Configuration

You can view your network configuration through the USM Anywhere Sensor console, allowing you to diagnose your network issues and set a static management IP address.

To set a static management IP address

1. Open your virtualization management console and connect to the USM Anywhere Sensor virtual machine (VM).

   Alternatively, you can open an SSH session to the sensor VM. When using an SSH session, the default username is *sysadmin*.

Important: If you are accessing a Microsoft Azure Sensor through SSH and you specified a username other than the default (*sysadmin*) for your SSH access, you must use the following commands at the command line to "sudo up" and access the sensor console:

```
# sudo su - sysadmin
```

2. From the USM Anywhere Sensor console System menu, select **Network Configuration** and press **Enter**.

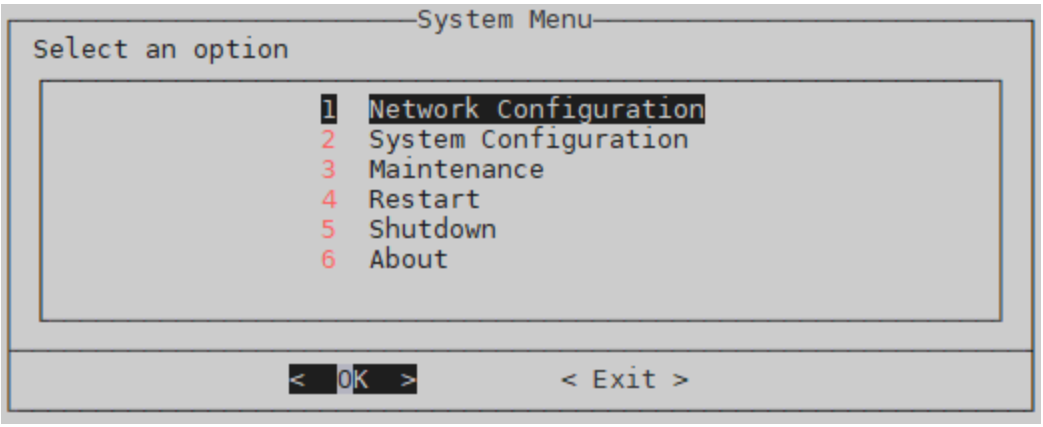

3. From the Network Configuration menu, select **View Network Configuration** and press **Enter**.

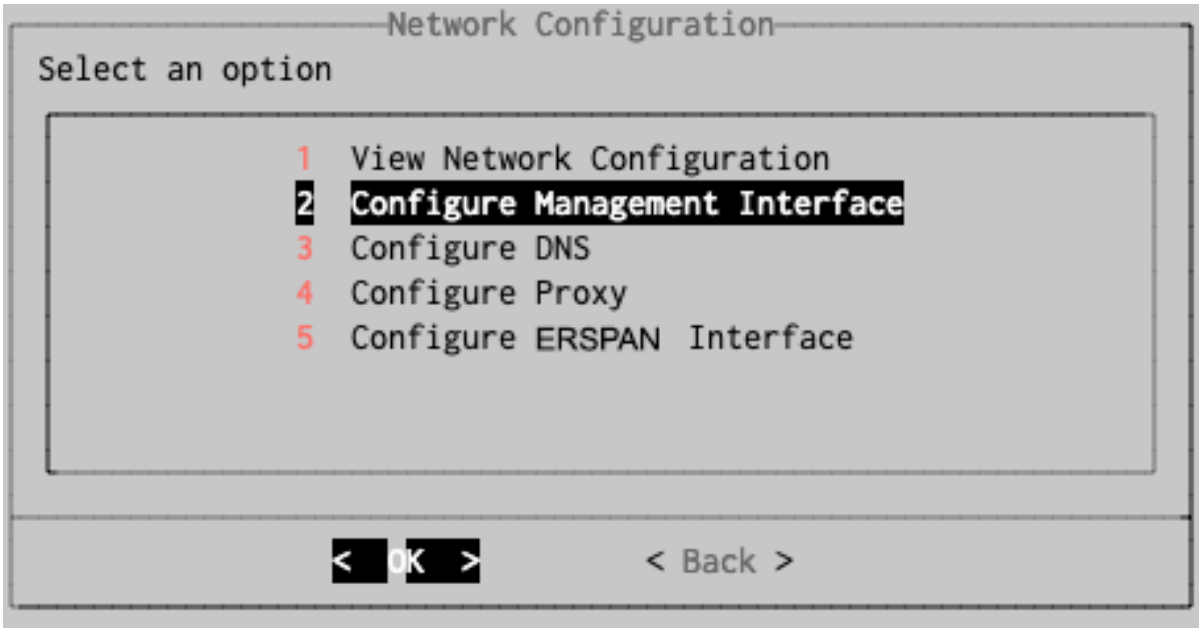

4. The system displays the View Network Configuration screen, press **Enter**.

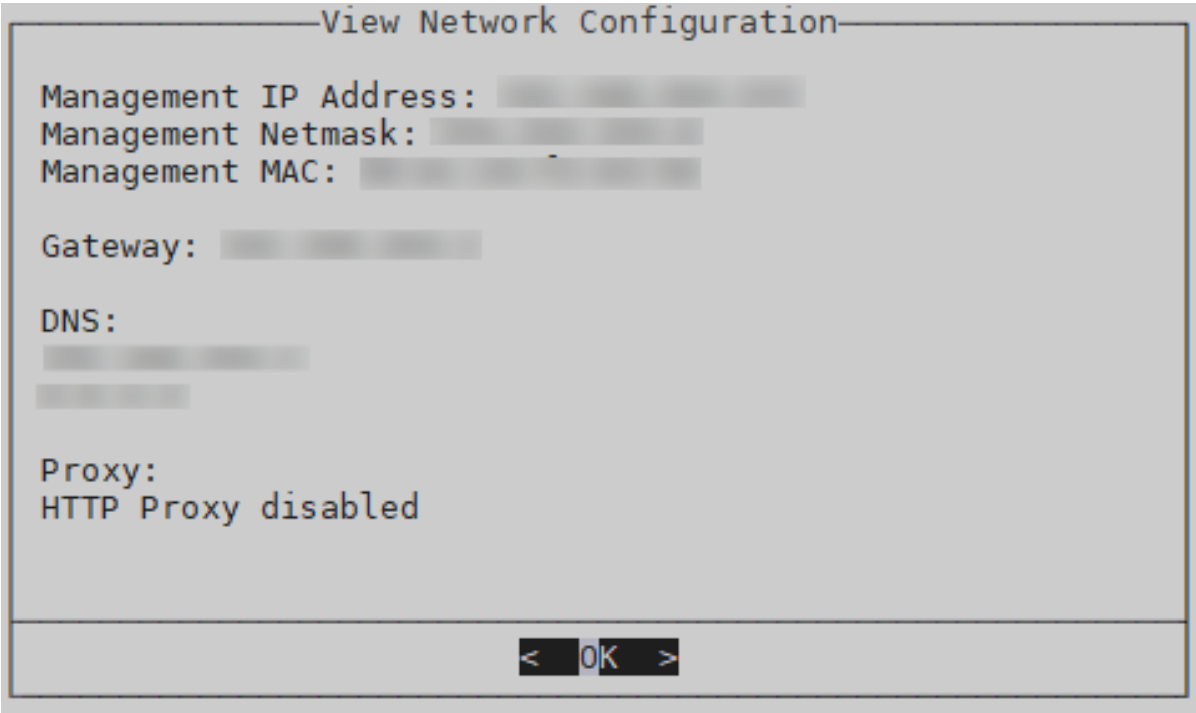

5. Select **Set a Static Management IP Address** and press **Enter**.

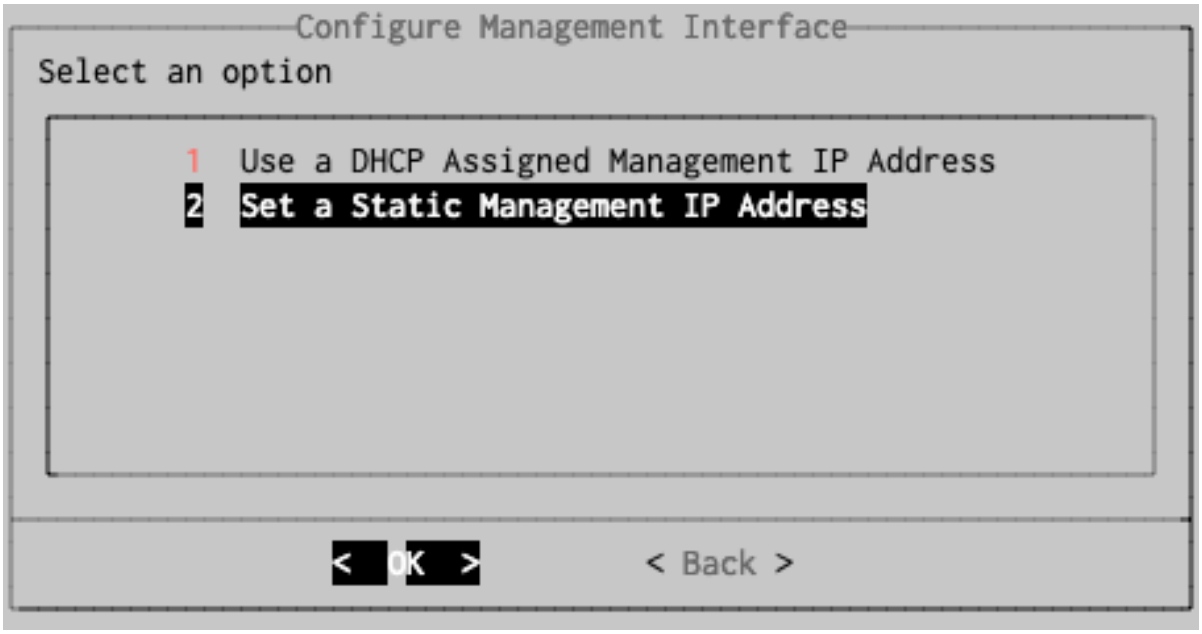

6. Enter the IP address and press **Enter**.

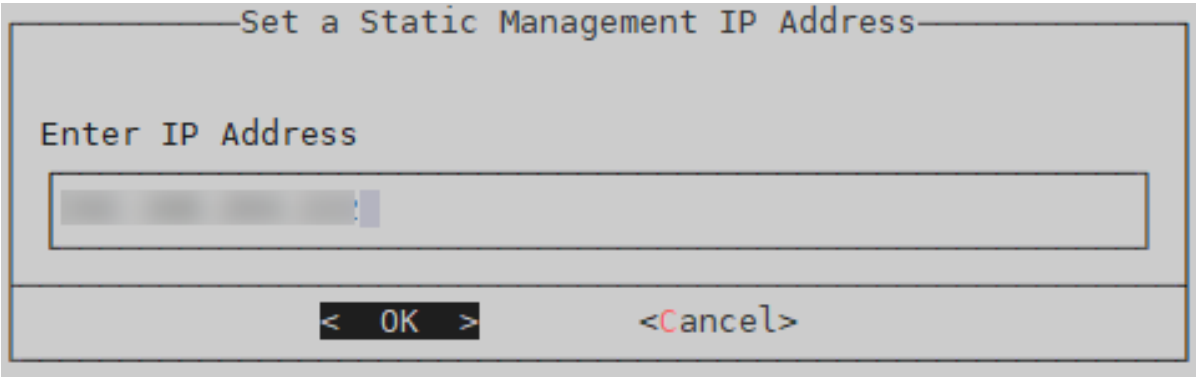

7. Enter the netmask and press **Enter**.

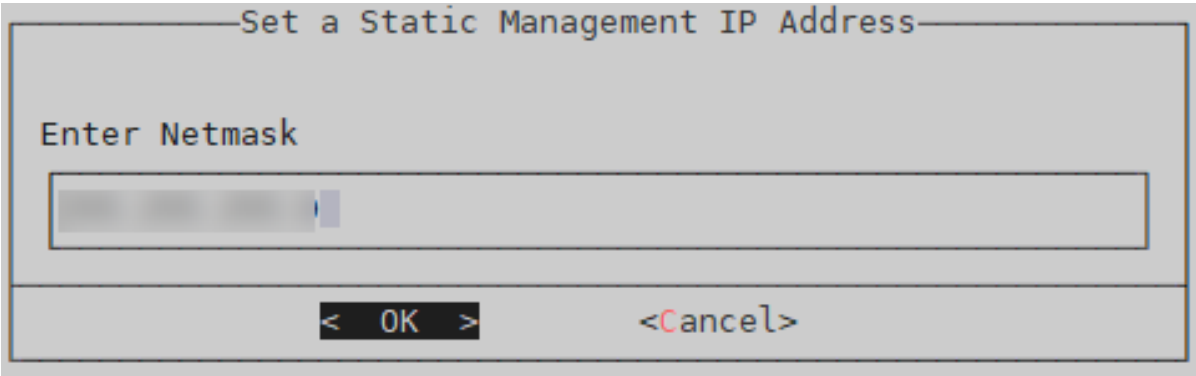

8. When a summary of your changes displays, review them for accuracy. Press **Enter** if they are accurate, or select **No** to edit your entries.

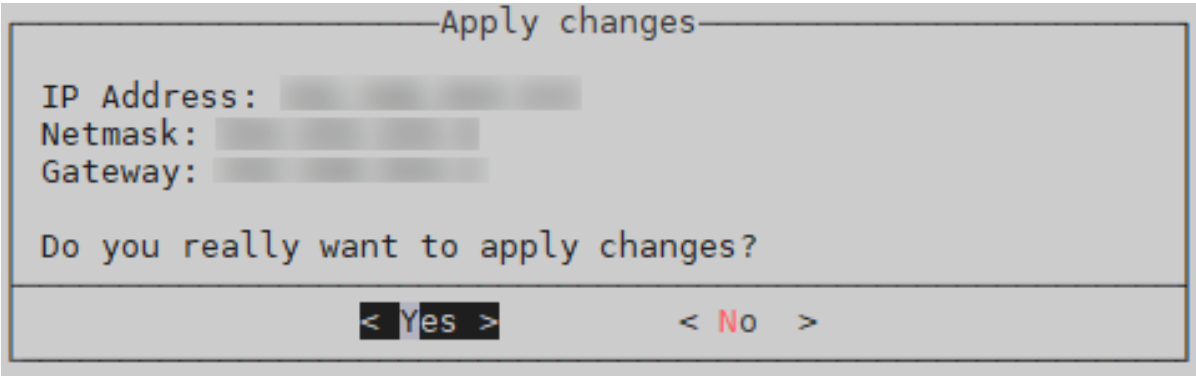

9. When your changes have been applied the system will display a message saying 'Changes applied'. Press **Enter**.

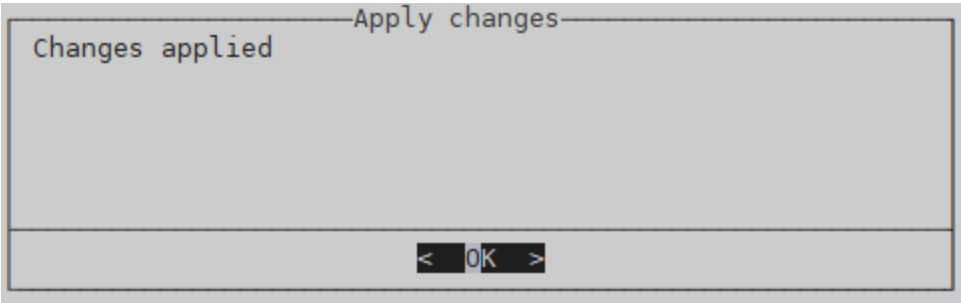

10. After the system displays a notice to configure your USM Anywhere Sensor, press **Enter**.

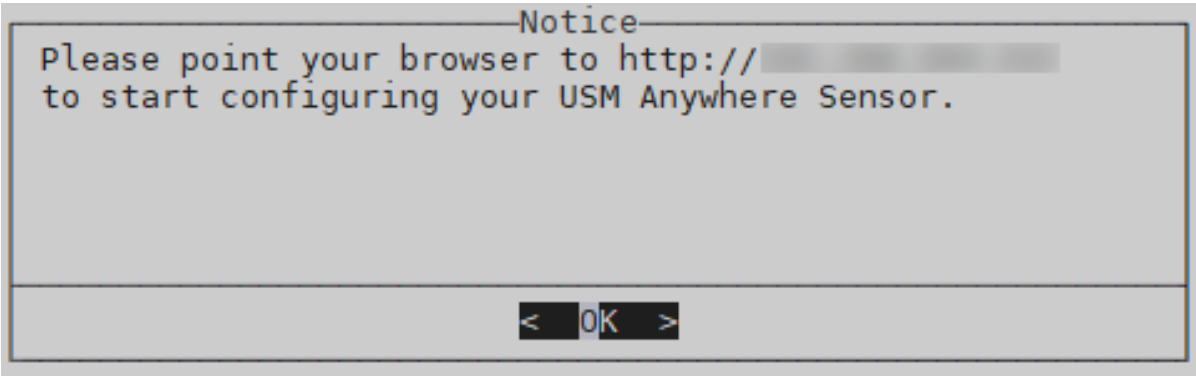

To see other options of the Network Configuration menu

1. From the USM Anywhere Sensor console System menu, select **Network Configuration** and press **Enter**.

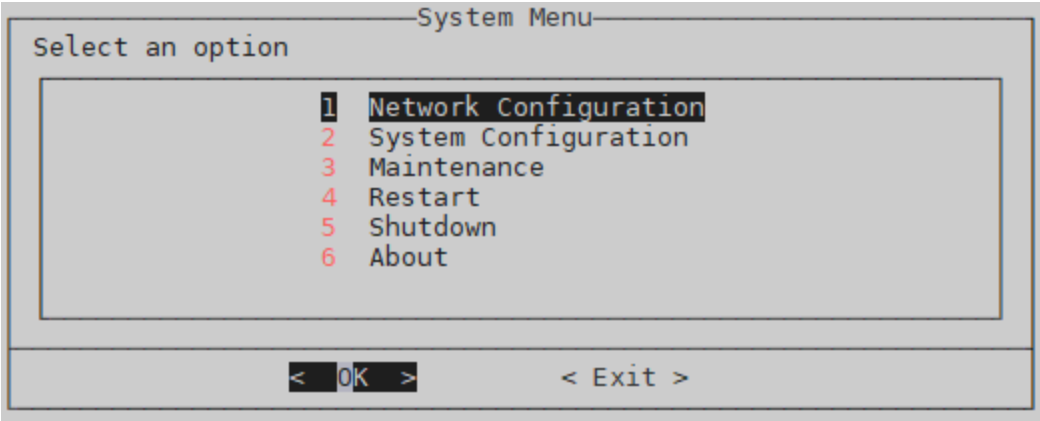

2. From the Network Configuration menu, select **Configure Management Interface** and press **Enter**.

The management interface is the primary IP address used to connect to a sensor node. The user sets this IP when the sensor node virtual image is created. This value doesn't change unless, after created, you set it with Dynamic Host Configuration Protocol (DHCP). AT&T Cybersecurity recommends that all sensors are configured with a static IP address and that you not change the value.

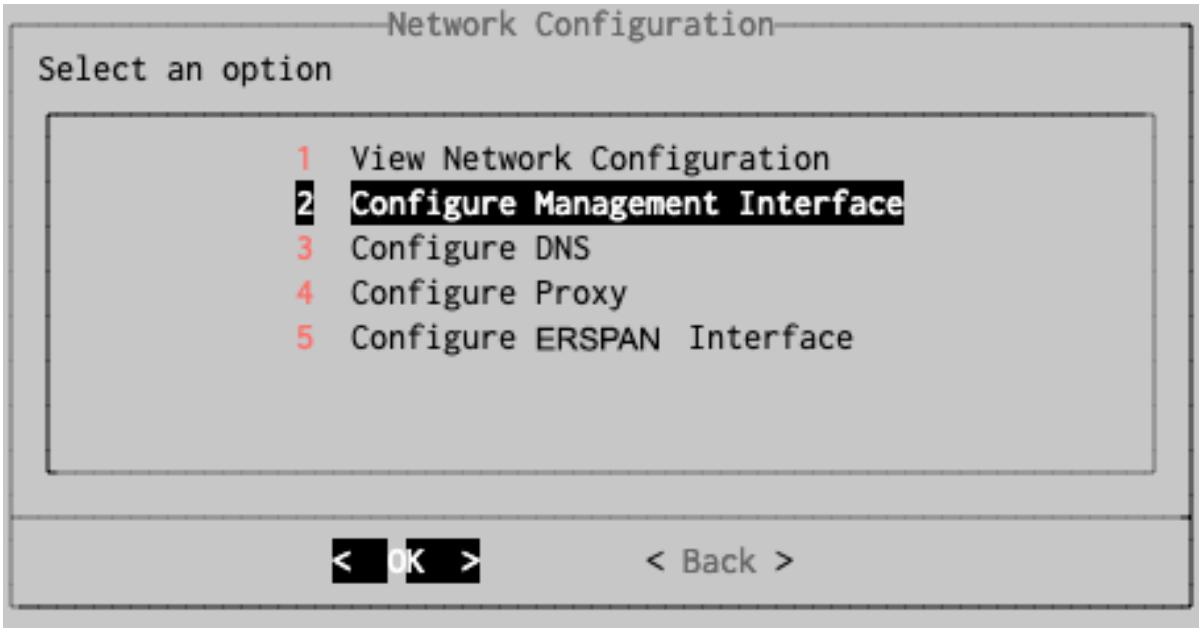

3. Go back to the Network Configuration menu, select **Configure DNS**, and press **Enter**.

   This option displays the current Domain Name System (DNS) server settings for the sensor. Use this option to modify them. You can specify a primary and secondary DNS server.

4. Go back to the Network Configuration menu, select **Configure Proxy**, and press **Enter**. This is for an HTTP port 80 proxy. AT&T Cybersecurity recommends not to create one.
5. Go back to the Network Configuration menu, select **Configure ERSPAN Interface**, and press **Enter**.

   See Configure USM Anywhere to Receive ERSPAN Traffic for more information.

System Configuration

This option enables you to configure the sensor keyboard layout and the sensor working time zone.

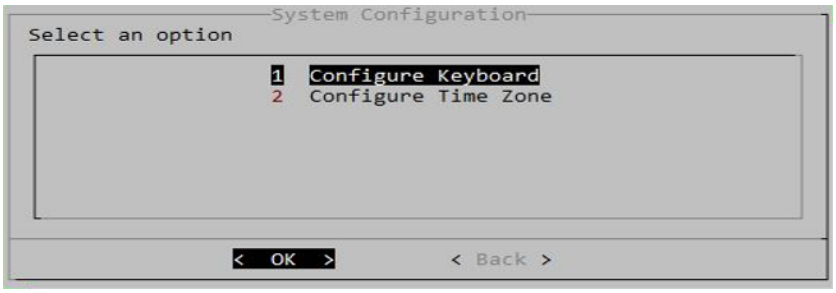

Maintenance

The Maintenance menu includes several useful options that enable you to perform debug and research of the sensor node.

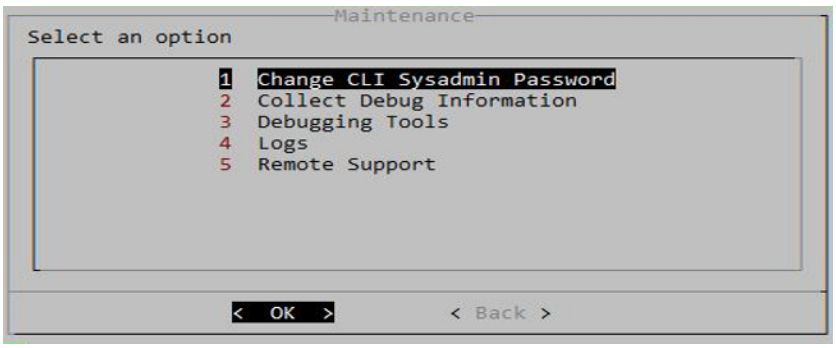

Change CLI Sysadmin Password

You always access to the sensor using the sensor username of “sysadmin”. There is no access to the sensor node command line and no user access to a "root"-level username. The initial password is set when the sensor node is created and initially configured. Be extremely careful to save this password. There is no “Forgot Password” function available for sensor access. If the password is lost or forgotten, a sensor redeploy action is the only possible recovery.

Important: This option is not displayed for Amazon Web Services (AWS) Sensors. These sensors must use AWS console actions to these modifications.

Collect Debug Information

This option enables you to gather and download all of the debug logs created and that are available on a sensor node. When you select this option, you will see several new views. The first is a verification of the debug data request.

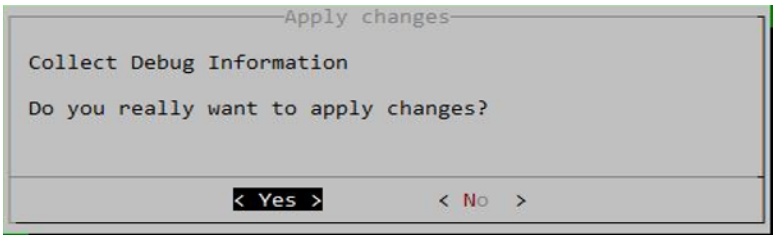

If you enter “No”, the panel escapes back to the previous panel view. If you enter “Yes”, the sensor will operate to collect the debug files into a password protected `zip` format file. When it finishes, the new panel displays to collect the file. You need the password to extract the data.

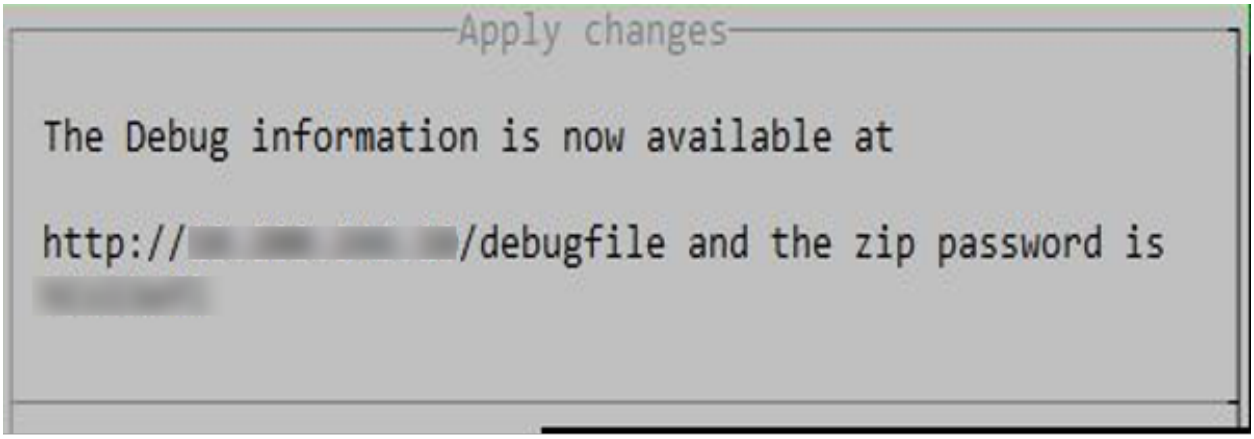

Debugging Tools

The user interface (UI) comes packaged with useful system level and networking level debugging tools to assist with diagnosing sensor node problems. Use these tools for internal sensor troubleshooting. This information can be different than what is seen in the management consoles of the different sensor types. These are the debugging tools:

- **System Debugging Tools**: The “System” selection uses standard system-level diagnostic tools, like htop. After a confirmation screen is displayed, you will see the output of the tool. As noted in the screens, use “q” or “F10” to revert back.
- **Network Debugging Tools**: The Networking Debugging Tools UI view displays available tools to monitor network traffic on the sensor.

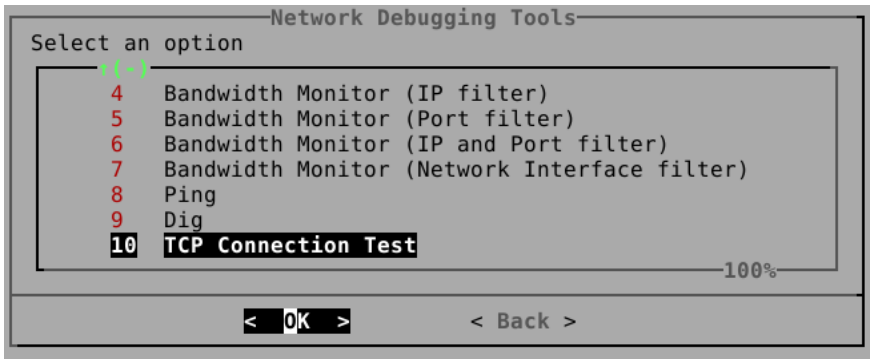

These are the Network Debugging Tools:

- **Network Test**: This option performs several probing commands to verify sensor connectivity to needed external servers and ports. If there are external connectivity issues, this option will display them.
- **Network Monitor**: This option displays network traffic between server endpoints associated with the sensor.
- **Bandwidth Monitor**: This option displays all endpoint network activity with IP address and port information. For each entry, it shows BPS data for transmits, receipts, and total activity.
- **Bandwidth Monitor With Filtering**: These options enable you to target filtering of the data based on the filter type noted in the option. When selected, a new panel displays to set the values for the filter.
- **Ping**: This option enables you to test the reachability of a given endpoint. This does not guarantee that the required TCP ports are open to allow connectivity (see Dig below).
- **Dig**: This option enables you to find the IP address of a given endpoint using its FQDN (such as *prod-usm-saas-tractorbeam.alienvault.cloud*). Then, you can use that IP address and Ping or the TCP connection test with a port number to test the reachability and connectivity to the endpoint, respectively.
- **TCP Connection Test**: This option enables you to test the TCP connection for successful data transmission to a given endpoint using an IP address and a port number.

> A TCP connection test can only be performed using IP addresses.
>
> When deep packet inspection (DPI) is implemented, this TCP Connectivity test may still succeed while the secure connection to the endpoints fails.

See View Network Testing Information for more information.

Logs

This option enables you to access a read-only view of three debug logs being captured on the sensor node. A confirmation view pane is displayed ahead of the data. The unimatrix, unimatrix-stats, and hermes stats logs are part of the Debug Logs `zip` file. The `zip` file may contain additional log files of the same data in compressed "gz" format. When the active file grows to its maximum size, the file is compressed, given a numerically tagged new name, and a fresh log file created.

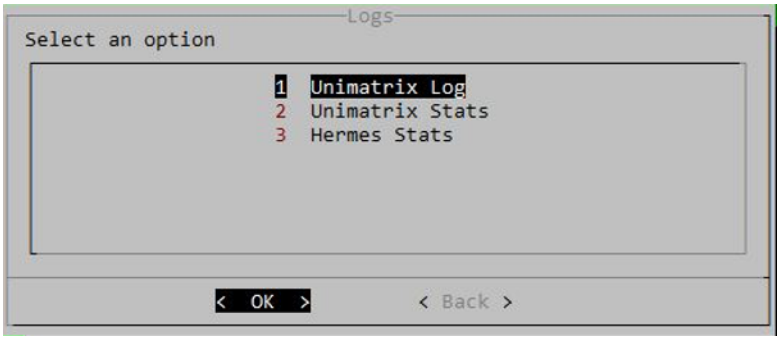

Remote Support

As needed, AT&T Cybersecurity Technical Support may require access to the sensor node for research and debug purposes. It does so through a special outbound portal that must be initiated by the user through this option. To perform this action, ports 22 and 443 must be open.

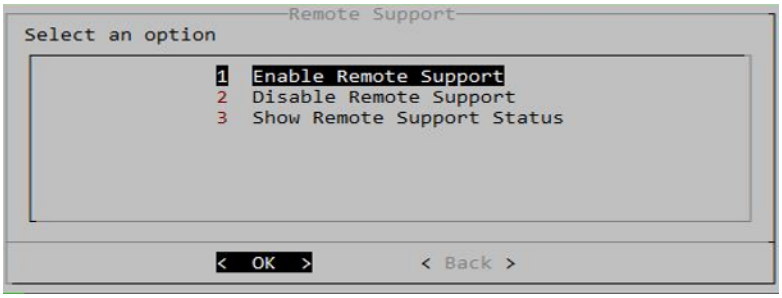

Restart

This option does a reboot and restart action on the sensor node.

Shutdown

This option performs a graceful shutdown of the sensor node image.

About

This option displays the current running sensor version. It should match the version running on the Control Node. If not, this can cause side effects on how the sensor node interacts with the control node.

Exiting the Sensor UI View

When all activity is completed, exit from the UI by selecting the **Exit** option from the main screen. This logs off users from the sensor node.

Collecting Debug Information

When you open a ticket with AT&T Cybersecurity Technical Support, you can include collected debug information to assist the support engineer with diagnosing your issue. The USM Anywhere Sensor console provides a function that you can use to collect this information. When enabled, the sensor rotates the debug logs when they exceed 100 megabytes (MB), and keeps up to 7 rotated files afterward.

To collect debug information for the sensor

1. Open your virtualization management console and connect to the USM Anywhere Sensor virtual machine (VM).

   Alternatively, you can open an SSH session to the sensor VM. When using an SSH session, the default username is *sysadmin*.

Important: If you are accessing a Microsoft Azure Sensor through SSH and you specified a username other than the default (*sysadmin*) for your SSH access, you must use the following commands at the command line to "sudo up" and access the sensor console:

```
# sudo su - sysadmin
```

2. From the USM Anywhere Sensor console System Menu, select **Maintenance** and press **Enter**.

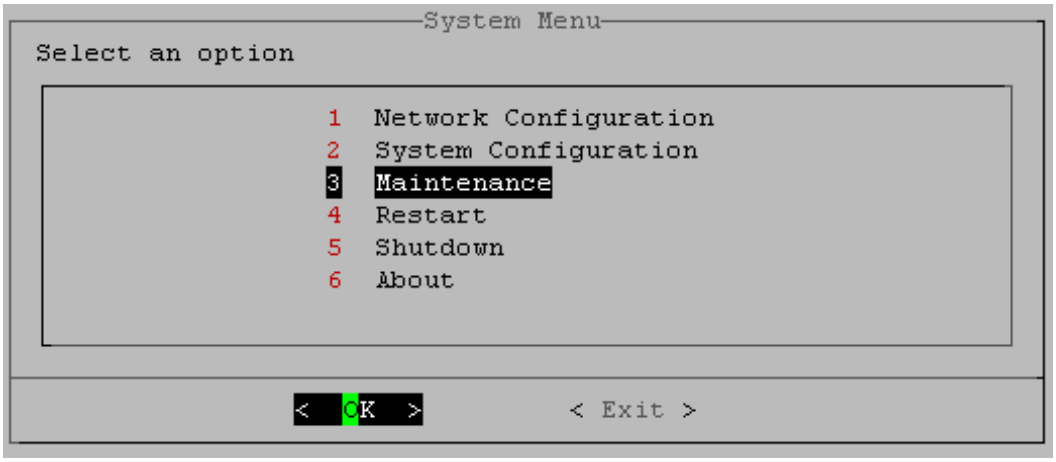

3. From the Maintenance menu, select **Collect Debug Information** and press **Enter**.

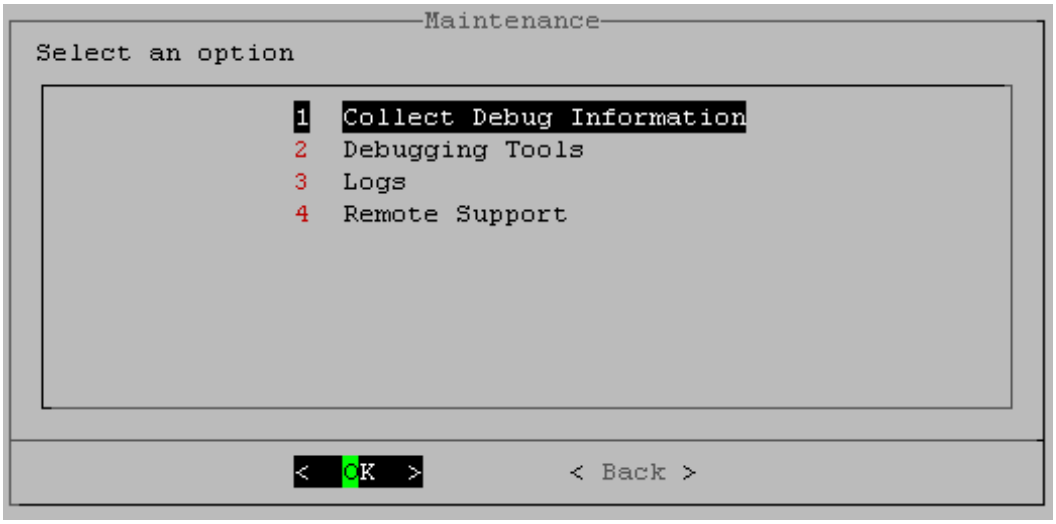

4. In the confirmation screen, select **Yes** and press **Enter**.

When the collection process is complete, you see an alert message. This provides the URL for the file and the password.

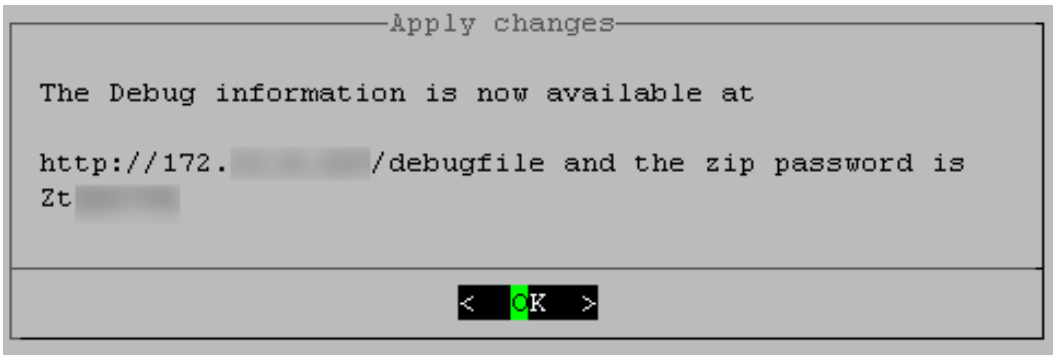

5. Press **Enter**.
6. Download the debug file and attach it to your support case.

   Make sure to update the support case information to include the file password.

View Network Testing Information

When you open a ticket with AT&T Cybersecurity Technical Support, you may be required to test the sensor's network connectivity with the system debugging tool to assist the support engineer with diagnosing your issue. The USM Anywhere Sensor console provides a function that you can use to collect this information.

To view sensor network test information for the sensor

1. Open your virtualization management console and connect to the USM Anywhere Sensor virtual machine (VM).

   Alternatively, you can open an SSH session to the sensor VM. When using an SSH session, the default username is *sysadmin*.

   Important: If you are accessing a Microsoft Azure Sensor through SSH and you specified a username other than the default (*sysadmin*) for your SSH access, you must use the following commands at the command line to "sudo up" and access the sensor console:

   ```
   # sudo su - sysadmin
   ```

2. From the USM Anywhere Sensor console System Menu, select **Maintenance** and press **Enter**.

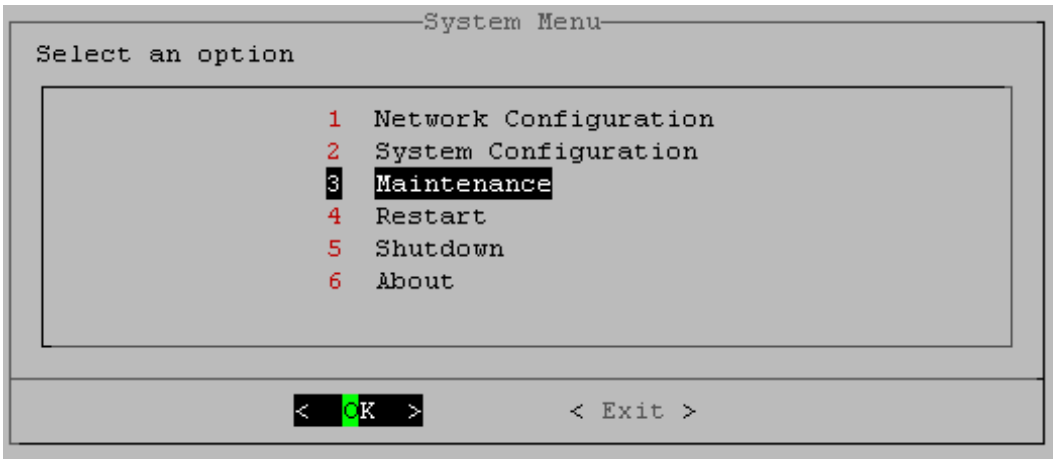

3. From the Maintenance menu, select **Debugging Tools** and press **Enter**.

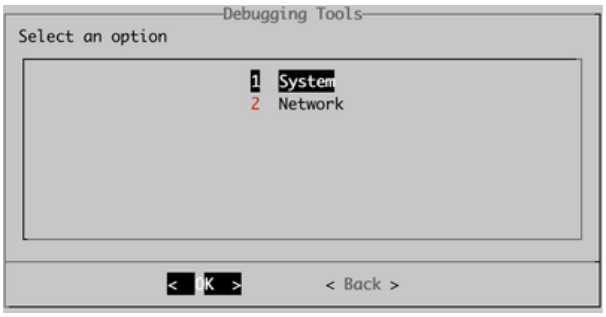

4. From the Debugging Tools menu, select **Network** and press **Enter**.

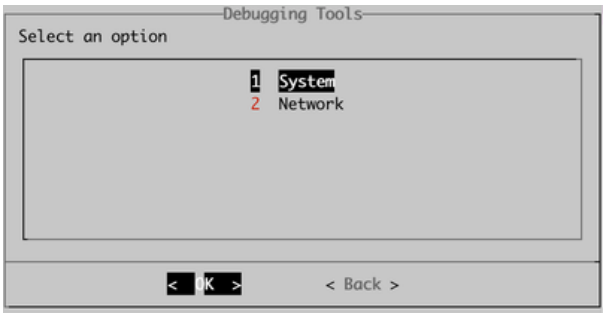

5. From the Network menu, select **Network Test** and press **Enter**.

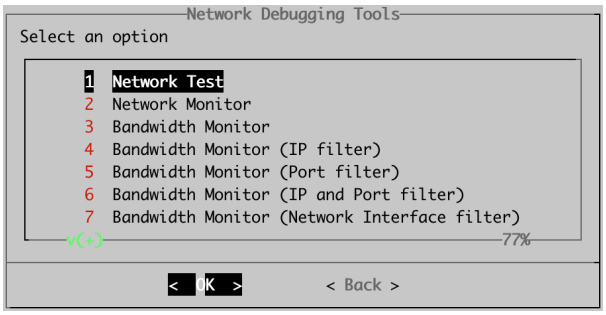

The network test runs and displays the test results.

Network Test Results

The test displays a screen with the results of the sensor's network test. There are seven tests that are displayed. Each test displays a SUCCESS or ERROR result.

```
--------------- NETWORK TESTS ---------------
-------- Mon Sep 27 13:56:11 UTC 2021 -------
----------------------------------------------

---------------- Sensor Info ----------------
Sensor Name     = MySensor
Deployment Type = AZURE
Sensor Version  = 7.11.830
Sensor Stats    = 13:56:11 up  6:37,  2 users,  load average: 0.34, 0.18, 0.17
RAM Capacity    = 7.1G
vCPUs           = 2
----------------------------------------------

------------- Ping Gateway Test -------------
 X  Ping to default gateway (10.0.0.1): ERROR
----------------------------------------------

------------ DNS resolution Test ------------
 V  DNS test to Control Node SUBDOMAIN.alienvault.cloud (1.2.3.4): SUCCESS
----------------------------------------------

----------- SSL Certificate Test ------------
 V  Testing SSL Certificate of the SUBDOMAIN.alienvault.cloud on port 7100: SUCCESS
----------------------------------------------

----------- TCP Connections Test ------------
 V  Testing connection to SUBDOMAIN.alienvault.cloud on port 7100: SUCCESS
 V  Testing connection to SUBDOMAIN.alienvault.cloud on port 443: SUCCESS
 V  Testing connection to update.alienvault.cloud on port 443: SUCCESS
 V  Testing connection to reputation.alienvault.com on port 443: SUCCESS
 V  Testing connection to otx.alienvault.com on port 443: SUCCESS
----------------------------------------------

------ Remote Support Connection Test -------
 V  Testing connection to prod-usm-saas-tractorbeam.alienvault.cloud on port 22: SUCCESS
 V  Testing connection to prod-usm-saas-tractorbeam.alienvault.cloud on port 443: SUCCESS
----------------------------------------------

Press any key to continue...
```

This table lists the individual tests with a potential diagnosis for a test failure.

Individual Tests with a Potential Diagnosis for a Test Failure

Test	Purpose	Failure Diagnosis
`Ping to default gateway`	This test determines if the sensor can ping its default gateway or router.	If this test fails, confirm that the sensor is using the correct default gateway and subnet. An error for this test results in a sensor connection failure. **Important:** Sometimes this test can fail because some providers don't allow users to ping their gateway.
`DNS test to Control Node <your subdomain>`	This test determines if the sensor can resolve the IP of the USM Anywhere subdomain.	If this test fails, confirm that the sensor is using the correct Domain Name System (DNS) server and can resolve the IP address of the domain. This can be tested from another machine using the following command: `nslookup <SUBDOMAIN> <DNS-IP-Address>` An error for this test results in a sensor connection failure.

Individual Tests with a Potential Diagnosis for a Test Failure (Continued)

Test	Purpose	Failure Diagnosis
`Testing connection to Control Node port 443`	This test determines whether a full TCP connection is possible to the domain on port 443.	If all previous tests are successful, this test can fail due to a firewall or a similar device blocking the connection. An error for this test results in a sensor connection failure. See USM Anywhere Deployment Requirements for sensor configuration requirements.
`Testing connection to Control Node port 7100`	This test determines whether a full TCP connection is possible to the domain on port 7100.	If all previous tests are successful, this test can fail due to a firewall or similar device blocking the connection. An error for this test results in a sensor connection failure. See USM Anywhere Deployment Requirements for sensor configuration requirements.

Individual Tests with a Potential Diagnosis for a Test Failure (Continued)

Test	Purpose	Failure Diagnosis
`Testing SSL Certificate of the Control Node`	This test determines whether the OpenSSL certificate is being returned.	If this test fails, it is most likely due to a firewall or proxy duplicating the OpenSSL certificate. An error for this test results in a sensor connection failure. See USM Anywhere Deployment Requirements for sensor configuration requirements.
`Testing connection to update.alienvault.cloud on por 443`	This test determines whether a full TCP connection is possible to update.alienvault.com on port 443.	If all previous tests are successful, this test can fail due to a firewall or similar device blocking the connection. A failure means that the initial setup will fail and future updates of the sensor will also fail. See USM Anywhere Deployment Requirements for sensor configuration requirements.

Individual Tests with a Potential Diagnosis for a Test Failure (Continued)

Test	Purpose	Failure Diagnosis
`Testing connection to reputation.alienvault.com on port 443`	This test determines whether a full TCP connection is possible to reputation.alienvault.com on port 443.	If all previous tests are successful, this test can fail due to a firewall or a similar device blocking the connection. A failure means that communication with the AT&T Alien Labs™ team threat intelligence can't be successful. See USM Anywhere Deployment Requirements for sensor configuration requirements.
`Testing connection to otx.alienvault.com on port 443`	This test determines whether a full TCP connection is possible to otx.alienvault.com on port 443.	If all previous tests are successful, this test can fail due to a firewall or a similar device blocking the connection. A failure means that communication with AT&T Alien Labs™ Open Threat Exchange® (OTX™) and that OTX threat intelligence can't be downloaded. See USM Anywhere Deployment Requirements for sensor configuration requirements.

Individual Tests with a Potential Diagnosis for a Test Failure (Continued)

Test	Purpose	Failure Diagnosis
`Testing connection to prod-usm-saas-tractorbeam.alienvault.cloud on port 22`	This test determines whether a full TCP connection is possible to usm-saas-tractorbeam.alienvault.cloud on port 22.	If all previous tests are successful, this test can fail due to a firewall or similar device blocking the connection. A failure means that a support session can't be open to the sensor. See USM Anywhere Deployment Requirements for sensor configuration requirements.
`Testing connection to prod-usm-saas-tractorbeam.alienvault.cloud on port 443`	This test determines whether a full TCP connection is possible to usm-saas-tractorbeam.alienvault.cloud on port 443.	If all previous tests are successful, this test can fail due to a firewall or similar device blocking the connection. A failure means that a support session can't be open to the sensor. See USM Anywhere Deployment Requirements for sensor configuration requirements

Retrieve Unique Identifier Information

In some instances, you may need to find the Unique Identifier (UID) of an asset, alarm, or event for a support investigation ticket. Each asset, alarm, and event has its own UID locatied in the URL of the specific alarm.

To locate the UID of an Asset

1. Go to **Environment > Assets.**
2. Search or filter the results on the page to the locate the asset for which you need the UID.
3. Click the ⌄ next to the desired Asset Name to open the drop-down menu, and then click **Full Details**.
4. In the URL for this page, you will see a string of characters at the end of the link. This string is the UID.

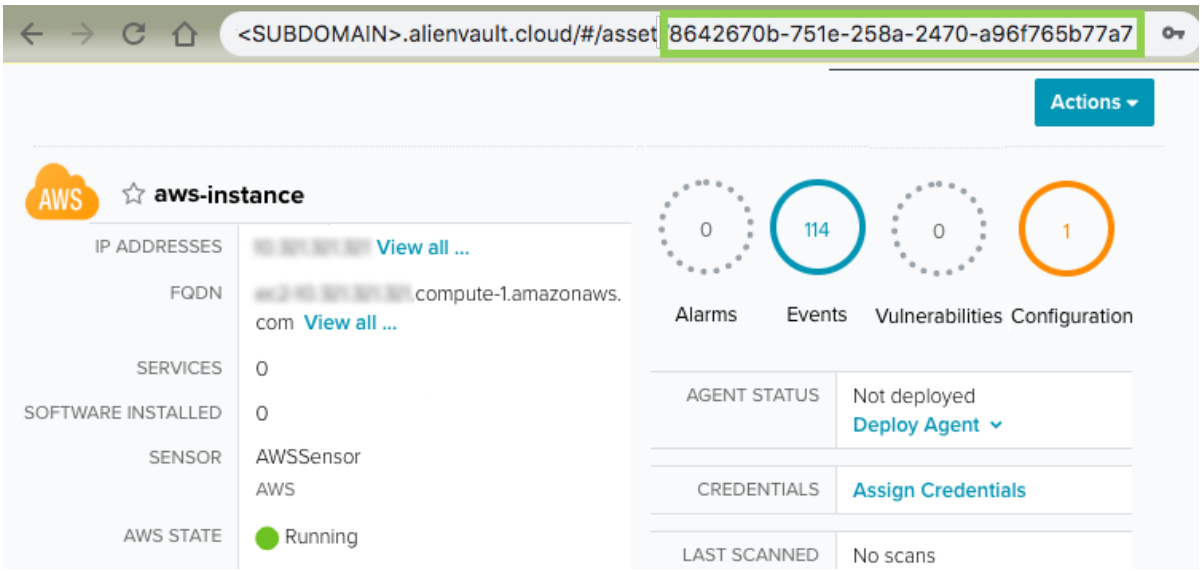

To locate the UID of an Alarm

1. Go to **Activity > Alarms.**
2. Search or filter the results on the page to the locate the alarm for which you need the UID.
3. Click the desired alarm to bring up the alarm summary view.
4. In the summary view click on the name of the alarm at the top of the page to open the full alarm details page.

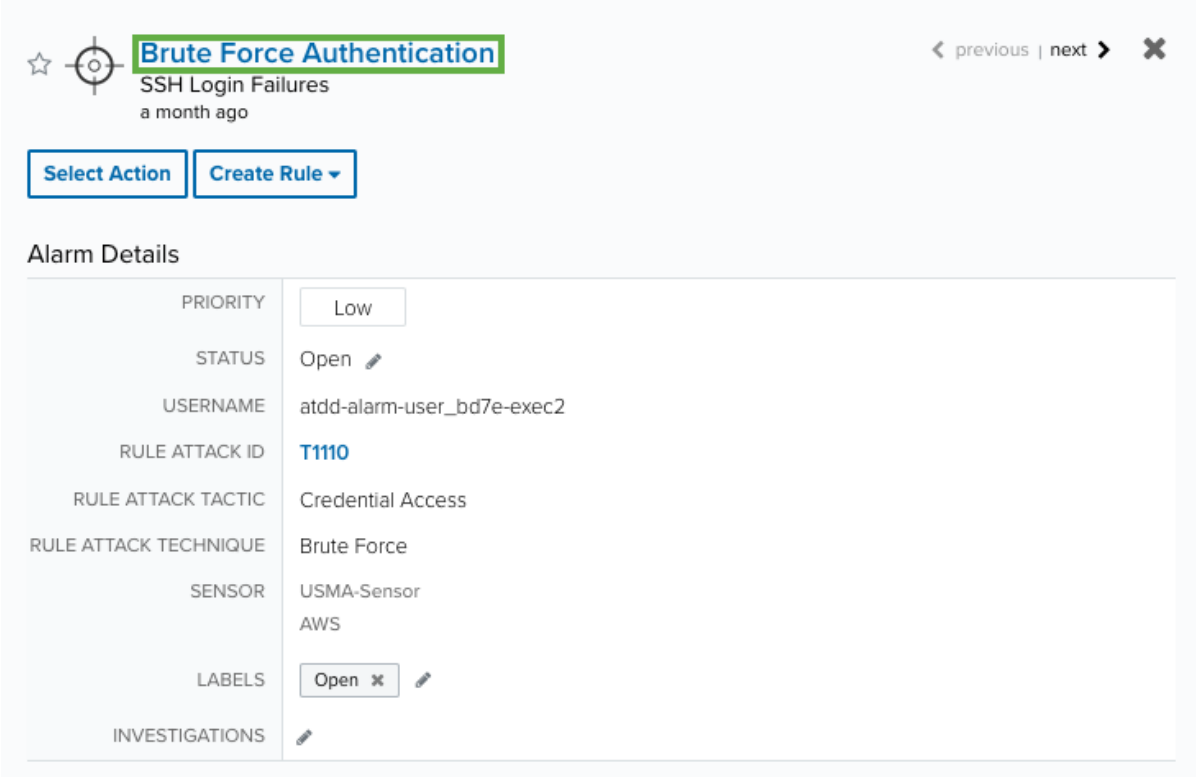

5. In the URL for this page, you will see a string of characters at the end of the link. This string is the UID.

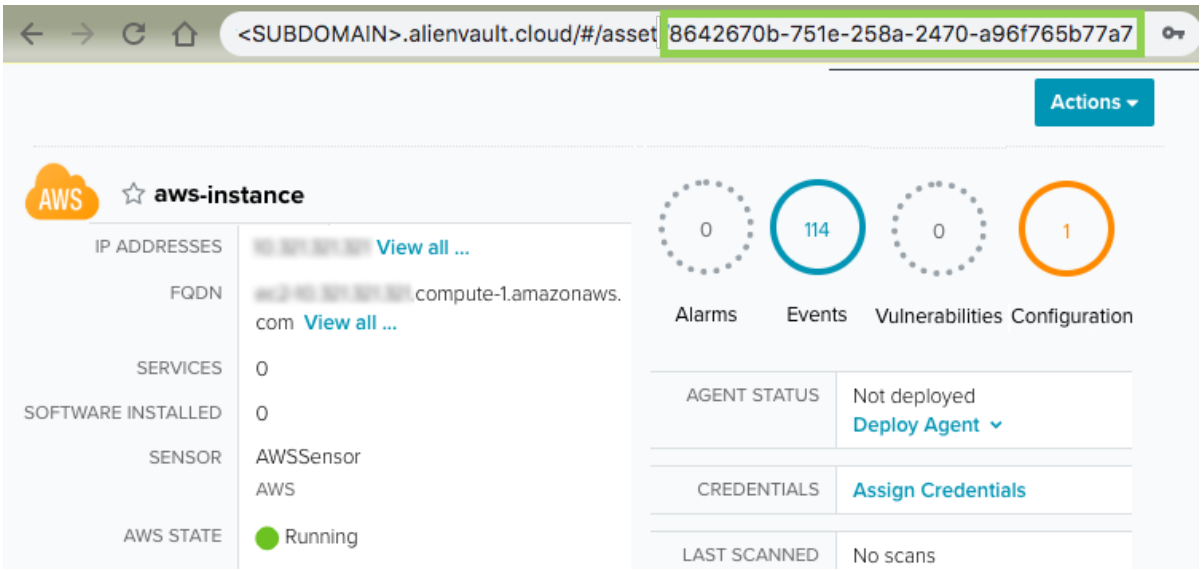

To locate the UID of an Event

1. Go to **Activity > Events.**
2. Search or filter the results on the page to the locate the event for which you need the UID.
3. Click the desired event to bring up the event summary view.
4. In the summary view click on the name of the event at the top of the page to open the full event details page.

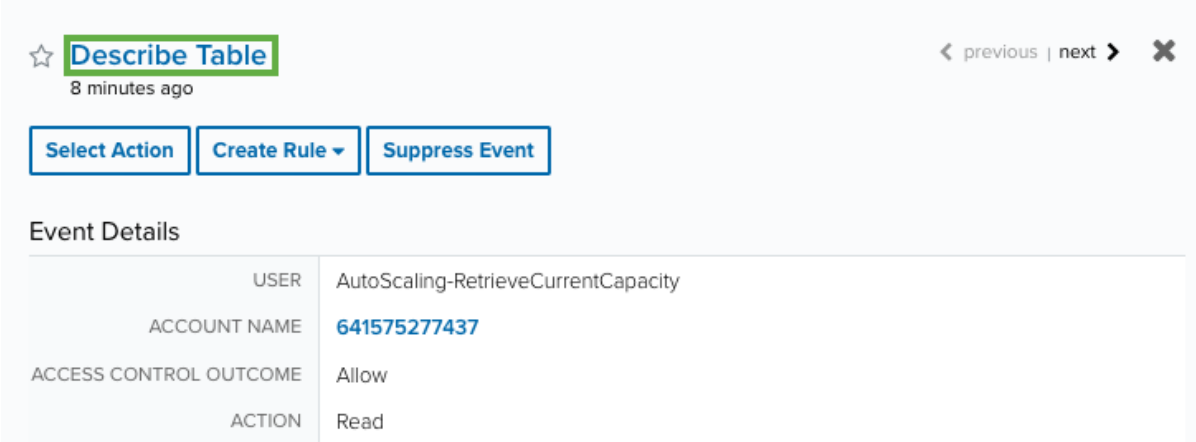
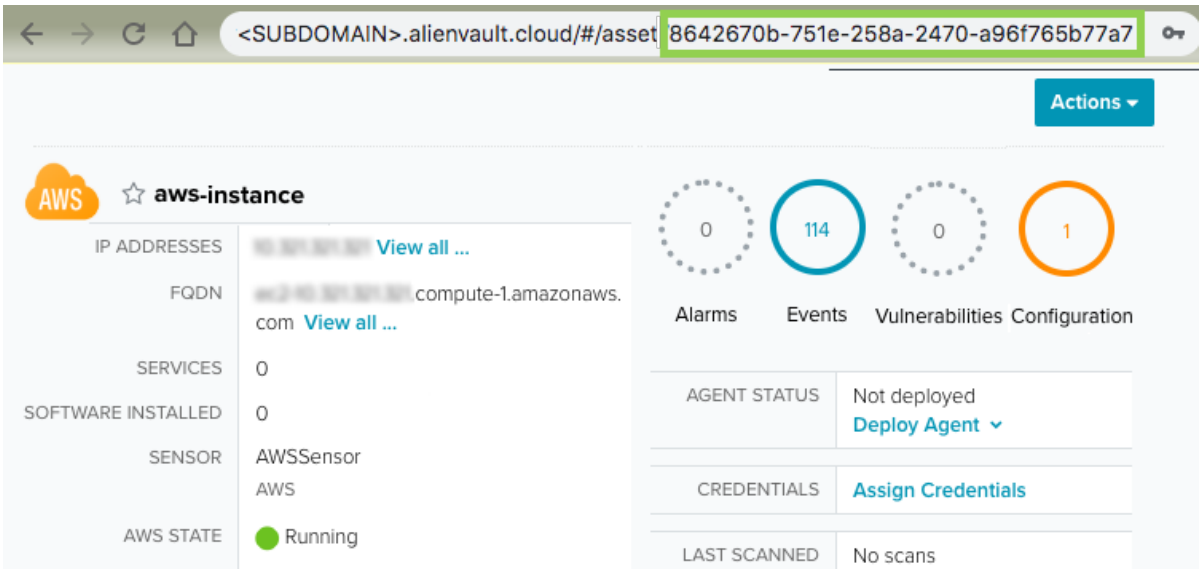
5. In the URL for this page, you will see a string of characters at the end of the link. This string is the UID.