



Workload Security

Data Infrastructure Insights

NetApp
February 17, 2026

Table of Contents

Workload Security	1
About Storage Workload Security	1
Visibility	1
Protection	1
Compliance	1
Getting Started	1
Getting Started with Workload Security	1
Workload Security Agent Requirements	2
Deploy Workload Security Agents	6
Deleting a Workload Security Agent	14
Configuring an Active Directory (AD) User Directory Collector	15
Configuring the ONTAP SVM Data Collector	20
Troubleshooting the ONTAP SVM Data Collector	30
Configuring the Cloud Volumes ONTAP and Amazon FSx for NetApp ONTAP collector	36
User Management	38
Event Rate Checker: Agent Sizing Guide	39
Understanding and Investigating Alerts	42
Alert	43
Filter Options	44
The Alert Details page	44
<i>Take a Snapshot</i> Action	46
Alert Notifications	47
Retention Policy	47
Troubleshooting	48
Forensics	48
Forensics - All Activity	48
Forensic User Overview	57
Automated Response Policies	58
Allowed File Types Policies	60
Integration with ONTAP Autonomous Ransomware Protection	61
Prerequisites	61
User permissions required	62
Sample Alert	62
Limitations	63
Troubleshooting	63
Integration with ONTAP Access Denied	64
Prerequisites	64
User permissions required	65
Access Denied events	65
Blocking User Access to Stop Attacks	66
Prerequisites for User Access Blocking	66
How to enable the feature?	67
How to set up Automatic user access blocking?	67

How to know if there are blocked users in the system?	67
Restrict and manage user access manually	67
User Access Limitation History	67
How to disable the feature?	68
Manually Restore IPs for NFS	68
Manually Restore Users for SMB	69
Troubleshooting	70
Workload Security: Simulating File Tampering	72
Things to note before you begin	72
Guidelines:	72
Steps:	72
Generate the sample files programmatically:	73
Resume the collector	74
Generate the sample files programmatically:	74
Generate an Alert in Workload Security	75
Triggering alert multiple times	75
Configuring Email Notifications for Alerts, Warnings, and Agent/Data Source Collector health	76
Potential Attack Alerts and Warnings	76
Agent and Data Collector Health monitoring	76
Receiving Agent And Data Collector Upgrade Notifications	76
Troubleshooting	76
Webhook Notifications	77
Workload Security notifications using webhooks	77
Workload Security Webhook Example for Discord	82
Workload Security Webhook Example for PagerDuty	85
Workload Security Webhook Example for Slack	89
Workload Security Webhook Example for Microsoft Teams	92
Workload Security API	98
API Documentation (Swagger)	98
API Access Tokens	98
Script to extract data via the API	99
Troubleshooting the ONTAP SVM Data Collector	99

Workload Security

About Storage Workload Security

Data Infrastructure Insights Storage Workload Security (formerly Cloud Secure) helps protect your data with actionable intelligence on insider threats. It provides centralized visibility and control of all corporate data access across hybrid cloud environments to ensure security and compliance goals are met.

Visibility

Gain centralized visibility and control of user access to your critical corporate data stored on-premise or in the cloud.

Replace tools and manual processes that fail to provide timely and accurate visibility into data access and control. Workload Security uniquely operates on both cloud and on-premise storage systems to give you real-time alerts of malicious user behavior.

Protection

Protect organizational data from being misused by malicious or compromised users through advanced machine learning and anomaly detection.

Alerts you to any abnormal data access through advanced machine learning and anomaly detection of user behavior.

Compliance

Ensure corporate compliance by auditing user data access to your critical corporate data stored on-premise or in the cloud.

Getting Started

Getting Started with Workload Security

Workload Security helps you monitor user activity and detect potential security threats in your storage environment. Before you can begin monitoring, you need to configure agents, data collectors, and directory services to establish the foundation for comprehensive security monitoring.

The Workload Security system uses an agent to collect access data from storage systems and user information from Directory Services servers.

You need to configure the following before you can start collecting data:

Task	Related information
------	---------------------

Configure an Agent	Agent Requirements Add Agent
Configure a User Directory Connector	Add User Directory Connector
Configure data collectors	Click Workload Security > Collectors Click the data collector you want to configure. See the Data Collector Vendor Reference section of the documentation for collector information.
Create Users Accounts	Manage User Accounts

Workload Security can integrate with other tools as well. For example, [see this guide](#) on integration with Splunk.

Workload Security Agent Requirements

Deploy Workload Security Agents on dedicated servers that meet minimum OS, CPU, memory, and disk space requirements to ensure optimal monitoring and threat detection performance. This guide specifies the hardware and network requirements needed before [installing your Workload Security Agent](#), including supported Linux distributions, network connectivity rules, and system sizing guidance.

Component	Linux Requirement
Operating system	A computer running a licensed version of one of the following: <ul style="list-style-type: none"> * AlmaLinux 9.4 (64bit) thru 9.5 (64bit), 10 (64bit), including SELinux * CentOS Stream 9 (64-bit) * Debian 11 (64-bit), 12 (64-bit), including SELinux * OpenSUSE Leap 15.3 (64-bit) through 15.6 (64-bit) * Oracle Linux 8.10 (64-bit), 9.1 (64-bit) through 9.6 (64-bit), including SELinux * Red Hat Enterprise Linux 8.10 (64bit), 9.1 (64bit) thru 9.6 (64bit), 10 (64bit), including SELinux * Rocky 9.4 (64bit) thru 9.6 (64bit), including SELinux * SUSE Linux Enterprise Server 15 SP4 (64-bit) through 15 SP6 (64-bit), including SELinux * Ubuntu 20.04 LTS (64-bit), 22.04 LTS (64-bit), 24.04 LTS (64-bit) This computer should be running no other application-level software. A dedicated server is recommended.
Commands	'unzip' is required for installation. Additionally, the 'sudo su -' command is required for installation, running scripts, and uninstall.
CPU	4 CPU cores
Memory	16 GB RAM

Component	Linux Requirement
Available disk space	<p>Disk space should be allocated in this manner: <code>/opt/netapp</code> 36 GB (minimum 35 GB free space after filesystem creation)</p> <p>Note: It is recommended to allocate a little extra disk space to allow for the creation of the filesystem. Ensure that there is at least 35 GB free space in the filesystem.</p> <p>If <code>/opt</code> is a mounted folder from a NAS storage, make sure that local users have access to this folder. Agent or Data collector may fail to install if local users do not have permission to this folder. see the troubleshooting section for more details.</p>
Network	<p>100 Mbps to 1 Gbps Ethernet connection, static IP address, IP connectivity to all devices, and a required port to the Workload Security instance (80 or 443).</p>

Please note: The Workload Security agent can be installed in the same machine as a Data Infrastructure Insights acquisition unit and/or agent. However, it is a best practice to install these in separate machines. In the event that these are installed on the same machine, please allocate disk space as shown below:

Available disk space	<p>50-55 GB For Linux, disk space should be allocated in this manner: <code>/opt/netapp</code> 25-30 GB <code>/var/log/netapp</code> 25 GB</p>
----------------------	---

Additional recommendations

- It is strongly recommended to synchronize the time on both the ONTAP system and the Agent machine using **Network Time Protocol (NTP)** or **Simple Network Time Protocol (SNTP)**.

Cloud Network Access Rules

For **US-based** Workload Security environments:

Protocol	Port	Source	Destination	Description
TCP	443	Workload Security Agent	<code><site_name>.cs01.cloudinsights.netapp.com</code> <code><site_name>.c01.cloudinsights.netapp.com</code> <code><site_name>.c02.cloudinsights.netapp.com</code>	Access to Data Infrastructure Insights
TCP	443	Workload Security Agent	<code>agentlogin.cs01.cloudinsights.netapp.com</code>	Access to authentication services

For **Europe-based** Workload Security environments:

Protocol	Port	Source	Destination	Description
TCP	443	Workload Security Agent	<site_name>.cs01-eu-1.cloudinsights.netapp.com <site_name>.c01-eu-1.cloudinsights.netapp.com <site_name>.c02-eu-1.cloudinsights.netapp.com	Access to Data Infrastructure Insights
TCP	443	Workload Security Agent	agentlogin.cs01-eu-1.cloudinsights.netapp.com	Access to authentication services

For **APAC-based** Workload Security environments:

Protocol	Port	Source	Destination	Description
TCP	443	Workload Security Agent	<site_name>.cs01-ap-1.cloudinsights.netapp.com <site_name>.c01-ap-1.cloudinsights.netapp.com <site_name>.c02-ap-1.cloudinsights.netapp.com	Access to Data Infrastructure Insights
TCP	443	Workload Security Agent	agentlogin.cs01-ap-1.cloudinsights.netapp.com	Access to authentication services

In-network rules

Protocol	Port	Source	Destination	Description
TCP	389(LDAP) 636 (LDAPs / start-tls)	Workload Security Agent	LDAP Server URL	Connect to LDAP
TCP	443	Workload Security Agent	Cluster or SVM Management IP Address (depending on SVM collector configuration)	API communication with ONTAP

Protocol	Port	Source	Destination	Description
TCP	35000 - 55000	SVM data LIF IP Addresses	Workload Security Agent	<p>Communication from ONTAP to the Workload Security Agent for Fpolicy events. These ports must be opened towards the Workload Security Agent in order for ONTAP to send events to it, including any firewall on the Workload Security Agent itself (if present).</p> <p>NOTE that you do not need to reserve all of these ports, but the ports you reserve for this must be within this range. It is recommended to start by reserving ~100 ports, and increasing if necessary.</p>

Protocol	Port	Source	Destination	Description
TCP	35000-55000	Cluster Management IP	Workload Security Agent	<p>Communication from ONTAP Cluster Management IP to the Workload Security Agent for EMS events. These ports must be opened towards the Workload Security Agent in order for ONTAP to send EMS events to it, including any firewall on the Workload Security Agent itself (if present).</p> <p>NOTE that you do not need to reserve all of these ports, but the ports you reserve for this must be within this range. It is recommended to start by reserving ~100 ports, and increasing if necessary.</p>
SSH	22	Workload Security Agent	Cluster management	Needed for CIFS/SMB user blocking.

System Sizing

See the [Event Rate Checker](#) documentation for information about sizing.

Deploy Workload Security Agents

Workload Security agents are essential for monitoring user activity and detecting potential security threats across your storage infrastructure. This guide provides step-by-step installation instructions, best practices for agent management (including pause/resume and pin/unpin capabilities), and post-deployment configuration requirements. Before you begin, ensure your agent server meets the [system requirements](#).

Before You Begin

- The sudo privilege is required for installation, running scripts, and uninstall.
- While installing the agent, a local user `cssys` and a local group `cssys` are created on the machine. If permission settings do not allow creation of a local user, and instead require Active Directory, a user with the username `cssys` must be created in the Active Directory server.

- You can read about Data Infrastructure Insights security [here](#).

Best Practices

Keep the following in mind before configuring your Workload Security agent.

Pause and Resume	Pause: Removes fpolicies from ONTAP. Typically used when customers perform extended maintenance activities which may take significant time, such as agent VM reboots or storage replacements. Resume: Adds fpolicies back to ONTAP.
Pin and Unpin	Unpin immediately fetches the latest version (if available) and upgrades the agent and collector. During this upgrade, fpolicies will disconnect and reconnect. This feature is designed for customers who want to control the timing of automatic upgrades. See below for pin/unpin instructions .
Recommended Approach	For large configurations, it is advisable to use Pin and Unpin rather than pausing collectors. There is no need to pause and resume while using pin and unpin. Customers can keep their agents and collectors pinned, and upon receiving an email notification about a new version, have a 30-day window to selectively upgrade agents one by one. This approach minimizes latency impact on fpolicies and provides greater control over the upgrade process.

Steps to Install Agent

1. Log in as Administrator or Account Owner to your Workload Security environment.
2. Select **Collectors > Agents > +Agent**

The system displays the Add an Agent page:

Cloud Secure collects device and user data using one or more Agents installed on local servers. Each Agent can host multiple Data Collectors, which send data to Cloud Secure for analysis.

Which Operating system are you using ?

CentOS RHEL

Close

3. Verify that the agent server meets the minimum system requirements.
4. To verify that the agent server is running a supported version of Linux, click *Versions Supported (i)*.
5. If your network is using proxy server, please set the proxy server details by following the instructions in the Proxy section.

Add an Agent



Cloud Secure collects device and user data using one or more Agents installed on local servers. Each Agent can host multiple Data Collectors, which send data to Cloud Secure for analysis.

Agent Server Requirements

Linux Versions Supported: [?](#) Minimum Server Requirements: [?](#)

Installation Instructions

Need Help?

Open up a terminal window and run the following commands:

1. If a proxy server is used, please enter these proxy server settings after editing in your proxy variables. [?](#)

```
export https proxy='USER:PASSWORD@PROXY SERVER:PORT'
```



2. Enter this agent installation command.

This snippet has a unique key valid for 2 hours and for one Agent only.

6. Click the Copy to Clipboard icon to copy the installation command.
7. Run the installation command in a terminal window.
8. The system displays the following message when the installation completes successfully:



After You Finish

1. You need to configure a [User Directory Collector](#) .
2. You need to configure one or more Data Collectors.

Network Configuration

Run the following commands on the local system to open ports that will be used by Workload Security. If there is a security concern regarding the port range, you can use a lesser port range, for example 35000:35100. Each SVM uses two ports.

Steps

1. `sudo firewall-cmd --permanent --zone=public --add-port=35000-55000/tcp`
2. `sudo firewall-cmd --reload`

Follow the next steps according to your platform:

CentOS 7.x / RHEL 7.x:

1. `sudo iptables-save | grep 35000`

Sample output:

```
-A IN_public_allow -p tcp -m tcp --dport 35000:55000 -m conntrack -ctstate NEW,UNTRACKED -j ACCEPT
```

CentOS 8.x / RHEL 8.x:

1. `sudo firewall-cmd --zone=public --list-ports | grep 35000` (for CentOS 8)

Sample output:

```
35000-55000/tcp
```

"Pinning" an Agent at the current version

By default, Data Infrastructure Insights Workload Security updates agents automatically. Some customers may wish to pause automatic updating, which leaves an Agent at its current version until one of the following occurs:

- The customer resumes automatic Agent updates.
- 30 days have passed. Note that the 30 days starts on the day of the most recent Agent update, not at the day the Agent is paused.

In each of these cases, the agent will be updated at the next Workload Security refresh.

To pause or resume automatic agent updates, use the `cloudsecure_config.agents` APIs:

cloudsecure_config.agents		
GET	/v1/cloudsecure/agents	Retrieve all agents.
POST	/v1/cloudsecure/agents/configuration	Pin all agents under tenant
DELETE	/v1/cloudsecure/agents/configuration	Unpin all agents under tenant
POST	/v1/cloudsecure/agents/{agentId}/configuration	Pin an agent under tenant
DELETE	/v1/cloudsecure/agents/{agentId}/configuration	Unpin an agent under tenant
GET	/v1/cloudsecure/agents/{agentUuid}	Retrieve an agent by agentUuid.

Note that it may take up to five minutes for the pause or resume action to take effect.

You can view your current Agent versions on the **Workload Security > Collectors** page, in the **Agents** tab.

Installed Agents (15)

Name ↑	IP Address	Version	Status
agent-1396	10.128.218.124	1.625.0	Connected

Troubleshooting Agent Errors

Known problems and their resolutions are described in the following table.

Problem:	Resolution:
Agent installation fails to create the /opt/netapp/cloudsecure/agent/logs/agent.log folder and the install.log file provides no relevant information.	This error occurs during bootstrapping of the agent. The error is not logged in log files because it occurs before logger is initialized. The error is redirected to standard output, and is visible in the service log using the journalctl -u cloudsecure-agent.service command. This command can be used for troubleshooting the issue further.
Agent installation fails with 'This linux distribution is not supported. Exiting the installation'.	This error appears when you attempt to install the Agent on an unsupported system. See Agent Requirements .
Agent Installation failed with the error: "-bash: unzip: command not found"	Install unzip and then run the installation command again. If Yum is installed on the machine, try "yum install unzip" to install unzip software. After that, re-copy the command from the Agent installation UI and paste it in the CLI to execute the installation again.

Problem:	Resolution:
Agent was installed and was running. However agent has stopped suddenly.	<p>SSH to the Agent machine. Check the status of the agent service via <code>sudo systemctl status cloudsecure-agent.service</code>.</p> <ol style="list-style-type: none"> 1. Check if the logs shows a message “Failed to start Workload Security daemon service” . 2. Check if <code>cssys</code> user exists in the Agent machine or not. Execute the following commands one by one with root permission and check if the <code>cssys</code> user and group exists. <pre> sudo id cssys sudo groups cssys </pre> <ol style="list-style-type: none"> 3. If none exists, then a centralized monitoring policy may have deleted the <code>cssys</code> user. 4. Create <code>cssys</code> user and group manually by executing the following commands. <pre> sudo useradd cssys sudo groupadd cssys </pre> <ol style="list-style-type: none"> 5. Restart the agent service after that by executing the following command: <pre>sudo systemctl restart cloudsecure-agent.service</pre> <ol style="list-style-type: none"> 6. If it is still not running, please check the other troubleshooting options.
Unable to add more than 50 Data collectors to an Agent.	Only 50 Data collectors can be added to an Agent. This can be a combination of all the collector types, for example, Active Directory, SVM and other collectors.
UI shows Agent is in NOT_CONNECTED state.	<p>Steps to restart the Agent.</p> <ol style="list-style-type: none"> 1. SSH to the Agent machine. 2. Restart the agent service after that by executing the following command: <pre>sudo systemctl restart cloudsecure-agent.service</pre> <ol style="list-style-type: none"> 3. Check the status of the agent service via <code>sudo systemctl status cloudsecure-agent.service</code>. 4. Agent should go to CONNECTED state.
Agent VM is behind Zscaler proxy and the agent installation is failing. Because of Zscaler proxy's SSL inspection, the Workload Security certificates are presented as it is signed by Zscaler CA so the agent is not trusting the communication.	Disable SSL inspection in the Zscaler proxy for the <code>*.cloudinsights.netapp.com</code> url. If Zscaler does SSL inspection and replaces the certificates, Workload Security will not work.

Problem:	Resolution:
While installing the agent, the installation hangs after unzipping.	<p>“chmod 755 -Rf” command is failing. The command fails when the agent installation command is being run by a non-root sudo user that has files in the working directory, belonging to another user, and permissions of those files cannot be changed. Because of the failing chmod command, the rest of the installation does not execute.</p> <ol style="list-style-type: none"> 1. Create a new directory named “cloudsecure”. 2. Go to that directory. 3. Copy and paste the full “token=..... ./cloudsecure-agent-install.sh” installation command and press enter. 4. Installation should be able to proceed.
If the Agent is still not able to connect to SaaS, please open a case with NetApp Support. Provide the Data Infrastructure Insights serial number to open a case, and attach logs to the case as noted.	<p>To attach logs to the case:</p> <ol style="list-style-type: none"> 1. Execute the following script with root permission and share the output file (cloudsecure-agent-symptoms.zip). a. /opt/netapp/cloudsecure/agent/bin/cloudsecure-agent-symptom-collector.sh 2. Execute the following commands one by one with root permission and share the output. a. id cssys b. groups cssys c. cat /etc/os-release
<p>The cloudsecure-agent-symptom-collector.sh script fails with the following error.</p> <pre data-bbox="132 1178 817 1695"> [root@machine tmp]# /opt/netapp/cloudsecure/agent/bin/cloudsecure-agent-symptom-collector.sh Collecting service log Collecting application logs Collecting agent configurations Taking service status snapshot Taking agent directory structure snapshot /opt/netapp/cloudsecure/agent/bin/cloudsecure-agent-symptom-collector.sh: line 52: zip: command not found ERROR: Failed to create /tmp/cloudsecure-agent-symptoms.zip </pre>	<p>Zip tool is not installed.. Install the zip tool by running the command “yum install zip”.</p> <p>Then run the cloudsecure-agent-symptom-collector.sh again.</p>

Problem:	Resolution:
<p>Agent installation Fails with useradd: cannot create directory /home/cssys</p>	<p>This error can occur if user's login directory cannot be created under /home, due to lack of permissions.</p> <p>The workaround would be to create cssys user and add its login directory manually using the following command:</p> <p><code>sudo useradd user_name -m -d HOME_DIR</code></p> <p>-m :Create the user's home directory if it does not exist.</p> <p>-d : The new user is created using HOME_DIR as the value for the user's login directory.</p> <p>For instance, <code>sudo useradd cssys -m -d /cssys</code>, adds a user cssys and creates its login directory under root.</p>
<p>Agent is not running after installation. <code>Systemctl status cloudsecure-agent.service</code> shows the following:</p> <pre>[root@demo ~]# systemctl status cloudsecure-agent.service agent.service – Workload Security Agent Daemon Service Loaded: loaded (/usr/lib/systemd/system/cloudsecure-agent.service; enabled; vendor preset: disabled) Active: activating (auto-restart) (Result: exit-code) since Tue 2021-08-03 21:12:26 PDT; 2s ago Process: 25889 ExecStart=/bin/bash /opt/netapp/cloudsecure/agent/bin/cloudsecure-agent (code=exited status=126) Main PID: 25889 (code=exited, status=126), Aug 03 21:12:26 demo systemd[1]: cloudsecure-agent.service: main process exited, code=exited, status=126/n/a Aug 03 21:12:26 demo systemd[1]: Unit cloudsecure-agent.service entered failed state. Aug 03 21:12:26 demo systemd[1]: cloudsecure-agent.service failed.</pre>	<p>This can be failing because cssys user may not have permission to install.</p> <p>If /opt/netapp is an NFS mount and if cssys user does not have access to this folder, installation will fail. cssys is a local user created by the Workload Security installer that may not have permission to access the mounted share.</p> <p>You can check this by attempting to access /opt/netapp/cloudsecure/agent/bin/cloudsecure-agent using cssys user. If it returns "Permission denied", installation permission is not present.</p> <p>Instead of a mounted folder, install on a directory local to the machine.</p>

Problem:	Resolution:
<p>Agent was initially connected via a proxy server and the proxy was set during Agent installation. Now the proxy server has changed. How can the Agent's proxy configuration be changed?</p>	<p>You can edit the <code>agent.properties</code> to add the proxy details. Follow these steps:</p> <ol style="list-style-type: none"> 1. Change to the folder containing the properties file: <code>cd /opt/netapp/cloudsecure/conf</code> 2. Using your favorite text editor, open the <code>agent.properties</code> file for editing. 3. Add or modify the following lines: <pre> AGENT_PROXY_HOST=scspa1950329001.vm.netapp.com AGENT_PROXY_PORT=80 AGENT_PROXY_USER=pxuser AGENT_PROXY_PASSWORD=pass1234 </pre> <ol style="list-style-type: none"> 4. Save the file. 5. Restart the agent: <code>sudo systemctl restart cloudsecure-agent.service</code>

Deleting a Workload Security Agent

When you delete a Workload Security Agent, all the data collectors associated with the Agent must be deleted first.

Deleting an Agent



Deleting an Agent deletes all of the Data Collectors associated with the Agent. If you plan to configure the data collectors with a different agent you should create a backup of the Data Collector configurations before you delete the Agent.

Before you begin

1. Make sure all the data collectors associated with the agent are deleted from the Workload Security portal.

Note: Ignore this step if all the associated collectors are in STOPPED state.

Steps to delete an Agent:

1. SSH into the agent VM and execute the following command. When prompted, enter "y" to continue.

```

sudo /opt/netapp/cloudsecure/agent/install/cloudsecure-agent-
uninstall.sh
Uninstall CloudSecure Agent? [y|N] :

```

2. Click **Workload Security > Collectors > Agents**

The system displays the list of configured Agents.

3. Click the options menu for the Agent you are deleting.
4. Click **Delete**.

The system displays the **Delete Agent** page.

5. Click **Delete** to confirm the deletion.

Configuring an Active Directory (AD) User Directory Collector

Workload Security can be configured to collect user attributes from Active Directory servers.

Before you begin

- You must be a Data Infrastructure Insights Administrator or Account Owner to perform this task.
- You must have the IP address of the server hosting the Active Directory server.
- An Agent must be configured before you configure a User Directory connector.

Steps to Configure a User Directory Collector

1. In the Workload Security menu, click:
Collectors > User Directory Collectors > + User Directory Collector and select **Active Directory**

The system displays the Add User Directory screen.

Configure the User Directory Collector by entering the required data in the following tables:

Name	Description
Name	Unique name for the user directory. For example <i>GlobalADCollector</i>
Agent	Select a configured agent from the list
Server IP/Domain Name	IP address or Fully-Qualified Domain Name (FQDN) of server hosting the active directory

Forest Name	<p>Forest level of the directory structure. Forest name allows both of the following formats:</p> <p><i>x.y.z</i> ⇒ direct domain name as you have it on your SVM. [Example: <i>hq.companyname.com</i>]</p> <p><i>DC=x,DC=y,DC=z</i> ⇒ Relative distinguished names [Example: <i>DC=hq,DC= companyname,DC=com</i>]</p> <p>Or you can specify as the following:</p> <p><i>OU=engineering,DC=hq,DC= companyname,DC=com</i> [to filter by specific OU <i>engineering</i>]</p> <p><i>CN=username,OU=engineering,DC=companyname,DC=netapp, DC=com</i> [to get only specific user with <username> from OU <engineering>]</p> <p><i>CN=Acrobat Users,CN=Users,DC=hq,DC=companyname,DC=com,O= companyname,L=Boston,S=MA,C=US</i> [to get all Acrobat Users within the Users in that organization]</p> <p>Trusted Active Directory domains are also supported.</p>
Bind DN	<p>User permitted to search the directory. For example: <i>username@companyname.com</i> or <i>username@domainname.com</i></p> <p>In addition, Domain Read Only permission is required. User must be a member of the Security group <i>Read-only Domain Controllers</i>.</p>
BIND password	Directory server password (i.e. password for <i>username</i> used in Bind DN)
Protocol	ldap, ldaps, ldap-start-tls
Ports	Select port

Add to table once link is provided:

For more details about forest names, please refer to this xref:////

Enter the following Directory Server required attributes if the default attribute names have been modified in LDAP Directory Server. Most often these attributes names are *not* modified in LDAP Directory Server, in which case you can simply proceed with the default attribute name.

Attributes	Attribute name in Directory Server
Display Name	name
UNIXID	uidnumber
User Name	uid

Click Include Optional Attributes to add any of the following attributes:

Attributes	Attribute Name in Directory Server
Email Address	mail
Telephone Number	telephonenumber
Role	title
Country	co
State	state
Department	departmentnumber
Photo	photo
ManagerDN	manager
Groups	memberOf

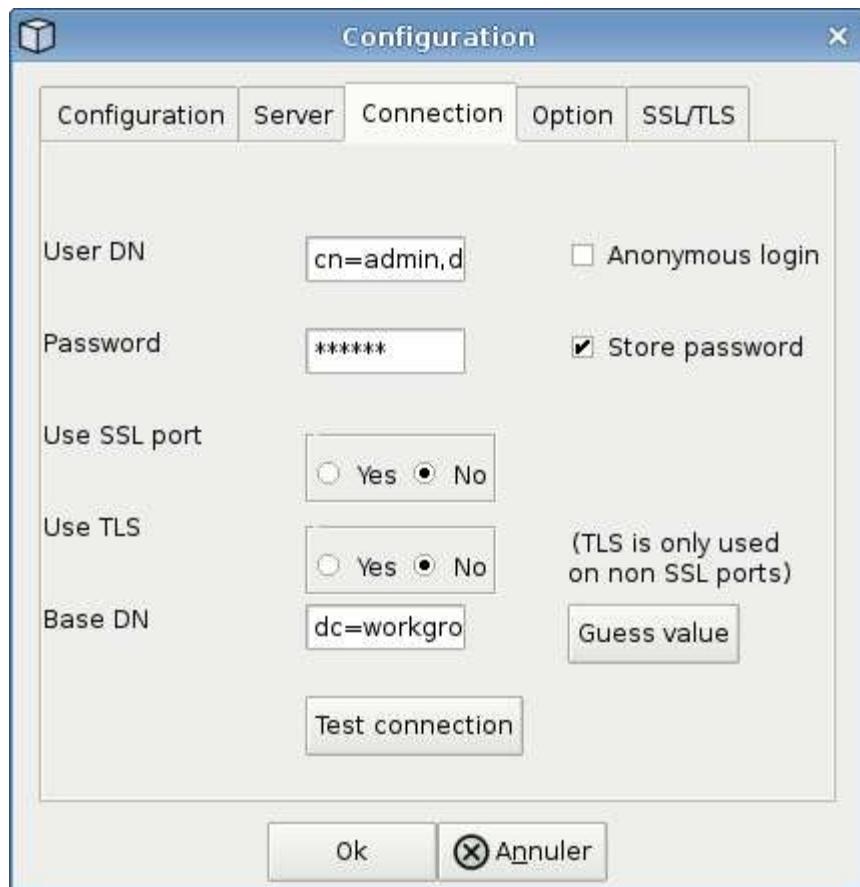
Testing Your User Directory Collector Configuration

You can validate LDAP User Permissions and Attribute Definitions using the following procedures:

- Use the following command to validate Workload Security LDAP user permission:

```
ldapsearch -D "uid=john,cn=users,cn=accounts,dc=dorp,dc=company,dc=com"
-W -x -LLL -o ldif-wrap=no -b "cn=accounts,dc=dorp,dc=company,dc=com" -H
ldap://vmwipaapp08.dorp.company.com
```

- Use LDAP Explorer to navigate an LDAP database, view object properties and attributes, view permissions, view an object's schema, execute sophisticated searches that you can save and re-execute.
 - Install LDAP Explorer (<http://ldaptool.sourceforge.net/>) or Java LDAP Explorer (<http://jxplorer.org/>) on any windows machine which can connect to the LDAP Server.
 - Connect to the LDAP server using the username/password of the LDAP directory server.



Troubleshooting LDAP Directory Collector Configuration Errors

The following table describes known problems and resolutions that can occur during collector configuration:

Problem:	Resolution:
Adding an LDAP Directory connector results in the 'Error' state. Error says, "Invalid credentials provided for LDAP server".	Incorrect Bind DN or Bind Password or Search Base provided. Edit and provide the correct information.
Adding an LDAP Directory connector results in the 'Error' state. Error says, "Failed to get the object corresponding to DN=DC=hq,DC=domainname,DC=com provided as forest name."	Incorrect Search Base provided. Edit and provide the correct forest name.
The optional attributes of domain user are not appearing in the Workload Security User Profile page.	This is likely due to a mismatch between the names of optional attributes added in CloudSecure and the actual attribute names in Active Directory. Fields are case sensitive. Edit and provide the correct optional attribute name(s).
Data collector in error state with "Failed to retrieve LDAP users. Reason for failure: Cannot connect on the server, the connection is null"	Restart the collector by clicking on the <i>Restart</i> button.

Problem:	Resolution:
Adding an LDAP Directory connector results in the 'Error' state.	<p>Ensure you have provided valid values for the required fields (Server, forest-name, bind-DN, bind-Password).</p> <p>Ensure bind-DN input is always provided as uid=ldapuser,cn=users,cn=accounts,dc=domain,dc=companyname,dc=com.</p>
Adding an LDAP Directory connector results in the 'RETRYING' state. Shows error "Failed to determine the health of the collector hence retrying again"	<p>Ensure correct Server IP and Search Base is provided</p> <p>////</p>
While adding LDAP directory the following error is shown: "Failed to determine the health of the collector within 2 retries, try restarting the collector again(Error Code: AGENT008)"	Ensure correct Server IP and Search Base is provided
Adding an LDAP Directory connector results in the 'RETRYING' state. Shows error "Unable to define state of the collector, reason Tcp command [Connect(localhost:35012,None,List(),Some(,seconds),true)] failed because of java.net.ConnectionException:Connection refused."	<p>Incorrect IP or FQDN provided for the AD Server. Edit and provide the correct IP address or FQDN.</p> <p>////</p>
Adding an LDAP Directory connector results in the 'Error' state. Error says, "Failed to establish LDAP connection".	<p>Incorrect IP or FQDN provided for the LDAP Server. Edit and provide the correct IP address or FQDN.</p> <p>Or</p> <p>Incorrect value for Port provided. Try using the default port values or the correct port number for the LDAP server.</p>
Adding an LDAP Directory connector results in the 'Error' state. Error says, "Failed to load the settings. Reason: Datasource configuration has an error. Specific reason: /connector/conf/application.conf: 70: ldap.ldap-port has type STRING rather than NUMBER"	Incorrect value for Port provided. Try using the default port values or the correct port number for the AD server.
I started with the mandatory attributes, and it worked. After adding the optional ones, the optional attributes data is not getting fetched from AD.	This is likely due to a mismatch between the optional attributes added in CloudSecure and the actual attribute names in Active Directory. Edit and provide the correct mandatory or optional attribute name.
After restarting the collector, when will the LDAP sync happen?	LDAP sync will happen immediately after the collector restarts. It will take approximately 15 minutes to fetch user data of approximately 300K users, and is refreshed every 12 hours automatically.
User Data is synced from LDAP to CloudSecure. When will the data be deleted?	User data is retained for 13months in case of no refresh. If the tenant is deleted then the data will be deleted.

Problem:	Resolution:
<p>LDAP Directory connector results in the 'Error' state. "Connector is in error state. Service name: usersLdap. Reason for failure: Failed to retrieve LDAP users. Reason for failure: 80090308: LdapErr: DSID-0C090453, comment: AcceptSecurityContext error, data 52e, v3839"</p>	<p>Incorrect forest name provided. See above on how to provide the correct forest name.</p>
<p>Telephone number is not getting populated in the user profile page.</p>	<p>This is most likely due to an attribute mapping problem with the Active Directory.</p> <ol style="list-style-type: none"> 1. Edit the particular Active Directory collector which is fetching the user's information from Active Directory. 2. Notice under optional attributes, there is a field name "Telephone Number" mapped to Active Directory attribute 'telephonenumber'. 4. Now, please use the Active Directory Explorer tool as described above to browse the LDAP Directory server and see the correct attribute name. 3. Make sure that in LDAP Directory there is an attribute named 'telephonenumber' which has indeed the telephone number of the user. 5. Let us say in LDAP Directory it has been modified to 'phonenumber'. 6. Then Edit the CloudSecure User Directory collector. In optional attribute section, replace 'telephonenumber' with 'phonenumber'. 7. Save the Active Directory collector, the collector will restart and get the telephone number of the user and display the same in the user profile page.
<p>If encryption certificate (SSL) is enabled on the Active Directory (AD) Server, the Workload Security User Directory Collector can not connect to the AD Server.</p>	<p>Disable AD Server encryption before Configuring a User Directory Collector. Once the user detail is fetched it will be there for 13 months. If the AD server gets disconnected after fetching the user details, the newly added users in AD won't get fetched. To fetch again the user directory collector needs to be connected to AD.</p>

Configuring the ONTAP SVM Data Collector

The ONTAP SVM Data Collector enables Workload Security to monitor file and user access activities on NetApp ONTAP storage virtual machines (SVMs). This guide walks you through the configuration and management of the SVM data collector to provide comprehensive security monitoring of your ONTAP environment.

Before you begin

- This data collector is supported with the following:
 - Data ONTAP 9.2 and later versions. For best performance, use a Data ONTAP version greater than 9.13.1.

- SMB protocol version 3.1 and earlier.
- NFS versions up to and including NFS 4.1 (Note that NFS 4.1 is supported with ONTAP 9.15 or later).
- Flexgroup is supported from ONTAP 9.4 and later versions
- FlexCache is supported for NFS with ONTAP 9.7 and later versions.
- FlexCache is supported for SMB with ONTAP 9.14.1 and later versions.
- ONTAP Select is supported
- Only data type SVMs are supported. SVMs with infinite volumes are not supported.
- SVM has several sub-types. Of these, only *default*, *sync_source*, and *sync_destination* are supported.
- An Agent [must be configured](#) before you can configure data collectors.
- Make sure that you have a properly configured User Directory Connector, otherwise events will show encoded user names and not the actual name of the user (as stored in Active Directory) in the “Activity Forensics” page.
- ONTAP Persistent Store is supported from 9.14.1.
- For optimal performance, you should configure the FPolicy server to be on the same subnet as the storage system.
- For comprehensive best practices and recommendations regarding Workload Security FPolicy configuration, see the [KB Article on FPolicy Best Practices](#).
- You must add an SVM using one of the following two methods:
 - By Using Cluster IP, SVM name, and Cluster Management Username and Password. ***This is the recommended method.***
 - SVM name must be exactly as is shown in ONTAP and is case-sensitive.
 - By Using SVM Vserver Management IP, Username, and Password
 - If you are not able or not willing to use the full Administrator Cluster/SVM Management Username and Password, you can create a custom user with lesser privileges as mentioned in the [“A note about permissions”](#) section below. This custom user can be created for either SVM or Cluster access.
 - You can also use an AD user with a role that has at least the permissions of csrole as mentioned in [“A note about permissions”](#) section below. Also refer to the [ONTAP documentation](#).
- Ensure the correct applications are set for the SVM by executing the following command:

```
clustershell:> security login show -vserver <vservername> -user-or-group  
-name <username>
```

Example output:

Authentication				Acct	Second Authentication
User/Group Name	Application	Method	Role Name	Locked	Method
vsadmin	http	password	vsadmin	no	none
vsadmin	ontapi	password	vsadmin	no	none
vsadmin	ssh	password	vsadmin	no	none

- Ensure that the SVM has a CIFS server configured:

```
clustershell:> vserver cifs show
```

The system returns the Vserver name, CIFS server name and additional fields.

- Set a password for the SVM vsadmin user. If using custom user or cluster admin user, skip this step.

```
clustershell:> security login password -username vsadmin -vserver svmname
```

- Unlock the SVM vsadmin user for external access. If using custom user or cluster admin user, skip this step.

```
clustershell:> security login unlock -username vsadmin -vserver svmname
```

- Ensure the firewall-policy of the data LIF is set to 'mgmt' (not 'data'). Skip this step if using a dedicated management lif to add the SVM.

```
clustershell:> network interface modify -lif <SVM_data_LIF_name> -firewall-policy mgmt
```

- When a firewall is enabled, you must have an exception defined to allow TCP traffic for the port using the Data ONTAP Data Collector.

See [Agent requirements](#) for configuration information. This applies to on-premise Agents and Agents installed in the Cloud.

- When an Agent is installed in an AWS EC2 instance to monitor a Cloud ONTAP SVM, the Agent and Storage must be in the same VPC. If they are in separate VPCs, there must be a valid route between the VPC's.

Test Connectivity for Data Collectors

The test connectivity feature (introduced March 2025) aims to help end users identify the specific causes of failures when setting up data collectors in Data Infrastructure Insights (DII) Workload Security. This allows the users to self-correct issues related to network communication or missing roles.

This feature will help users determine if all network-related checks are in place before setting up a data collector. Additionally, it will inform users about the features they can access based on the ONTAP version, roles, and permissions assigned to them in ONTAP.



Test connectivity is not supported for User Directory collectors

Prerequisites for Connection Testing

- Cluster level credentials are needed for this feature to work in full.
- Feature access check is not supported in SVM mode.
- If you are using cluster administration credentials, no new permissions are needed.
- If you are using a custom user (e.g., `csuser`), provide the mandatory permissions and feature specific permissions for the features you want to use.



Be sure to review the [Permissions](#) section below as well.

Test the Connection

The user can go to the add/edit collector page, enter the cluster level details (in Cluster Mode) or SVM level details (in SVM Mode), and click on the **Test Connection** button. Workload Security will then process the request and display an appropriate success or failure message.

Add ONTAP SVM

[Need Help?](#)

An Agent is required to fetch data from the ONTAP SVM in to Storage Workload Security

Network Checks:

Https: Connection successful on port 443 (AGENT -> ONTAP)

Ontap Version: 9.14.1

Data Lifs: Found 1 (10.██████) data interfaces in the SVM which contains service name data-fpolicy-client, admin/oper status as up.

Agent IP: Determined agent IP address to be used (10.██████)

Policy Server: Connection successful on Agent IP (10.██████), ports [35037, 35038, 35039] (ONTAP -> AGENT)

Features (User has permissions):

Snapshot, Ems, Access Denied, Persistent Store, Ontap ARP, User Blocking

Features (User does not have permissions):

Protobuf: Ontap version 9.14.1 is below minimum supported version 9.15.0

Things to note for ONTAP Multi Admin Verify (MAV)

Some features, such as the creation and deletion of snapshots or user blocking (SMB), may not work based on the MAV commands added in your version of ONTAP.

Follow the steps below to add exclusions to your MAV commands which allow Workload Security to create or delete snapshots and block users.

Commands to allow snapshot create and delete:

```
multi-admin-verify rule modify -operation "volume snapshot create" -query  
"-snapshot !*cloudsecure_*"  
multi-admin-verify rule modify -operation "volume snapshot delete" -query  
"-snapshot !*cloudsecure_*"
```

Command to allow user blocking:

```
multi-admin-verify rule delete -operation set
```

Prerequisites for User Access Blocking

Keep the following in mind for [User Access Blocking](#):

Cluster level credentials are needed for this feature to work.

If you are using cluster administration credentials, no new permissions are needed.

If you are using a custom user (for example, *csuser*) with permissions given to the user, then follow the steps in [User Access Blocking](#) to give permissions to Workload Security to block user.

A Note About Permissions

Permissions when adding via Cluster Management IP:

If you cannot use the Cluster management administrator user to allow Workload Security to access the ONTAP SVM data collector, you can create a new user named “csuser” with the roles as shown in the commands below. Use the username “csuser” and password for “csuser” when configuring the Workload Security data collector to use Cluster Management IP.

Note: You can create a single role to use for all feature permissions for a custom user. If there is an existing user then first delete the existing user and role using these commands:

```
security login delete -user-or-group-name csuser -application *
security login role delete -role csrole -cmddirname *
security login rest-role delete -role csrestrole -api *
security login rest-role delete -role arwrole -api *
```

To create the new user, log in to ONTAP with the Cluster management Administrator username/password, and execute the following commands on the ONTAP server:

```
security login role create -role csrole -cmddirname DEFAULT -access
readonly
```

```
security login role create -role csrole -cmddirname "vserver fpolicy"
-access all
security login role create -role csrole -cmddirname "volume snapshot"
-access all -query "-snapshot cloudsecure_*"
security login role create -role csrole -cmddirname "event catalog"
-access all
security login role create -role csrole -cmddirname "event filter" -access
all
security login role create -role csrole -cmddirname "event notification
destination" -access all
security login role create -role csrole -cmddirname "event notification"
-access all
security login role create -role csrole -cmddirname "security certificate"
-access all
security login role create -role csrole -cmddirname "cluster application-
record" -access all
security login create -user-or-group-name csuser -application ontapi
-authmethod password -role csrole
security login create -user-or-group-name csuser -application ssh
-authmethod password -role csrole
security login create -user-or-group-name csuser -application http
-authmethod password -role csrole
```

Permissions when adding via Vserver Management IP:

If you cannot use the Cluster management administrator user to allow Workload Security to access the ONTAP SVM data collector, you can create a new user named “csuser” with the roles as shown in the commands below. Use the username “csuser” and password for “csuser” when configuring the Workload Security data collector to use Vserver Management IP.

Note: You can create a single role to use for all feature permissions for a custom user. If there is an existing user then first delete the existing user and role using these commands:

```
security login delete -user-or-group-name csuser -application * -vserver
<vservername>
security login role delete -role csrole -cmddirname * -vserver
<vservername>
security login rest-role delete -role csrestrole -api * -vserver
<vservername>
```

To create the new user, log in to ONTAP with the Cluster management Administrator username/password, and execute the following commands on the ONTAP server. For ease, copy these commands to a text editor and replace the <vservername> with your Vserver name before and executing these commands on ONTAP:

```
security login role create -vserver <vservername> -role csrole -cmddirname
DEFAULT -access none
```

```
security login role create -vserver <vservername> -role csrole -cmddirname
"network interface" -access readonly
security login role create -vserver <vservername> -role csrole -cmddirname
version -access readonly
security login role create -vserver <vservername> -role csrole -cmddirname
volume -access readonly
security login role create -vserver <vservername> -role csrole -cmddirname
vserver -access readonly
```

```
security login role create -vserver <vservername> -role csrole -cmddirname
"vserver fpolicy" -access all
security login role create -vserver <vservername> -role csrole -cmddirname
"volume snapshot" -access all
```

```
security login create -user-or-group-name csuser -application ontapi
-authmethod password -role csrole -vserver <vservername>
security login create -user-or-group-name csuser -application http
-authmethod password -role csrole -vserver <vservername>
```

Protobuf Mode

Workload Security will configure the FPolicy engine in protobuf mode when this option is enabled in the collector's *Advanced Configuration* settings. Protobuf mode is supported in ONTAP version 9.15 and later.

More details on this feature can be found in the [ONTAP documentation](#).

Specific permissions are required for protobuf (some or all of these may already exist):

Cluster mode:

```
security login role create -role csrole -cmddirname "vserver fpolicy"  
-access all
```

Vserver mode:

```
security login role create -vserver <vservername> -role csrole -cmddirname  
"vserver fpolicy" -access all
```

Permissions for ONTAP Autonomous Ransomware Protection and ONTAP Access Denied

If you are using cluster administration credentials, no new permissions are needed.

If you are using a custom user (for example, *csuser*) with permissions given to the user, then follow the steps below to give permissions to Workload Security to collect ARP related information from ONTAP.

For more information, read about [Integration with ONTAP Access Denied](#)

and [Integration with ONTAP Autonomous Ransomware Protection](#)

Configure the data collector

Steps for Configuration

1. Log in as Administrator or Account Owner to your Data Infrastructure Insights environment.
2. Click **Workload Security > Collectors > +Data Collectors**

The system displays the available Data Collectors.

3. Hover over the **NetApp SVM tile** and click ***+Monitor**.

The system displays the ONTAP SVM configuration page. Enter the required data for each field.

Configuration

Field	Description
Name	Unique name for the Data Collector
Agent	Select a configured agent from the list.
Connect via Management IP for:	Select either Cluster IP or SVM Management IP

Cluster / SVM Management IP Address	The IP address for the cluster or the SVM, depending on your selection above.
SVM Name	The Name of the SVM (this field is required when connecting via Cluster IP)
Username	User name to access the SVM/Cluster When adding via Cluster IP the options are: 1. Cluster-admin 2. 'csuser' 3. AD-user having similar role as csuser. When adding via SVM IP the options are: 4. vsadmin 5. 'csuser' 6. AD-username having similar role as csuser.
Password	Password for the above user name
Filter Shares/Volumes	Choose whether to include or exclude Shares / Volumes from event collection
Enter complete share names to exclude/include	Comma-separated list of shares to exclude or include (as appropriate) from event collection
Enter complete volume names to exclude/include	Comma-separated list of volumes to exclude or include (as appropriate) from event collection
Monitor Folder Access	When checked, enables events for folder access monitoring. Note that folder create/ rename and delete will be monitored even without this option selected. Enabling this will increase the number of events monitored.
Set ONTAP Send Buffer size	Sets the ONTAP Fpolicy send buffer size. If an ONTAP version prior to 9.8p7 is used and performance issue is seen, then the ONTAP send buffer size can be altered to get improved ONTAP performance. Contact NetApp Support if you do not see this option and wish to explore it.

After you finish

- In the Installed Data Collectors page, use the options menu on the right of each collector to edit the data collector. You can restart the data collector or edit data collector configuration attributes.

Recommended Configuration for MetroCluster

The following is recommended for MetroCluster:

1. Connect two data collectors, one to the source SVM and another to the destination SVM.
2. The data collectors should be connected by *Cluster IP*.
3. At any moment in time, the current 'running' SVM's data collector will show as *Running*. The current 'stopped' SVM's data collector will show as *Stopped*.
4. Whenever there is a switchover, the state of the data collector will change from *Running* to *Stopped* and vice versa.
5. It will take up to two minutes for the data collector to move from *Stopped* state to *Running* state.

Service Policy

If using service policy with ONTAP **version 9.9.1 or newer**, in order to connect to the Data Source Collector, the *data-fpolicy-client* service is required along with the data service *data-nfs*, and/or *data-cifs*.

Example:

```
Testcluster-1:*> net int service-policy create -policy only_data_fpolicy
-allowed-addresses 0.0.0.0/0 -vserver aniket_svm
-services data-cifs,data-nfs,data,-core,data-fpolicy-client
(network interface service-policy create)
```

In versions of ONTAP prior to 9.9.1, *data-fpolicy-client* need not be set.

Play-Pause Data Collector

If the Data Collector is in *Running* state, you can Pause collection. Open the "three dots" menu for the collector and select PAUSE. While the collector is paused, no data is gathered from ONTAP, and no data is sent from the collector to ONTAP. This means no Fpolicy events will flow from ONTAP to the data collector, and from there to Data Infrastructure Insights.

Note that if any new volumes, etc. are created on ONTAP while the collector is Paused, Workload Security won't gather the data and those volumes, etc. will not be reflected in dashboards or tables.



A collector cannot be paused if it has restricted users. Restore the user access before pausing the collector.

Keep the following in mind:

- Snapshot purge won't happen as per the settings configured on a paused collector.
- EMS events (like ONTAP ARP) won't be processed on a paused collector. This means if ONTAP identifies a file tampering attack, Data Infrastructure Insights Workload Security won't be able to acquire that event.
- Health notifications emails will NOT be sent for a paused collector.
- Manual or Automatic actions (such as Snapshot or User Blocking) will not be supported on a paused collector.
- On agent or collector upgrades, agent VM restarts/reboots, or agent service restart, a paused collector will remain in *Paused* state.
- If the data collector is in *Error* state, the collector cannot be changed to *Paused* state. The Pause button will be enabled only if the state of the collector is *Running*.
- If the agent is disconnected, the collector cannot be changed to *Paused* state. The collector will go into *Stopped* state and the Pause button will be disabled.

Persistent Store

Persistent store is supported with ONTAP 9.14.1 and later. Note that volume name instructions vary from ONTAP 9.14 to 9.15.

Persistent Store can be enabled by selecting the checkbox in the collector edit/add page. After selecting the checkbox, a text field is displayed for accepting volume name. Volume name is a mandatory field for enabling

Persistent Store.

- For ONTAP 9.14.1, you must create the volume prior to enabling the feature, and provide the same name in the *Volume Name* field. The recommended volume size is 16GB.
- For ONTAP 9.15.1, the volume will be created automatically with 16GB size by the collector, using the name provided in the *Volume Name* field.

Specific permissions are required for Persistent Store (some or all of these may already exist):

Cluster mode:

```
security login role create -role csrole -cmddirname "vserver fpolicy"  
-access all  
security login role create -role csrole -cmddirname "job show" -access  
readonly
```

Vserver mode:

```
security login role create -vserver <vservername> -role csrole -cmddirname  
"vserver fpolicy" -access all  
security login role create -vserver <vservername> -role csrole -cmddirname  
"job show" -access readonly
```

Migrate Collectors

You can easily migrate a Workload Security collector from one agent to another, allowing for efficient load balancing of collectors across agents.

Prerequisites

- Source agent must be in *connected* state.
- Collector to be migrated must be in *running* state.

Note:

- Migrate is supported for both Data and User Directory collectors.
- Migration of a collector is not supported for manually managed tenants.

Migrate collector

To migrate a collector, follow these steps:

1. Go to the "Edit Collector" page.
2. Select a destination agent from the agent dropdown.
3. Click on the "Save Collector" button.

Workload Security will process the request. On successful migration, the user will be redirected to the collectors list page. In case of failure, an appropriate message will be displayed on the edit page.

Note: Any configuration changes previously made on the "Edit Collector" page will be remain applied when the collector is successfully migrated to the destination agent.

Edit ONTAP SVM

Name*	Agent
<input type="text" value="CI_SVM"/>	<input type="text" value="fp-cs-1-agent (CONNECTED)"/>
Connect via Management IP for:	
<input checked="" type="radio"/> Cluster	agent-1537 (CONNECTED)
<input type="radio"/> SVM	agent-jptsc (CONNECTED)
	fp-cs-1-agent (CONNECTED)
	fp-cs-2-agent (CONNECTED)
	GSSC_girton (CONNECTED)

Troubleshooting

See the [Troubleshooting the SVM Collector](#) page for troubleshooting tips.

Troubleshooting the ONTAP SVM Data Collector

Workload Security uses data collectors to collect file and user access data from devices. Here you can find tips for troubleshooting issues with this collector.

See the [Configuring the SVM Collector](#) page for instructions on configuring this collector.

In the case of an error, you can click on *more detail* in the *Status* column of the Installed Data Collectors page for detail about the error.

Installed Data Collectors

Name	Status	Type	Agent
9.8_vs1	! Error more detail	ONTAP SVM	agent-11

Known problems and their resolutions are described below.

Problem: Data Collector runs for some time and stops after a random time, failing with: "Error message: Connector is in error state. Service name: audit. Reason for failure: External fpolicy server overloaded."

Try This:

The event rate from ONTAP was much higher than what the Agent box can handle. Hence the connection got terminated.

Check the peak traffic in CloudSecure when the disconnection happened. This you can check from the [CloudSecure > Activity Forensics > All Activity](#) page.

If the peak aggregated traffic is higher than what the Agent Box can handle, then please refer to the Event

Rate Checker page on how to size for Collector deployment in an Agent Box.

If the Agent was installed in the Agent box prior to 4 March 2021, run the following commands in the Agent box:

```
echo 'net.core.rmem_max=8388608' >> /etc/sysctl.conf
echo 'net.ipv4.tcp_rmem = 4096 2097152 8388608' >> /etc/sysctl.conf
sysctl -p
```

Restart the collector from the UI after resizing.

Problem: Collector reports Error Message: “No local IP address found on the connector that can reach the data interfaces of the SVM”.

Try This: This is most likely due to a networking issue on the ONTAP side. Please follow these steps:

1. Ensure that there are no firewalls on the SVM data lif or the management lif which are blocking the connection from the SVM.
2. When adding an SVM via a cluster management IP, please ensure that the data lif and management lif of the SVM are pingable from the Agent VM. In case of issues, check the gateway, netmask and routes for the lif.

You can also try logging in to the cluster via ssh using the cluster management IP, and ping the Agent IP. Make sure that the agent IP is pingable:

```
network ping -vserver <vserver name> -destination <Agent IP> -lif <Lif Name> -show-detail
```

If not pingable, make sure the network settings in ONTAP are correct, so that the Agent machine is pingable.

3. If you have tried connecting via Cluster IP and it is not working, try connecting directly via SVM IP. Please see above for the steps to connect via SVM IP.
4. While adding the collector via SVM IP and vsadmin credentials, check if the SVM Lif has Data plus Mgmt role enabled. In this case ping to the SVM Lif will work, however SSH to the SVM Lif will not work. If yes, create an SVM Mgmt Only Lif and try connecting via this SVM management only Lif.
5. If it is still not working, create a new SVM Lif and try connecting through that Lif. Make sure that the subnet mask is correctly set.
6. Advanced Debugging:
 - a. Start a packet trace in ONTAP.
 - b. Try to connect a data collector to the SVM from CloudSecure UI.
 - c. Wait till the error appears. Stop the packet trace in ONTAP.
 - d. Open the packet trace from ONTAP. It is available at this location

```
https://<cluster_mgmt_ip>/spi/<clustername>/etc/log/packet_traces/
```

- e. Make sure there is a SYN from ONTAP to the Agent box.
- f. If there is no SYN from ONTAP then it is an issue with firewall in ONTAP.
- g. Open the firewall in ONTAP, so that ONTAP is able to connect the agent box.

7. If it is still not working, please consult the networking team to make sure that no external firewall is blocking the connection from ONTAP to the Agent box.

8. If none of the above solves the issue, open a case with [Netapp Support](#) for further assistance.

Problem: Message: "Failed to determine ONTAP type for [hostname: <IP Address>. Reason: Connection error to Storage System <IP Address>: Host is unreachable (Host unreachable)"

Try this:

1. Verify that the correct SVM IP Management address or Cluster Management IP has been provided.
2. SSH to the SVM or the Cluster to which you are intending to connect. Once you are connected ensure that the SVM or the Cluster name is correct.

Problem: Error Message: "Connector is in error state. Service.name: audit. Reason for failure: External fpolicy server terminated."

Try this:

1. It is most likely that a firewall is blocking the necessary ports in the agent machine. Verify the port range 35000-55000/tcp is opened for the agent machine to connect from the SVM. Also ensure that there are no firewalls enabled from the ONTAP side blocking communication to the agent machine.
2. Type the following command in the Agent box and ensure that the port range is open.

```
sudo iptables-save | grep 3500*
```

Sample output should look like:

```
-A IN_public_allow -p tcp -m tcp --dport 35000 -m conntrack -ctstate NEW  
-j ACCEPT
```

3. Login to SVM, enter the following commands and check that no firewall is set to block the communication with ONTAP.

```
system services firewall show  
system services firewall policy show
```

[Check firewall commands](#) on the ONTAP side.

4. SSH to the SVM/Cluster which you want to monitor. Ping the Agent box from the SVM data lif (with CIFS, NFS protocols support) and ensure that ping is working:

```
network ping -vserver <vserver name> -destination <Agent IP> -lif <Lif Name> -show-detail
```

If not pingable, make sure the network settings in ONTAP are correct, so that the Agent machine is pingable.

5. If a single SVM is added twice added to a tenant via 2 data collectors, then this error will be shown. Delete one of the data collectors thru the UI. Then restart the other data collector thru the UI. Then the data collector will show "RUNNING" status and will start receiving events from SVM.

Basically, in a tenant, 1 SVM should be added only once, via 1 data collector. 1 SVM should not added twice via 2 data collectors.

6. In instances where the same SVM was added in two different Workload Security environments (tenants), the last one will always succeed. The second collector will configure fpolicy with its own IP address and kick out the first one. So the collector in the first one will stop receiving events and its "audit" service will enter into error state.

To prevent this, configure each SVM on a single environment.

7. This error may also occur if service policies are not configured correctly. With ONTAP 9.8 or later, in order to connect to the Data Source Collector, the data-fpolicy-client service is required along with the data service data-nfs, and/or data-cifs. Additionally, the data-fpolicy-client service must be associated with the data lif(s) for the monitored SVM.

Problem: No events seen in activity page.

Try this:

1. Check if ONTAP collector is in "RUNNING" state. If yes, then ensure that some cifs events are being generated on the cifs client VMs by opening some files.
2. If no activities are seen, please login to the SVM and enter the following command.

```
<SVM>event log show -source fpolicy
```

Please ensure that there are no errors related to fpolicy.

3. If no activities are seen, please login to the SVM. Enter the following command:

```
<SVM>fpolicy show
```

Check if the fpolicy policy named with prefix "cloudsecure_" has been set and status is "on". If not set, then most likely the Agent is unable to execute the commands in the SVM. Please ensure all the prerequisites

as described in the beginning of the page have been followed.

Problem: SVM Data Collector is in error state and Error message is "Agent failed to connect to the collector"
Try this:

1. Most likely the Agent is overloaded and is unable to connect to the Data Source collectors.
2. Check how many Data Source collectors are connected to the Agent.
3. Also check the data flow rate in the "All Activity" page in the UI.
4. If the number of activities per second is significantly high, install another Agent and move some of the Data Source Collectors to the new Agent.

Problem: SVM Data Collector shows error message as "fpolicy.server.connectError: Node failed to establish a connection with the FPolicy server "12.195.15.146" (reason: "Select Timed out")"

Try this: Firewall is enabled in SVM/Cluster. So fpolicy engine is unable to connect to fpolicy server. CLIs in ONTAP which can be used to get more information are:

```
event log show -source fpolicy which shows the error
event log show -source fpolicy -fields event,action,description which
shows more details.
```

[Check firewall commands on the ONTAP side.](#)

Problem: Error Message: "Connector is in error state. Service name:audit. Reason for failure: No valid data interface (role: data,data protocols: NFS or CIFS or both, status: up) found on the SVM."

Try this: Ensure there is an operational interface (having role as data and data protocol as CIFS/NFS).

Problem: The data collector goes into Error state and then goes into RUNNING state after some time, then back to Error again. This cycle repeats.

Try this: This typically happens in the following scenario:

1. There are multiple data collectors added.
2. The data collectors which show this kind of behavior will have 1 SVM added to these data collectors. Meaning 2 or more data collectors are connected to 1 SVM.
3. Ensure 1 data collector connects to only 1 SVM.
4. Delete the other data collectors which are connected to the same SVM.

Problem: Connector is in error state. Service name: audit. Reason for failure: Failed to configure (policy on SVM svmname. Reason: Invalid value specified for 'shares-to-include' element within 'fpolicy.policy.scope-modify: "Federal"

Try this: *The share names need to be given without any quotes. Edit the ONTAP SVM DSC configuration to correct the share names.

Include and exclude shares is not intended for a long list of share names. Use filtering by volume instead if you have a large number of shares to include or exclude.

Problem: There are existing fpolicies in the Cluster which are unused. What should be done with those prior to installation of Workload Security?

Try this: It is recommended to delete all existing unused fpolicy settings even if they are in disconnected state. Workload Security will create fpolicy with the prefix "cloudsecure_". All other unused fpolicy configurations can be deleted.

CLI command to show fpolicy list:

```
fpolicy show
```

Steps to delete fpolicy configurations:

```
fpolicy disable -vserver <svmname> -policy-name <policy_name>
fpolicy policy scope delete -vserver <svmname> -policy-name <policy_name>
fpolicy policy delete -vserver <svmname> -policy-name <policy_name>
fpolicy policy event delete -vserver <svmname> -event-name <event_list>
fpolicy policy external-engine delete -vserver <svmname> -engine-name
<engine_name>
```

Problem: After enabling Workload Security, ONTAP performance is impacted: Latency becomes sporadically high, IOPs become sporadically low.

Try this: While using ONTAP with Workload Security sometimes latency issues can be seen in ONTAP. There are a number of possible reasons for this as noted in the following: [1372994](#), [1415152](#), [1438207](#), [1479704](#), [1354659](#). All of these issues are fixed in ONTAP 9.13.1 and later; it is strongly recommended to use one of these later versions.

Problem: Data Collector shows the error message:

"Error: Failed to determine the health of the collector within 2 retries, try restarting the collector again (Error Code: AGENT008)".

Try this:

1. On the Data Collectors page, scroll to the right of the data collector giving the error and click on the 3 dots

menu. Select *Edit*.

Enter the password of the data collector again.

Save the data collector by pressing on the *Save* button.

Data Collector will restart and the error should be resolved.

2. The Agent machine may not have enough CPU or RAM headroom, that is why the DSCs are failing.

Please check the number of Data Collectors which are added to the Agent in the machine.

If it is more than 20, please increase the CPU and RAM capacity of the Agent machine.

Once the CPU and RAM is increased, the DSCs will get into Initializing and then to Running state automatically.

Look into the sizing guide on [this page](#).

Problem: The Data Collector is erroring out when SVM mode is selected.

Try this: While connecting in SVM mode, If cluster management IP is used to connect instead of SVM management IP, then the connection will error out. Make sure that the correct SVM IP is used.

Problem: Data collector shows an error message when Access Denied feature is enabled:

"Connector is in error state. Service name: audit. Reason for failure: Failed to configure fpolicy on SVM test_svm. Reason: User is not authorized."

Try this: The user might be missing the REST permissions needed for the Access Denied feature. Please follow the instructions on [this page](#) to set the permissions.

Restart the collector once the permissions are set.

Problem: Collector is in Error state with the message:

Connector is in error state. Service name: audit. Reason for failure: Failed to configure persistent store on SVM <SVM Name>. Reason: Unable to find a suitable aggregate for volume "<volumeName>" in SVM "<SVM Name>". Reason: Performance information for aggregate "<aggregateName>" is currently not available. Wait a few minutes and try the command again.

Try this: Wait a few minutes and then restart the Collector.

If you are still experiencing problems, reach out to the support links mentioned in the **Help > Support** page.

Configuring the Cloud Volumes ONTAP and Amazon FSx for NetApp ONTAP collector

Monitor file and user access across your cloud storage infrastructure by configuring Workload Security data collectors for Cloud Volumes ONTAP and Amazon FSx for NetApp ONTAP. This guide provides step-by-step instructions for deploying Agents in AWS and connecting them to your cloud storage instances.

Cloud Volumes ONTAP Storage Configuration

See the OnCommand Cloud Volumes ONTAP Documentation to configure a single-node / HA AWS instance to host the Workload Security Agent:

<https://docs.netapp.com/us-en/cloud-manager-cloud-volumes-ontap/index.html>

After the configuration is complete, follow the steps to setup your SVM:

https://docs.netapp.com/us-en/cloudinsights/task_add_collector_svm.html

Supported Platforms

- Cloud Volumes ONTAP, supported in all the available cloud service providers wherever available. For example: Amazon, Azure, Google Cloud.
- ONTAP Amazon FSx

Agent Machine Configuration

The agent machine must be configured in the respective subnets of the cloud Service providers. Read more about network access in the [Agent Requirements].

Below are the steps for Agent installation in AWS. Equivalent steps, as applicable to the cloud service provider, can be followed in Azure or Google Cloud for the installation.

In AWS, use the following steps to configure the machine to be used as a Workload Security Agent:

Use the following steps to configure the machine to be used as a Workload Security Agent:

Steps

1. Log in to the AWS console and navigate to EC2-Instances page and select *Launch instance*.
2. Select a RHEL or CentOS AMI with the appropriate version as mentioned in this page:
https://docs.netapp.com/us-en/cloudinsights/concept_cs_agent_requirements.html
3. Select the VPC and Subnet that the Cloud ONTAP instance resides in.
4. Select *t2.xlarge* (4 vcpus and 16 GB RAM) as allocated resources.
 - a. Create the EC2 instance.
5. Install the required Linux packages using the YUM package manager:
 - a. Install *wget* and *unzip* native Linux packages.

Install the Workload Security Agent

1. Log in as Administrator or Account Owner to your Data Infrastructure Insights environment.
2. Navigate to Workload Security **Collectors** and click the **Agents** tab.
3. Click **+Agent** and specify RHEL as the target platform.
4. Copy the Agent Installation command.
5. Paste the Agent Installation command into the RHEL EC2 instance you are logged in to.
This installs the Workload Security agent, providing all of the [Agent Prerequisites](#) are met.

For detailed steps please refer to this xref:./

https://docs.netapp.com/us-en/cloudinsights/task_cs_add_agent.html#steps-to-install-agent

Troubleshooting

Known problems and their resolutions are described in the following table.

Problem	Resolution
<p>“Workload Security: Failed to determine ONTAP type for Amazon FxSN data collector” error is shown by the Data Collector.</p> <p>Customer is unable to add new Amazon FSxN data collector into Workload Security. Connection to FSxN cluster on port 443 from the agent is timing out. Firewall and AWS security groups have the required rules enabled to allow communication. An agent is already deployed and is in the same AWS account as well. This same agent is used to connect and monitor the remaining NetApp devices (and all of them are working).</p>	<p>Solve this issue by adding fsxadmin LIF network segment to agent’s security rule.</p> <p>Allowed all ports if you are not sure about the ports.</p>

User Management

Workload Security user accounts are managed through Data Infrastructure Insights.

Data Infrastructure Insights provides four user account levels: Account Owner, Administrator, User, and Guest. Each account is assigned specific permission levels. A User account that has Administrator privileges can create or modify users, and assign each user one of the following Workload Security roles:

Role	Workload Security Access
Administrator	Can perform all Workload Security functions, including those for Alerts, Forensics, data collectors, automated response policies, and APIs for Workload Security. An Administrator can also invite other users but can only assign Workload Security roles.
User	Can view and manage Alerts and view Forensics. User role can change alert status, add a note, take snapshots manually, and restrict user access.
Guest	Can view Alerts and Forensics. Guest role cannot change alert status, add a note, take snapshots manually, or restrict user access.

Steps

1. Log into Workload Security
2. In the menu, click **Admin > User Management**

You will be forwarded to Data Infrastructure Insights’s User Management page.

3. Select the desired role for each user.

While adding a new user, simply select the desired role (usually User or Guest).

More information on User accounts and roles can be found in the Data Infrastructure Insights [User Role](#)

documentation.

Event Rate Checker: Agent Sizing Guide

Determine optimal Agent machine sizing by measuring NFS and SMB event rates generated by your SVMs before deploying data collectors. The Event Rate Checker script helps you understand capacity limits (maximum 50 data collectors per Agent) and ensures your Agent infrastructure can handle your expected event volume for reliable threat detection.

Requirements:

- Cluster IP
- Cluster admin username and password



When running this script no ONTAP SVM Data Collector should be running for the SVM for which event rate is being determined.

Steps:

1. Install the Agent by following the instructions in CloudSecure.
2. Once the agent is installed, run the `server_data_rate_checker.sh` script as a sudo user:

```
/opt/netapp/cloudsecure/agent/install/svm_event_rate_checker.sh
```

3. This script requires `sshpass` to be installed in the linux machine. There are two ways to install it:
 - a. Run the following command:

```
linux_prompt> yum install sshpass
```

- b. If that does not work, then download `sshpass` to the linux machine from the web and run the following command:

```
linux_prompt> rpm -i sshpass
```

4. Provide the correct values when prompted. See below for an example.
5. The script will take approximately 5 minutes to run.
6. After the run is complete, the script will print the event rate from the SVM. You can check Event rate per SVM in the console output:

```
"Svm svm_rate is generating 100 events/sec".
```

Each Ontap SVM Data Collector can be associated with a single SVM, which means each data collector will be

able to receive the number of events which a single SVM generates.

Keep the following in mind:

A) Use this table as a general sizing guide. You can increase the number of cores and/or memory to increase the number of data collectors supported, up to a maximum of 50 data collectors:

Agent Machine Configuration	Number of SVM Data Collectors	Max event Rate which the Agent Machine can handle
4 core, 16GB	10 data collectors	20K events/sec
4 core, 32GB	20 data collectors	20K events/sec

B) To calculate your total events, add the Events generated for all SVMs for that agent.

C) If the script is not run during peak hours or if peak traffic is difficult to predict, then keep an event rate buffer of 30%.

B + C Should be less than A, otherwise the Agent machine will fail to monitor.

In other words, the number of data collectors which can be added to a single agent machine should comply to the formula below:

Sum of all Event rate of all Data Source Collectors + Buffer Event rate of 30% < 20000 events/second

See the [Agent Requirements](#) page for additional pre-requisites and requirements.

Example

Let us say we have three SVMs generating event rates of 100, 200, and 300 events per second, respectively.

We apply the formula:

$(100+200+300) + [(100+200+300) * 30\%] = 600 + 180 = 780 \text{ events/sec}$
780 events/second is < 20000 events/second, so the 3 SVMs can be monitored via one agent box.

Console output is available in the Agent machine in the file name *fpolicy_stat_<SVM Name>.log* in the present working directory.

The script may give erroneous results in the following cases:

- Incorrect credentials, IP, or SVM name are provided.
- An already-existing fpolicy with same name, sequence number, etc. will give error.
- The script is stopped abruptly while running.

An example script run is shown below:

```
[root@ci-cs-data agent]#  
/opt/netapp/cloudsecure/agent/install/svm_event_rate_checker.sh
```

```
Enter the cluster ip: 10.192.139.166  
Enter the username to SSH: admin  
Enter the password:  
Getting event rate for NFS and SMB events.  
Available SVMs in the Cluster  
-----  
QA_SVM  
Stage_SVM  
Qa-fas8020  
Qa-fas8020-01  
Qa-fas8020-02  
audit_svm  
svm_rate  
vs_new  
vs_new2
```

```
-----  
Enter [1/5] SVM name to check (press enter to skip): svm_rate  
Enter [2/5] SVM name to check (press enter to skip): audit_svm  
Enter [3/5] SVM name to check (press enter to skip):  
Enter [4/5] SVM name to check (press enter to skip):  
Enter [5/5] SVM name to check (press enter to skip):  
Running check for svm svm_rate...  
Running check for svm audit_svm...  
Waiting 5 minutes for stat collection  
Stopping sample svm_rate_sample  
Stopping sample audit_svm_sample  
fpolicy stats of svm svm_rate is saved in fpolicy_stat_svm_rate.log  
Svm svm_rate is generating 100 SMB events/sec and 100 NFS events/sec  
Overall svm svm_rate is generating 200 events/sec  
fpolicy stats of svm audit_svm is saved in fpolicy_stat_audit_svm.log  
Svm audit_svm is generating 200 SMB events/sec and 100 NFS events/sec  
Overall svm audit_svm is generating 300 events/sec
```

```
[root@ci-cs-data agent]#
```

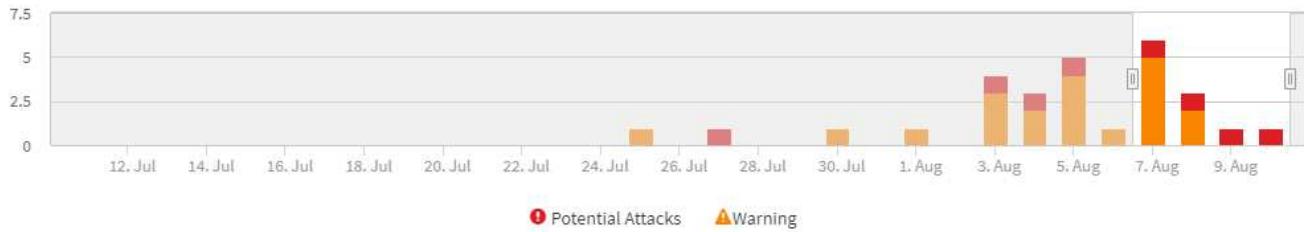
Troubleshooting

Question	Answer
If I run this script on an SVM that is already configured for Workload Security, does it just use the existing fpolicy config on the SVM or does it setup a temporary one and run the process?	The Event Rate Checker can run fine even for an SVM already configured for Workload Security. There should be no impact.
Can I increase the number of SVMs on which the script can be run?	Yes. Simply edit the script and change the max number of SVMs from 5 to any desirable number.
If I increase the number of SVMs, will it increase the time of running of the script?	No. The script will run for a max of 5 minutes, even if the number of SVMs is increased.
Can I increase the number of SVMs on which the script can be run?	Yes. You need to edit the script and change the max number of SVMs from 5 to any desirable number.
If I increase the number of SVMs, will it increase the time of running of the script?	No. The script will run for a max of 5mins, even if the number of SVMs are increased.
What happens if I run the Event Rate Checker with an existing agent?	Running the Event Rate Checker against an already-existing agent may cause an increase in latency on the SVM. This increase will be temporary in nature while the Event rate Checker is running.

Understanding and Investigating Alerts

The Workload Security Alerts page provides a comprehensive timeline of detected threats and warnings with detailed investigation tools. View alert details, manage status updates, filter by criteria, and track user activities to efficiently investigate and respond to security incidents.

Filter By Status New ✖️ +



● Potential Attacks ⚠️ Warning

● Potential Attacks (3)

Potential Attacks	Detected ↓	Status	User	Evidence	Action Taken
Ransomware Attack	5 hours ago Aug 10, 2020 4:38 AM	New		Iris McIntosh	> 700 Files Encrypted
Ransomware Attack	a day ago Aug 9, 2020 3:51 AM	New		Christy Santos	> 500 Files Encrypted
Ransomware Attack	2 days ago Aug 8, 2020 4:29 AM	New		Safwan Langley	> 700 Files Encrypted

⚠️ Warnings (7)

Abnormal Behaviour	Detected ↓	Status	User	Change	Action Taken	
User Activity Rate	2 days ago Aug 8, 2020 7:49 PM	New		Iris McIntosh	↑ 192.46%	None
User Activity Rate	2 days ago Aug 8, 2020 7:32 PM	New		Jenny Bryan	↑ 73.64%	None
User Activity Rate	3 days ago Aug 7, 2020 8:07 PM	New		Szymon Owen	↑ 189.88%	None

Alert

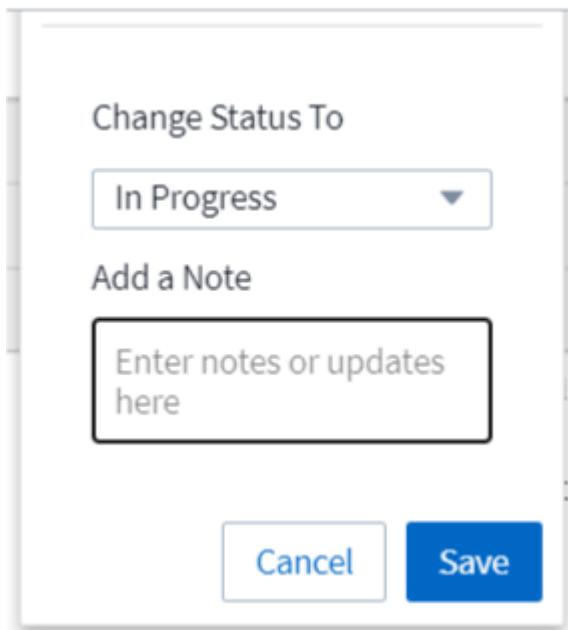
The Alert list displays a graph showing the total number of Potential Attacks and/or Warnings that have been raised in the selected time range, followed by a list of the attacks and/or warnings that occurred in that time range. You can change the time range by adjusting the start time and end time sliders in the graph.

The following is displayed for each alert:

Potential Attacks:

- The *Potential Attack* type (for example, File Tampering or Sabotage)
- The date and time the potential attack was *Detected*
- The *Status* of the alert:
 - **New:** This is the default for new alerts.
 - **In Progress:** The alert is under investigation by a team member or members.
 - **Resolved:** The alert has been marked as resolved by a team member.
 - **Dismissed:** The alert has been dismissed as false positive or expected behavior.

An administrator can change the status of the alert and add a note to assist with investigation.



The form contains the following fields and buttons:

- Change Status To:** A dropdown menu currently set to "In Progress".
- Add a Note:** A text area with the placeholder "Enter notes or updates here".
- Buttons:** "Cancel" (light blue background) and "Save" (blue background).

- The *User* whose behavior triggered the alert
- *Evidence* of the attack (for example, a large number of files was encrypted)
- The *Action Taken* (for example, a snapshot was taken)

Warnings:

- The *Abnormal Behavior* that triggered the warning
- The date and time the behavior was *Detected*
- The *Status* of the alert (New, In progress, etc.)
- The *User* whose behavior triggered the alert
- A description of the *Change* (for example, an abnormal increase in file access)
- The *Action Taken*

Filter Options

You can filter Alerts by the following:

- The *Status* of the alert
- Specific text in the *Note*
- The type of *Attacks/Warnings*
- The *User* whose actions triggered the alert/warning

The Alert Details page

You can click an alert link on the Alerts list page to open a detail page for the alert. Alert details may vary according to the type of attack or alert. For example, a file tampering attack detail page may show the following information:

Summary section:

- Attack type (File Tampering, Sabotage) and Alert ID (assigned by Workload Security)
- Date and Time the attack was detected
- Action Taken (for example, an automatic snapshot was taken. Time of snapshot is shown immediately below the summary section))
- Status (New, In Progress, etc.)

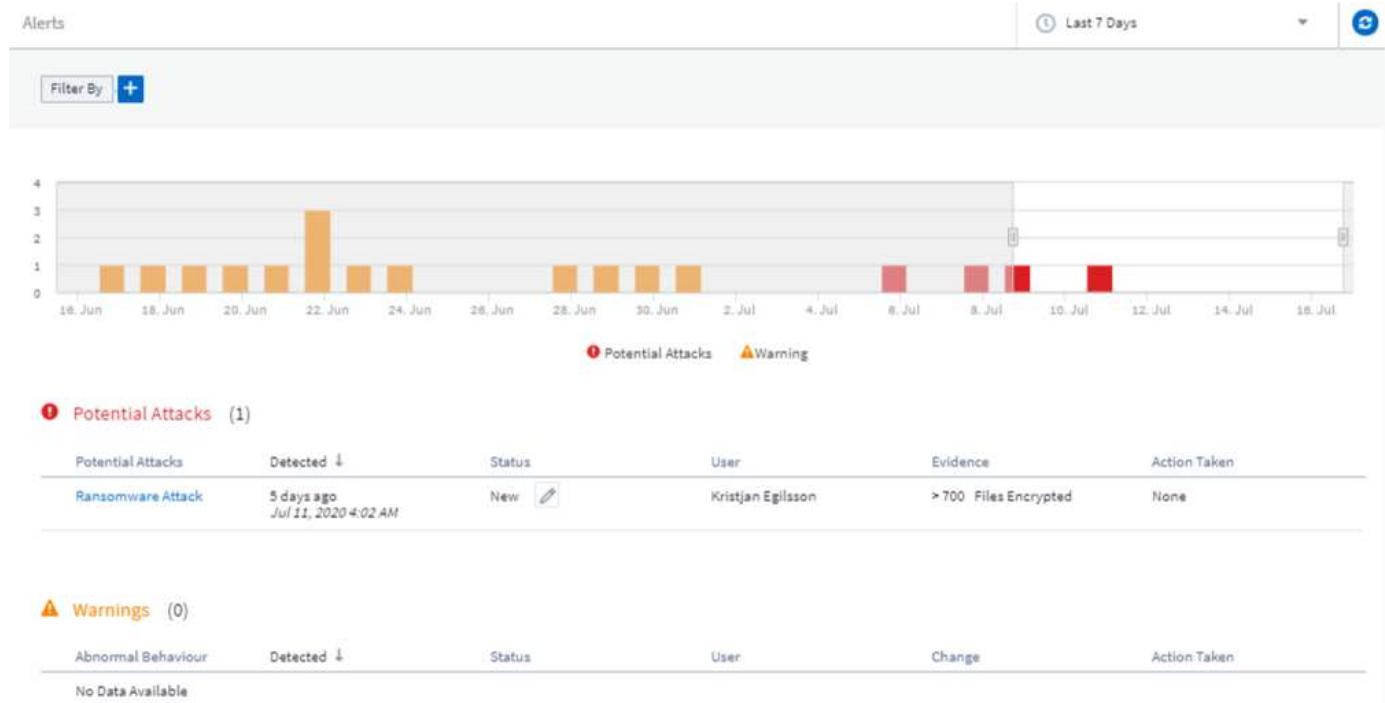
Attack Results section:

- Counts of Affected Volumes and Files
- An accompanying summary of the detection
- A graph showing file activity during the attack

Related Users section:

This section shows details about the user involved in the potential attack, including a graph of Top Activity for the user.

Alerts page (this example shows a potential file tampering attack):



Detail page (this example shows a potential file tampering attack):



POTENTIAL ATTACK: AL_305
Ransomware Attack

Detected
5 days ago
Jul 11, 2020 4:02 AM

Action Taken
None

Status
New

Total Attack Results

1 Affected Volumes | 0 Deleted Files | 4173 Encrypted Files

4173 Files have been copied, deleted, and potentially encrypted by 1 user account.

This is potentially a sign of ransomware attack.
The extension "crypt" was added to each file.

Encrypted Files

Activity per minute



Related Users



Kristjan Egilsson
Accountant
Finance

4173
Encrypted Files

Detected
5 days ago
Jul 11, 2020 4:02 AM

Action Taken
None



Username
us035

Email
Egilsson@netapp.com

Phone
387224312607

Department
Finance

Manager
Lyndsey Maddox

[View Activity Detail](#)

Top Activity Types
Activity per minute
Last access location: 10.197.144.115

Write Read Metadata Others



Take a Snapshot Action

Workload Security protects your data by automatically taking a snapshot when malicious activity is detected, ensuring that your data is safely backed up.

You can define [automated response policies](#) that take a snapshot when file tampering attack or other abnormal user activity is detected.

You can also take a snapshot manually from the alert page.

Automatic Snapshot taken:

POTENTIAL ATTACK: AL_307
Ransomware Attack

Detected
4 days ago
Jul 26, 2020 3:38 AM

Action Taken
Snapshots Taken

Status
In Progress

Last snapshots taken by
Amit Schwartz
Jul 30, 2020 2:54 PM

How To:
Restore Entities

Re-Take Snapshots

Total Attack Results

1	0	5148
Affected Volumes	Deleted Files	Encrypted Files

5148 Files have been copied, deleted, and potentially encrypted by 1 user account.

This is potentially a sign of ransomware attack.
The extension "crypt" was added to each file.

Encrypted Files
Activity per minute

0 100 200

03:00 03:30 04:00 04:30 05:00 05:30

Related Users

Ewen Hall
Developer
Engineering

5148
Encrypted Files

Detected
4 days ago
Jul 26, 2020 3:38 AM

Action Taken
Snapshots Taken

Manual Snapshot:

Cloud Insights

MONITOR & OPTIMIZE Alerts / *Nabilah Howell* had an abnormal change in activity rate

Jul 23, 2020 - Jul 26, 2020 1:44 AM 1:44 AM

CLOUD SECURE

ALERTS

FORENSICS

ADMIN

HELP

Alert Detail

WARNING: AL_306
Nabilah Howell had an abnormal change in activity rate.

Detected
5 days ago
Jul 25, 2020 1:44 PM

Action Taken
None

Status
New

Recommendation: Setup an Automated Response Policy
An Automated Response Policy will trigger measures to contain the damage automatically when a future attack is detected. Try it now.

Take Snapshots How To: Restore Entities

Nabilah Howell's Activity Rate Change

Typical	Alert
122.8 Activities Per Minute	210 Activities Per Minute

Activity Rate
Activity per 5 minutes

0 200 400 600

12:00 24. Jul 12:00 25. Jul 12:00 26. Jul

Alert Notifications

Email notifications of alerts are sent to an alert recipient list for every action on the alert. To configure alert recipients, click on **Admin > Notifications** and enter an email addresses for each recipient.

Retention Policy

Alerts and Warnings are retained for 13 months. Alerts and Warnings older than 13 months will be deleted. If the Workload Security environment is deleted, all data associated with the environment is also deleted.

Troubleshooting

Problem:	Try This:
There is a situation where, ONTAP takes hourly snapshots per day. Will Workload Security (WS) snapshots affect it? Will WS snapshot take the hourly snapshot place? Will the default hourly snapshot get stopped?	Workload Security snapshots will not affect the hourly snapshots. WS snapshots will not take the hourly snapshot space and that should continue as before. The default hourly snapshot will not get stopped.
What will happen if the maximum snapshot count is reached in ONTAP?	<p>If the maximum Snapshot count is reached, subsequent Snapshot taking will fail and Workload Security will show an error message noting that Snapshot is full.</p> <p>User needs to define Snapshot policies to delete the oldest snapshots, otherwise snapshots will not be taken.</p> <p>In ONTAP 9.3 and earlier, a volume can contain up to 255 Snapshot copies. In ONTAP 9.4 and later, a volume can contain up to 1023 Snapshot copies.</p> <p>See the ONTAP Documentation for information on setting Snapshot deletion policy.</p>
Workload Security is unable to take snapshots at all.	<p>Make sure that the role being used to create snapshots has proper rights assigned.</p> <p>Make sure <code>csrole</code> is created with proper access rights for taking snapshots:</p> <pre>security login role create -vserver <vservername> -role csrole -cmddirname "volume snapshot" -access all</pre>
Snapshots are failing for older alerts on SVMs which were removed from Workload Security and subsequently added back again. For new alerts which occur after SVM is added again, snapshots are taken.	This is a rare scenario. In the event you experience this, log in to ONTAP and take the snapshots manually for the older alerts.
In the <i>Alert Details</i> page, the message “Last attempt failed” error is seen below the <i>Take Snapshot</i> button. Hovering over the error displays “Invoke API command has timed out for the data collector with id”.	This can happen when a data collector is added to Workload Security via SVM Management IP, if the LIF of the SVM is in <i>disabled</i> state in ONTAP. Enable the particular LIF in ONTAP and trigger <i>Take Snapshot manually</i> from Workload Security. The Snapshot action will then succeed.

Forensics

Forensics - All Activity

The All Activity page helps you understand the actions performed on entities in the Workload Security environment.

Examining All Activity Data

Click **Forensics > Activity Forensics** and click the **All Activity** tab to access the All Activity page. This page provides an overview of activities on your tenant, highlighting the following information:

- A graph showing *Activity History* (based on selected global time range)

You can zoom the graph by dragging out a rectangle in the graph. The entire page will be loaded to display the zoomed time range. When zoomed in, a button is displayed that lets the user zoom out.

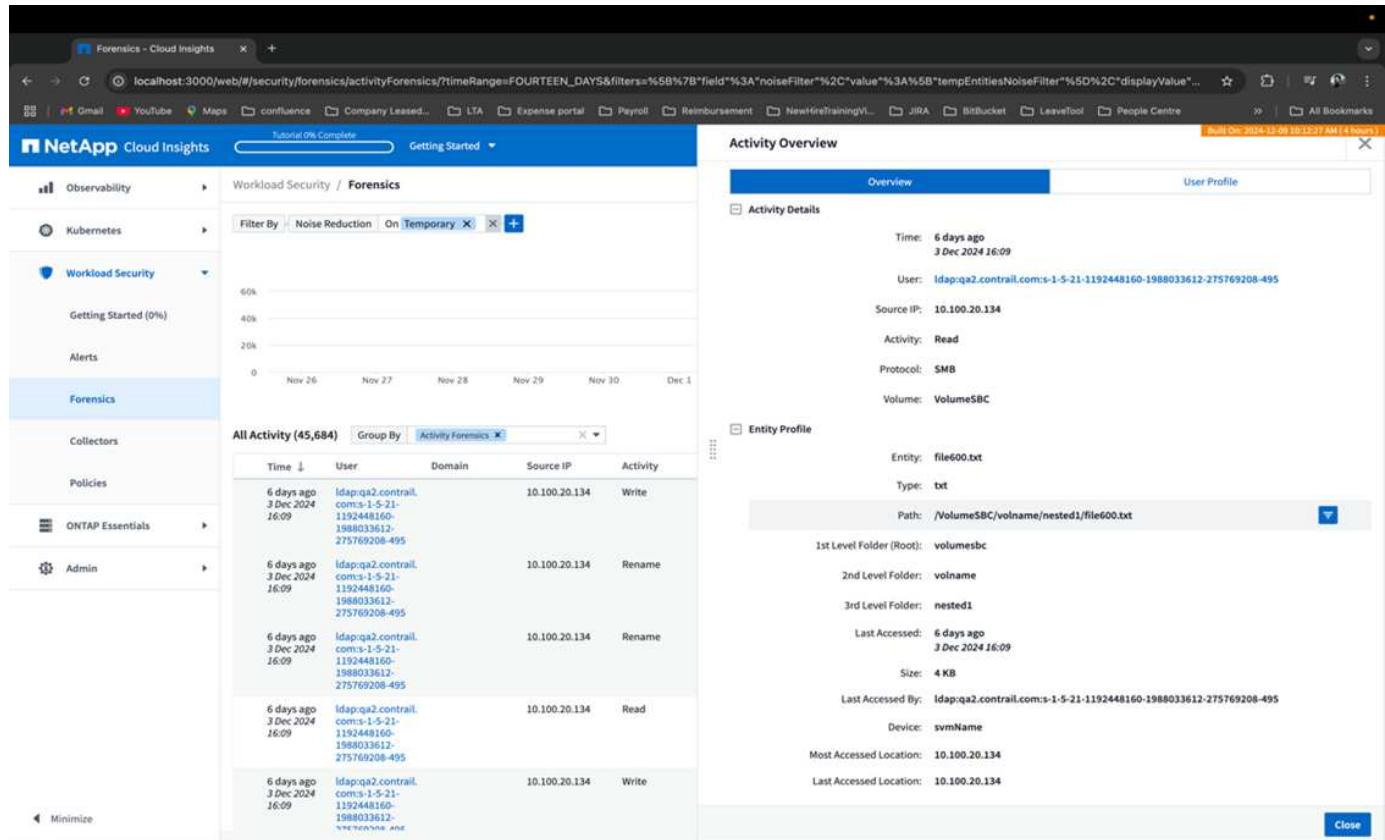
- A list of the *All Activity* data.
- A group by dropdown will provide the option to group the activity by users, folders, entity type, etc.
- A common path button will be available above the table on click of which we can get slide out panel with entity path details.

The **All Activity** table shows the following information. Note that not all of these columns are displayed by default. You can select columns to display by clicking on the "gear" icon.

- The **time** an entity was accessed including the year, month, day, and time of the last access.
- The **user** that accessed the entity with a link to the [User information](#) as a slide-out panel.
- The **activity** the user performed. Supported types are:
 - **Change Group Ownership** - Group Ownership is of file or folder is changed. For more details about group ownership please see [this link](#).
 - **Change Owner** - Ownership of file or folder is changed to another user.
 - **Change Permission** - File or folder permission is changed.
 - **Create** - Create file or folder.
 - **Delete** - Delete file or folder. If a folder is deleted, *delete* events are obtained for all the files in that folder and subfolders.
 - **Read** - File is read.
 - **Read Metadata** - Only on enabling folder monitoring option. Will be generated on opening a folder on Windows or Running "ls" inside a folder in Linux.
 - **Rename** - Rename file or folder.
 - **Write** - Data is written to a file.
 - **Write Metadata** - File metadata is written, for example, permission changed.
 - **Other Change** - Any other event which are not described above. All unmapped events are mapped to "Other Change" activity type. Applicable to files and folders.
- The **Path** is *entity path*. This should be either exact entity path (e.g., "/home/userX/nested1/nested2/abc.txt") OR directory portion of path for recursive search (e.g., "/home/userX/nested1/nested2"). NOTE: regex path patterns (e.g., *nested*) are NOT allowed here. Alternatively, individual path folder level filters as mentioned below can also be specified for path filtering.
- The **1st Level Folder (Root)** is the root directory of entity path in lower case.
- The **2nd Level Folder** is the second level directory of entity path in lower case.
- The **3rd Level Folder** is the third level directory of entity path in lower case.
- The **4th Level Folder** is the forth level directory of entity path in lower case.

- The **Entity Type**, including entity (i.e. file) extension (.doc, .docx, .tmp, etc.).
- The **Device** where the entities reside.
- The **Protocol** used to fetch events.
- The **Original Path** used for rename events when the original file was renamed. This column is not visible in the table by default. Use the column selector to add this column to the table.
- The **Volume** where the entities reside. This column is not visible in the table by default. Use the column selector to add this column to the table.
- The **Entity Name** is the last component of the entity path; For the Entity Type as file, it is the file name.

Selecting a table row opens a slide-out panel with the user profile in one tab, and the activity and entity overview in another tab.



The screenshot shows the NetApp Cloud Insights Forensics interface. On the left, there's a navigation sidebar with sections like Observability, Kubernetes, Workload Security, Forensics, and Admin. The Forensics section is selected and shows a table of activity logs. The table has columns for Time, User, Domain, Source IP, and Activity. One row is selected, showing details in a slide-out panel on the right. The slide-out panel is titled 'Activity Overview' and contains two tabs: 'Overview' (selected) and 'User Profile'. The 'Overview' tab shows 'Activity Details' with fields like Time (6 days ago, 3 Dec 2024 16:09), User (ldap:qa2.contrail.com:s-1-5-21-1192448160-1988033612-275769208-495), Source IP (10.100.20.134), Activity (Read), Protocol (SMB), and Volume (VolumeSBC). The 'Entity Profile' tab shows details for the selected file: Entity (file600.txt), Type (txt), Path (/VolumeSBC/volume/nested1/file600.txt), and 1st Level Folder (Root) (volumesbc). It also shows Last Accessed (6 days ago, 3 Dec 2024 16:09), Size (4 KB), Last Accessed By (ldap:qa2.contrail.com:s-1-5-21-1192448160-1988033612-275769208-495), Device (svmName), and Most Accessed Location (10.100.20.134). A 'Close' button is at the bottom right of the slide-out panel.

The default *Group by* method is *Activity forensics*. If you select a different *Group By* method—for example, Entity Type—the entity *Group By* table will be displayed. If no selection is made then *Group By all* is displayed.

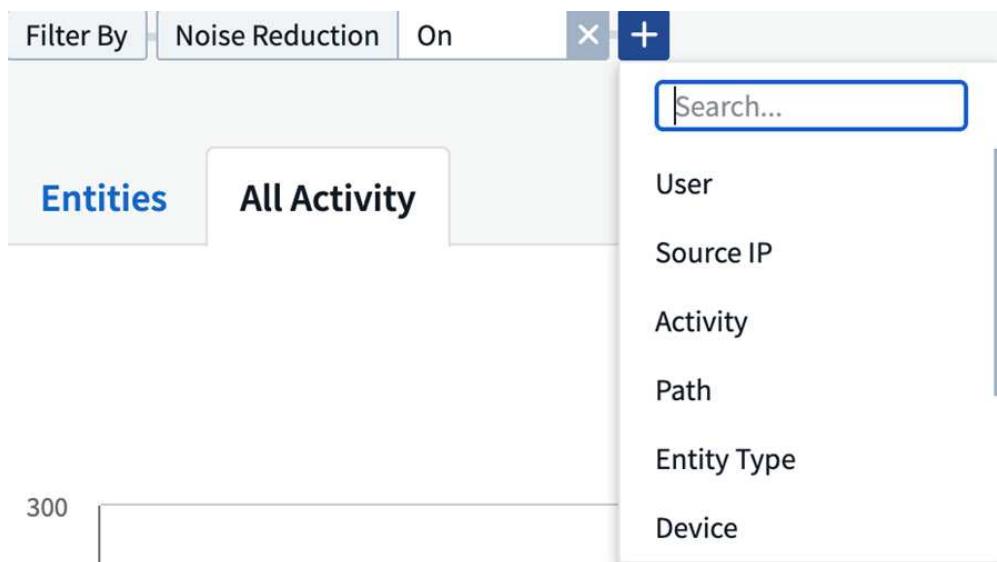
- Activity count is displayed as a hyperlink; selecting this will add the selected grouping as a filter. The table of activity will update based on that filter.
- Note that if you change the filter, alter the time range, or refresh the screen, you will not be able to return to the filtered results without setting the filter again.
- Please note that when Entity Name is selected as filter, the Group by dropdown will be disabled; Also, when the user is already on the Group By screen, the Entity Name as filter will be disabled.

Filtering Forensic Activity History Data

There are two methods you can use to filter data.

- The Filter can be added from the slide-out panel. The value is added to the appropriate filters in the top *Filter By* list.
- Filter data by typing in the *Filter By* field:

Select the appropriate filter from the top 'Filter By' widget by clicking the **[+]** button:



Enter the search text

Press Enter or click outside of the filter box to apply the filter.

You can filter Forensic Activity data by the following fields:

- The **Activity** type.
- **Protocol** to fetch protocol-specific activities.
- **Username** of the user performing the activity. You need to provide the exact Username to filter. Search with partial username, or partial username prefixed or suffixed with '*' will not work.
- **Noise Reduction** to filter files which are created in the last 2 hours by the user. It is also used to filter temporary files (for example, .tmp files) accessed by the user.
- **Domain** of the user performing the activity. You need to provide the **exact domain** to filter. Searching for partial domain, or partial domain prefixed or suffixed with wildcard ('*'), will not work. **None** can be specified to search for missing domain.

The following fields are subject to special filtering rules:

- **Entity Type**, using entity (file) extension - it is preferable to specify exact entity type within quotes. For example "txt".
- **Path** of the entity - This should be either exact entity path(e.g., "/home/userX/nested1/nested2/abc.txt") OR directory portion of path for recursive search(e.g., "/home/userX/nested1/nested2/"). NOTE: regex path patterns (e.g., *nested*) are NOT allowed here. Directory Path filters (path string ending with '/') up to 4 directories deep are recommended for faster results. For example, "/home/userX/nested1/nested2/". See the table below for more details.
- 1st Level Folder (Root) - root directory of entity Path as filters.
For example, if entity path is /home/userX/nested1/nested2/, then home OR "home" can be used.

- 2nd Level Folder - 2nd level directory of entity Path filters.
For example, if entity path is /home/userX/nested1/nested2/, then userX OR "userX" can be used.
- 3rd Level Folder – 3rd level directory of entity Path filters.
- For example, if entity path is /home/userX/nested1/nested2/, then nested1 OR "nested1" can be used.
- 4th Level Folder - Directory 4th level directory of entity Path filters.
For example, if entity path is /home/userX/nested1/nested2/, then nested2 OR "nested2" can be used.
- **User** performing the activity - it is preferable to specify the exact user within quotes. For example, "Administrator".
- **Device (SVM)** where entities reside
- **Volume** where entities reside
- The **Original Path** used for rename events when the original file was renamed.
- **Source IP** from which the entity was accessed.
 - You can use wild-cards * and ?. For example: 10.0.0., **10.0?.0.10, 10.10**
 - If exact match is required then, you must provide a valid source IP address in double quotes, for example "10.1.1.1.". Incomplete IPs with double quotes such as "10.1.1.", "10.1..*", etc. will not work.
- The **Entity Name** - the file name of the Entity Path as filters.
For example, if the entity path is /home/userX/nested1/testfile.txt then, entity name is testfile.txt.
Please note that it is recommended to specify the exact file name within quotes; Try to avoid the wildcard searches. For example, "testfile.txt".
Also, note that this entity name filter is recommended for shorter time ranges (up to 3 days).

The preceding fields are subject to the following when filtering:

- Exact value should be within quotes: Example: "searchtext"
- Wildcard strings must contain no quotes: Example: searchtext, *searchtext*, will filter for any strings containing 'searchtext'.
- String with a prefix, Example: searchtext*, will search any strings which start with 'searchtext'.

Please note that all filter fields are case in-sensitive search. For example: if the applied filter is Entity Type with value as 'searchtext', it will return results with Entity Type as 'searchtext', 'SearchText', 'SEARCHTEXT'

Activity Forensics Filter Examples:

User applied Filter expression	Expected Outcome	Performance assessment	Comment
Path = "/home/userX/nested1/nested2/"	Recursive lookup of all files and folders under given directory	Fast	Directory searches up to 4 directories will be fast.
Path = "/home/userX/nested1/"	Recursive lookup of all files and folders under given directory	Fast	Directory searches up to 4 directories will be fast.
Path = "/home/userX/nested1/test"	Exact match where path value matches with /home/userX/nested1/test	Slower	Exact search will be slower to search on compared to Directory searches.

User applied Filter expression	Expected Outcome	Performance assessment	Comment
Path = "/home/userX/nested1/nested2/nested3/"	Recursive lookup of all files and folders under given directory	Slower	More than 4 directories searches are slower to search on.
Any other Non path based filters. User and Entity Type filters recommended to be in quotes e.g., User="Administrator" Entity Type="txt"		Fast	
Entity Name = "test.log"	Exact match where file name is test.log	Fast	As it is exact match
Entity Name = *test.log	File names ending with test.log	Slow	Due to wild card, it can be slow.
Entity Name = test*.log	File names starting with test and ends with .log	Slow	Due to wild card, it can be slow.
Entity Name = test.lo	File names starting with test.lo For example: it will match test.log, test.log.1, test.log1	Slower	Due to wild card at the end, it can be slow.
Entity Name = test	File names starting with test	Slowest	Due to wild card at the end and more generic value used, it can be slowest.

NOTE:

1. The Activity count displayed alongside the All Activity icon is rounded off to 30 mins when the selected time range spans more than 3 days. e.g., a time range of *Sept 1st 10:15 am to Sept 7th 10:15 am* will show Activity counts from Sept 1st 10:00 am to Sept 7th 10:30 am.
2. Likewise the count metrics shown in Activity History graph are rounded off to 30 mins when the selected time range spans more than 3 days.

Sorting Forensic Activity History Data

You can sort activity history data by *Time*, *User*, *Source IP*, *Activity*, *Entity Type*, 1st Level Folder (Root), 2nd Level Folder, 3rd Level Folder and 4th Level Folder. By default, the table is sorted by descending *Time* order, meaning the latest data will be displayed first. Sorting is disabled for *Device* and *Protocol* fields.

User Guide for Asynchronous Exports

Overview

The Asynchronous Exports feature in Storage Workload Security is designed to handle large data exports.

Step-by-Step Guide: Exporting Data with Asynchronous Exports

- 1. Initiate Export:** Select the desired time duration and filters for the export and click on the export button.
- 2. Wait for Export to Complete:** The processing time can range from a few minutes to a few hours. You may need to refresh the forensics page a few times. Once the export job is complete, the "Download last export CSV file" button will be enabled.
- 3. Download:** Click on the "Download last created export file" button to get the exported data in a .zip format. This data will be available for download until the user initiates another Asynchronous Export or 3 days have elapsed, whichever occurs first. The button will remain enabled until another Asynchronous Export is initiated.
- 4. Limitations:**
 - The number of asynchronous downloads is currently limited to 1 per user for each Activities and Activities Analytics Table and 3 per tenant.
 - The exported data is limited to a maximum of 1 million records for Activities Table; while for Group By, the limit is half million records.

A sample script to extract forensic data via API is present at `/opt/netapp/cloudsecure/agent/export-script/` on the agent. See the readme at this location for more details about the script.

Column Selection for All Activity

The *All activity* table shows select columns by default. To add, remove, or change the columns, click the gear icon on the right of the table and select from the list of available columns.



Activity History Retention

Activity history is retained for 13 months for active Workload Security environments.

Applicability of Filters in Forensics Page

Filter	What it does	Example	Applicable for these Filters	Not applicable for these filters	Result
* (Asterisk)	enables you to search for everything	Auto*03172022 If search text contains hyphen or underscore, give expression in brackets. e.g., (svm*) for searching svm-123	User, Entity Type, Device, Volume, Original Path, 1stLevel Folder, 2ndLevel Folder, 3rdLevel Folder, 4thLevel Folder, Entity Name, Source IP		Returns all resources that start with "Auto" and end with "03172022"

Filter	What it does	Example	Applicable for these Filters	Not applicable for these filters	Result
? (question mark)	enables you to search for a specific number of characters	AutoSabotageUser1_03172022?	User, Entity Type, Device, Volume, 1stLevel Folder, 2ndLevel Folder, 3rdLevel Folder, 4thLevel Folder, Entity Name, Source IP		returns AutoSabotageUser1_03172022A, AutoSabotageUser1_03172022B, AutoSabotageUser1_031720225, and so on
OR	enables you to specify multiple entities	AutoSabotageUser1_03172022 OR AutoRansomUser4_03162022	User, Domain, Entity Type, Original Path, Entity Name, Source IP		returns any of AutoSabotageUser1_03172022 OR AutoRansomUser4_03162022
NOT	allows you to exclude text from the search results	NOT AutoRansomUser4_03162022	User, Domain, Entity Type, Original Path, 1stLevel Folder, 2ndLevel Folder, 3rdLevel Folder, 4thLevel Folder, Entity Name, Source IP	Device	returns everything that does not start with "AutoRansomUser4_03162022"
None	searches for NULL values in all fields	None	Domain		returns results where the target field is empty

Path Search

Search results with and without / will be different

"/AutoDir1/AutoFile03242022"	Only Exact search works; returns all activities with exact path as /AutoDir1/AutoFile03242022 (case insensitively)
"/AutoDir1/ "	Works; returns all activities with 1st level directory matching with AutoDir1 (case insensitively)
"/AutoDir1/AutoFile03242022/"	Works; returns all activities with 1st level directory matching with AutoDir1 and 2nd level directory matching with AutoFile03242022 (case insensitively)
/AutoDir1/AutoFile03242022 OR /AutoDir1/AutoFile03242022	Doesn't work
NOT /AutoDir1/AutoFile03242022	Doesn't work
NOT /AutoDir1	Doesn't work
NOT /AutoFile03242022	Doesn't work
*	Doesn't work

Local root SVM user activity changes

If a local root SVM user is performing any activity, the IP of the client on which the NFS share is mounted is now considered in the username, which will be shown as `root@<ip-address-of-the-client>` in both forensic activity and user activity pages.

For example:

- If SVM-1 is monitored by Workload Security, and the root user of that SVM mounts the share on a client with IP address 10.197.12.40, the username shown in forensic activity page will be `root@10.197.12.40`.
- If the same SVM-1 is mounted into another client with IP address 10.197.12.41, the username shown in forensic activity page will be `root@10.197.12.41`.

*• This is done to segregate NFS root user activity by IP address. Previously, all the activity was considered to be done by `root` user only, with no IP distinction.

Troubleshooting

Problem	Try This
In the "All Activities" table, under the 'User' column, the user name is shown as: "ldap:HQ.COMPANYNAME.COM:S-1-5-21-3577637-1906459482-1437260136-1831817" or "ldap:default:80038003"	Possible reasons could be: 1. No User Directory Collectors have been configured yet. To add one, go to Workload Security > Collectors > User Directory Collectors and click on +User Directory Collector . Choose <i>Active Directory</i> or <i>LDAP Directory Server</i> . 2. A User Directory Collector has been configured, however it has stopped or is in error state. Please go to Collectors > User Directory Collectors and check the status. Refer to the User Directory Collector troubleshooting section of the documentation for troubleshooting tips. After configuring properly, the name will get automatically resolved within 24 hours. If it still does not get resolved, check if you have added the correct User Data Collector. Make sure that the user is indeed part of the added Active Directory/LDAP Directory Server.
Some NFS events are not seen in UI.	Check the following: 1. A user directory collector for AD server with POSIX attributes set should be running with the <code>unixid</code> attribute enabled from UI. 2. Any user doing NFS access should be seen when searched in the user page from UI 3. Raw events (Events for whom the user is not yet discovered) are not supported for NFS 4. Anonymous access to the NFS export will not be monitored. 5. Make sure NFS version used is version 4.1 or less. (Note that NFS 4.1 is supported with ONTAP 9.15 or later.)

<p>After typing some letters containing a wildcard character like asterisk (*) in the filters on the Forensics <i>All Activity</i> or <i>Entities</i> pages, the pages load very slowly.</p>	<p>An asterisk (*) in the search string searches for everything. However, leading wildcard strings like *<i><searchTerm></i> or *<i><searchTerm></i>* will result in a slow query.</p> <p>To get better performance, use prefix strings instead, in the format <i><searchTerm></i>* (in other words, append the asterisk (*) <i>after</i> a search term).</p> <p>Example: use the string <i>testvolume</i>*, rather than *<i>testvolume</i> or *<i>test</i>*<i>volume</i>.</p> <p>Use a directory search to see all activities underneath a given folder recursively (Hierarchical search). e.g., "/path1/path2/path3/" will list all the activities recursively under /path1/path2/path3.</p> <p>Alternatively use the "Add To Filter" option under the All Activity tab."</p>
<p>I am encountering a "Request failed with status code 500/503" error when using a Path filter.</p>	<p>Try using a smaller date range for filtering records.</p>
<p>Forensic UI is loading data slowly when using the <i>path</i> filter.</p>	<p>Directory Path filters (path string ending with /) up to 4 directories deep are recommended for faster results. e.g., If the directory path is /Aaa/Bbb/Ccc/Ddd, try searching for "/Aaa/Bbb/Ccc/Ddd/" to load data faster.</p>
<p>Forensics UI is loading data slowly and facing failures when using the entity name filter.</p>	<p>Please try with smaller time-ranges and with exact value search with double quotes. e.g., If the entityPath is "/home/userX/nested1/nested2/nested3/testfile.txt" then, try with "testfile.txt" as entity name filter.</p>

Forensic User Overview

Information for each user is provided in the User Overview. Use these views to understand user characteristics, associated entities, and recent activities.

User Profile

User Profile information includes contact information and location of the user. The profile provides the following information:

- Name of the user
- Email address of the user
- User's Manager
- Phone contact for the user
- Location of the user

User Behavior

The user behavior information identifies recent activities and operations performed by the user. This information includes:

- Recent activity

- Last access location
- Activity graph
- Alerts
- Operations for the last seven days
 - Number of operations

Refresh Interval

The User list is refreshed every 12 hours.

Retention Policy

If not refreshed again, the User list is retained for 13 months. After 13 months, the data will be deleted. If your Workload Security environment is deleted, all data associated with the environment is deleted.

Automated Response Policies

Response Policies trigger actions such as taking a snapshot or restricting user access in the event of an attack or abnormal user behavior.

You can set policies on specific devices or all devices. To set a response policy, select **Admin > Automated Response Policies** and click the appropriate **+Policy** button. You can create policies for Attacks or for Warnings.

Add Attack Policy



Policy Name*

For Attack Type(s) *

- Ransomware Attack
- Data Destruction - File Deletion

On Device

[+ Another Device](#)

Actions

- Take Snapshot [?](#)
- Block User File Access [?](#)

Time Period

Webhooks Notifications

[Cancel](#)[Save](#)

You must save the policy with a unique name.

To disable an automated response action (for example, Take Snapshot), simply un-check the action and save the policy.

When an alert is triggered against the specified devices (or all devices, if selected), the automated response policy takes a snapshot of your data. You can see snapshot status on the [Alert detail page](#).

See the [Restrict User Access](#) page for more details on restricting user access by IP.

You can attach one or more webhooks to a policy to get notified when an alert is created and action is taken. It is recommended to add no more than 10 webhooks to a policy. Keep in mind that if a policy is paused, webhook notifications will not be triggered.

You can modify or pause an Automated Response Policy by choosing the option in the policy's drop-down menu.

Workload Security will automatically delete snapshots once per day based on the Snapshot Purge settings.

Snapshot Purge Settings

Define purge periods to automatically delete snapshots taken by Cloud Secure.

Attack Automated Response

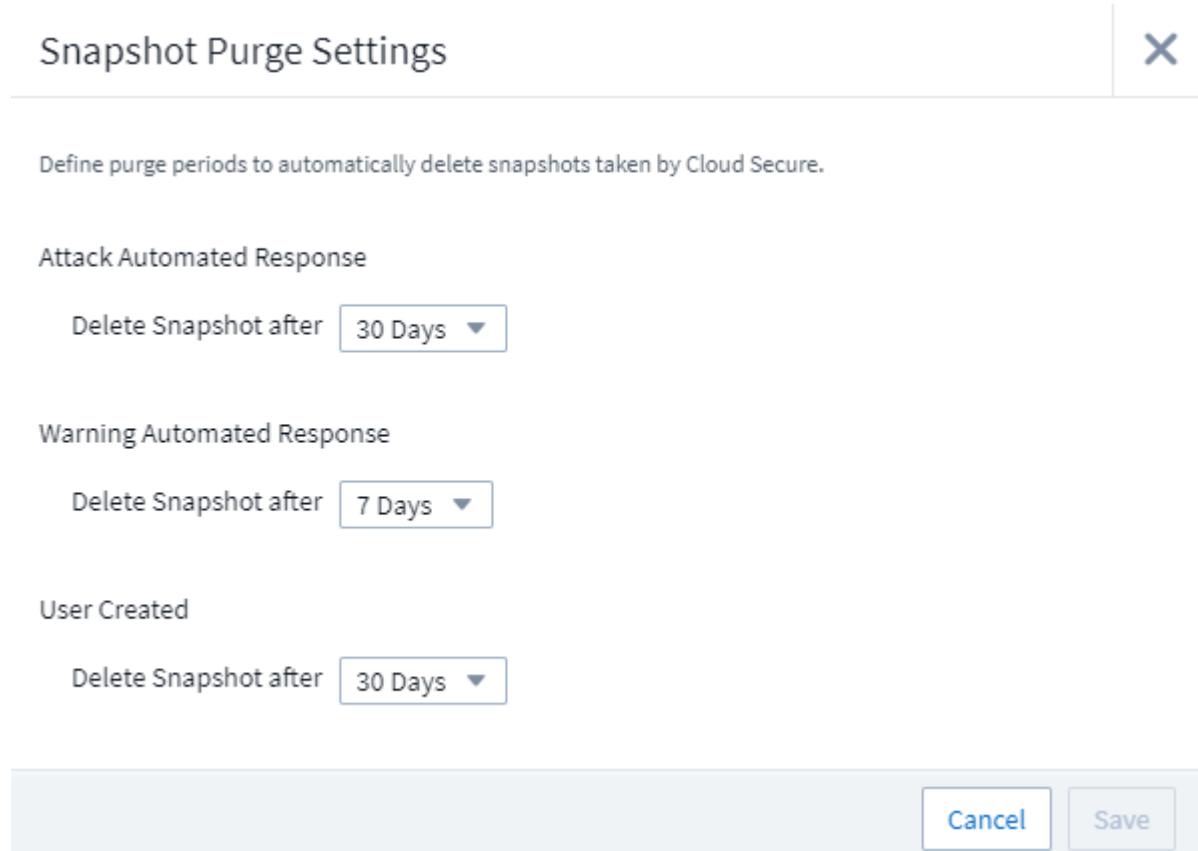
Delete Snapshot after

Warning Automated Response

Delete Snapshot after

User Created

Delete Snapshot after



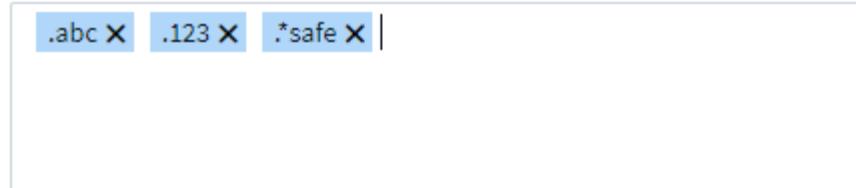
Allowed File Types Policies

If a file tampering attack is detected for a known file extension, and alerts are being generated on the Alerts screen, then that file extension can be added to an *allowed file types* list to prevent unnecessary alerting.

Navigate to **Workload Security > Policies** and go to the *Allowed File Type Policies* tab.

Allowed File Types Policies

Ransomware alerts will not be triggered for the following file types: 



Once added to the *allowed file types* list, no file tampering attack alert will be generated for that allowed file type. Note that the *Allowed File Types* policy is only applicable for file tampering detection.

For example, if a file named *test.txt* is renamed to *test.txt.abc* and Workload Security is detecting a file tampering attack because of the *.abc* extension, the *.abc* extension can be added to the *allowed file types* list. After being added to the list, file tampering attacks will no longer be generated against files with the *.abc* extension.

Allowed File Types can be exact matches (e.g., ".abc") or expressions (e.g., ".type", ".type", or "type"). Expressions of types ".a*c", ".p*f" are not supported.

Integration with ONTAP Autonomous Ransomware Protection

The ONTAP Autonomous Protection feature uses workload analysis in NAS (NFS and SMB) environments to proactively detect and warn about abnormal in-file activity that might indicate malicious attacks or unauthorized data modifications.

Additional details and license requirements about ARP can be found [here](#).

Workload Security integrates with ONTAP to receive ARP events and provide an additional analytics and automatic responses layer.

Workload Security receives the ARP events from ONTAP and takes the following actions:

1. Correlates volume encryption events with user activity to identify who is causing the damage.
2. Implements automatic response policies (if defined)
3. Provides forensics capabilities:
 - Allow customers to conduct data breach investigations.
 - Identify what files were affected, helping to recover faster and conduct data breach investigations.

Prerequisites

1. Minimum ONTAP version: 9.11.1

2. ARP enabled volumes. Details on enabling ARP can be found [here](#). ARP must be enabled via OnCommand System Manager. Workload Security cannot enable ARP.
3. Workload Security collector should be added via cluster IP.
4. Cluster level credentials are needed for this feature to work. In other words, cluster level credentials must be used when adding the SVM.

User permissions required

If you are using cluster administration credentials, no new permissions are needed.

If you are using a custom user (for example, *csuser*) with permissions given to the user, then follow the steps below to give permissions to Workload Security to collect ARP related information from ONTAP.

For *csuser* with cluster credentials, do the following from the ONTAP command line:

```
security login role create -role csrole -cmddirname "volume" -access
readonly
security login role create -role csrole -cmddirname "security anti-
ransomware volume" -access readonly
```

Read more about configuring other ONTAP permissions.

Sample Alert

A sample alert generated due to ARP event is shown below:

The screenshot shows a detailed alert interface for a 'POTENTIAL ATTACK: AL_1315 Ransomware Attack'. The alert was detected 5 months ago on Oct 20, 2022, 3:06 AM. It shows an icon of a face with a dollar sign for an eye. The 'Action Taken' section indicates 'Access Blocked on 5 SVMs' and 'Snapshots Taken'. The 'Status' is 'New'. Below this, there are buttons for 'Change Block Period', 'Re-Take Snapshots', and 'Unblock User'. A note says 'Blocked permanently by auto response policy'. A 'How To: Restore Entities' link is also present.

Total Attack Results

1	83	81
Affected Volumes	Deleted Files	Encrypted Files

81 Files have been copied, deleted, and potentially encrypted by 1 user account. The extension "osiris" was added to each file.

High Confidence Detection

Ransomware behavior and in-file encryption activities were detected.

Encrypted Files

Activity per minute

A line graph shows activity from 02:30 to 05:00. A sharp peak at 03:00 is labeled with a red circle and the letter 'I', representing 'Encryption activity in files'. A purple line labeled 'E' represents 'Encryption activity in files'.

Related Users

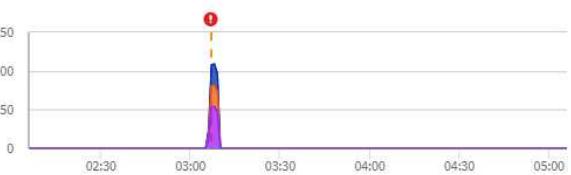


Jamelia Graham
Business Partner
HR

User/IP Access ?
Blocked

81 Encrypted Files
Detected 5 months ago
Oct 20, 2022 3:06 AM

View Activity Detail

Username us024	Department HR	Top Activity Types Activity per minute Last accessed from: 10.193.113.247
Domain cslab.netapp.com	Manager Iwan Holt	
Email Graham@netapp.com	Location WA	Create Read Others
Phone 9251140014		0 50 100 150 02:30 03:00 03:30 04:00 04:30 05:00

Access Limitation History for This User (3)

Time	Action	Duration	Action Taken by	Response	Blocked IPs on NFS
Oct 20, 2022 3:09 AM	⚠ Block more detail	Never Expires		Automatic	none
Mar 10, 2022 4:59 AM	Unblock		system	Blocking Expired	10.197.144.115
Mar 10, 2022 3:57 AM	⚠ Block more detail	1h		Automatic	10.197.144.115

Affected Devices/Volumes

Device ↑	Volume	Encrypted Files	Associated Snapshot Taken		
subprod_rtp	stargazer	81	Oct 20, 2022 3:09 AM	cloudsecure_attack_auto	Automatic Take Snapshot

A high confidence banner indicates the attack has shown file tampering behavior along with file encryption activities.

The encrypted files graph indicates the timestamp at which the volume encryption activity was detected by the ARP solution.

Limitations

In the case where an SVM is not monitored by Workload Security, but there are ARP events generated by ONTAP, the events will still be received and displayed by Workload Security. However, Forensic information related to the alert, as well as user mapping, will not be captured or shown.

Troubleshooting

Known problems and their resolutions are described in the following table.

Problem:	Resolution:
<p>Email alerts are received 24 hrs after an attack is detected. In the UI, the alerts are shown 24 hrs before that when the emails are received by Data Infrastructure Insights Workload Security.</p>	<p>When ONTAP sends the <i>Ransomware Detected</i> Event to Data Infrastructure Insights Workload Security (i.e. Workload Security), the email is sent. The Event contains a list of attacks and its timestamps. The Workload Security UI displays the alert timestamp of the first file attacked. ONTAP sends the <i>Ransomware Detected</i> Event to Data Infrastructure Insights when a certain number of files are encoded. Therefore, there may be a difference between the time the alert is displayed in the UI and the time the email is sent.</p>

Integration with ONTAP Access Denied

The ONTAP Access Denied feature uses workload analysis in NAS environments (NFS and SMB) to proactively detect and warn about failed file operations (i.e., a user trying to perform an operation for which they do not have permission). These failed file operation notifications—especially in cases of security-related failures—will further help in blocking insider attacks at early stages.

Data Infrastructure Insights Workload Security integrates with ONTAP to receive Access Denied events and provide an additional analytic and automatic response layer.

Prerequisites

- Minimum ONTAP version: 9.13.0.
- A Workload Security administrator must enable the Access Denied feature while adding a new collector or editing existing collector, by selecting the *Monitor Access Denied Events* checkbox under Advanced Configuration.

The screenshot shows the 'Add Data Collector' page in NetApp Cloud Insights. The left sidebar is titled 'Cloud Insights' and includes 'Observability', 'Kubernetes', 'Workload Security' (selected), 'Alerts', 'Forensics', 'Collectors' (selected), 'Policies', 'ONTAP Essentials', 'Cloud Cost', and 'Admin'. The main content area is titled 'Add Data Collector' and contains fields for 'Share Names' and 'Volume Names'. Under 'Advanced Configuration', there are checkboxes for 'Monitor Directory Read & Open Activity (SMB only)' and 'Monitor Access Denied Events' (which is checked). A note states: 'Note: This feature will be available from ONTAP 9.13 and above'. The 'Policy Server Send Buffer Size' is set to '1MB'. At the bottom are 'Cancel' and 'Save' buttons.

User permissions required

If the Data Collector is added using cluster administration credentials, no new permissions are needed.

If the Collector is added using a custom user (for example, *csuser*) with permissions given to the user, follow the steps below to give Workload Security the necessary permission to register for Access Denied events with ONTAP.

For *csuser* with *cluster* credentials, execute the following commands from the ONTAP command line. Note that this permission may already exist.

```
security login role create -role csrole -cmddirname "vserver fpolicy"
-access all
```

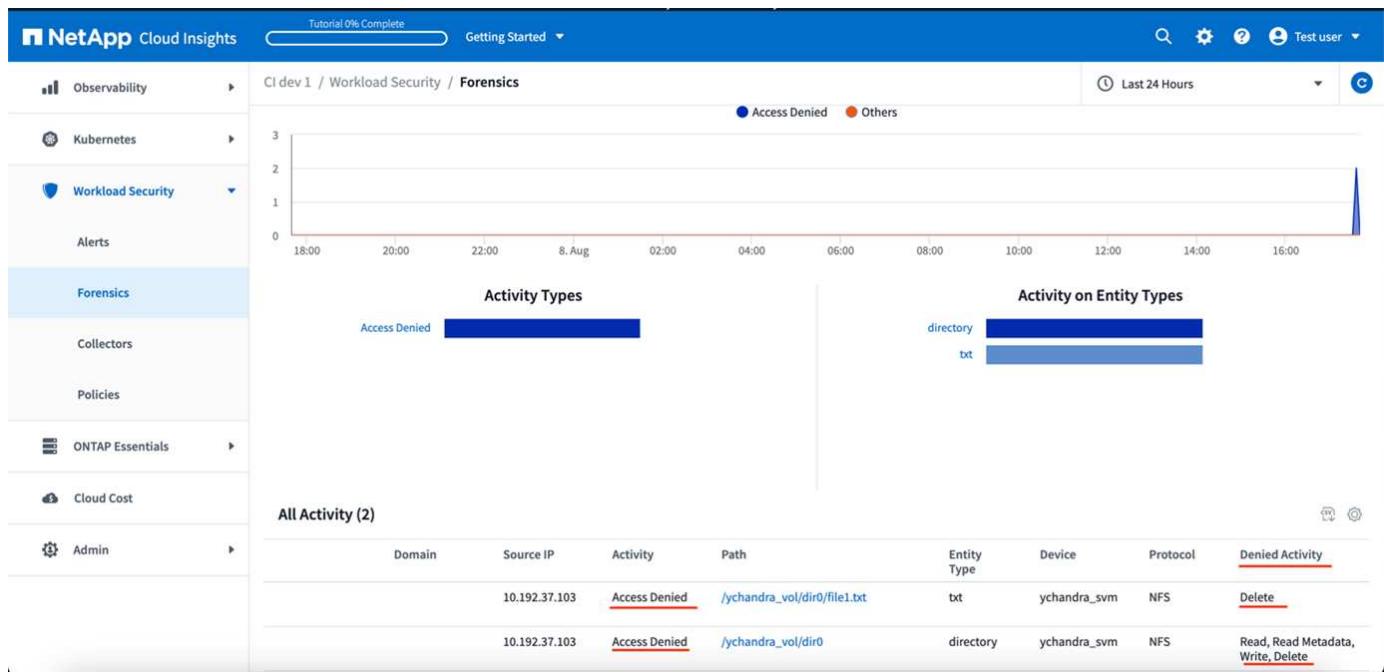
For *csuser* with *SVM* credentials, execute the following commands from the ONTAP command line. Note that this permission may already exist.

```
security login role create -vserver <vservername> -role csrole -cmddirname
"vserver fpolicy" -access all
```

Read more about configuring other [ONTAP permissions](#).

Access Denied events

Once events have been acquired from the ONTAP system, the Workload Security Forensics page will show Access Denied events. In addition to the information displayed, you can view the missing user permissions for a particular operation by adding the *Desired Activity* column to the table from the gear icon.



Blocking User Access to Stop Attacks

Immediately halt detected attacks by blocking compromised user access to prevent further data damage or exfiltration. Workload Security enables both automatic blocking through Automated Response Policies and manual intervention from alert or user details pages, giving you flexible control over your security response. Access restrictions apply automatically across all monitored storage volumes and are time-limited for automatic restoration.

User is directly blocked for SMB and IP address of the host machines causing the attack will be blocked for NFS. Those machine IP addresses will be blocked from accessing any of the Storage Virtual Machines (SVMs) monitored by Workload Security.

For example, let's say Workload Security manages 10 SVMs and the Automatic Response Policy is configured for four of those SVMs. If the attack originates in one of the four SVMs, the user's access will be blocked in all 10 SVMs. A Snapshot is still taken on the originating SVM.

If there are four SVMs with one SVM configured for SMB, one configured for NFS, and the remaining two configured for both NFS and SMB, all the SVMs will be blocked if the attack originates in any of the four SVMs.

Prerequisites for User Access Blocking

Cluster level credentials are needed for this feature to work.

If you are using cluster administration credentials, no new permissions are needed.

If you are using a custom user (for example, `csuser`) with permissions given to the user, then follow the steps below to give permissions to Workload Security to block user.

For `csuser` with cluster credentials, do the following from the ONTAP command line:

```
security login role create -role csrole -cmddirname "vserver export-policy rule" -access all
security login role create -role csrole -cmddirname set -access all
security login role create -role csrole -cmddirname "vserver cifs session" -access all
security login role create -role csrole -cmddirname "vserver services access-check authentication translate" -access all
security login role create -role csrole -cmddirname "vserver name-mapping" -access all
```

Be sure to review the Permissions section of the [Configuring the ONTAP SVM Data Collector](#) page as well.

How to enable the feature?

- In Workload Security, navigate to **Workload Security > Policies > Automated Response Policies**. Choose **+Attack Policy**.
- Select (check) *Block User File Access*.

How to set up Automatic user access blocking?

- Create a new Attack Policy or edit an existing Attack policy.
- Select the SVMs on which the attack policy should be monitored.
- Click on the checkbox “Block User File Access”. The feature will be enabled when this is selected.
- Under “Time Period” select the time until which the blocking should be applied.
- To test automatic user blocking,, you can simulate an attack via a [simulated script](#).

How to know if there are blocked users in the system?

- In the alert lists page, a banner on the top of screen will be displayed in case any user is blocked.
- Clicking on the banner will take you to the “Users” page, where the list of blocked users can be seen.
- In the “Users” page, there in a column named “User/IP Access”. In that column, the current state of user blocking will be displayed.

Restrict and manage user access manually

- You can go to the alert details or user details screen and then manually block or restore a user from those screens.

User Access Limitation History

In the alert details and user details page, in the user panel, you can view an audit of the user's access limitation history: Time, Action (Block, Unblock), duration, action taken by, manual/automatic, and affected IPs for NFS.

How to disable the feature?

At any time, you can disable the feature. If there are restricted users in the system, you must restore their access first.

- In Workload Security, navigate to **Workload Security > Policies > Automated Response Policies**. Choose **+Attack Policy**.
- De-select (uncheck) *Block User File Access*.

The feature will be hidden from all pages.

Manually Restore IPs for NFS

Use the following steps to manually restore any IPs from ONTAP if your Workload Security trial expires, or if the agent/collector is down.

1. List all export policies on an SVM.

```
contrail-qa-fas8020:> export-policy rule show -vserver <svm name>
      Policy          Rule  Access   Client          RO
Vserver   Name       Index  Protocol Match      Rule
-----  -----
svm0      default     1      nfs3,      cloudsecure_rule,  never
          nfs4,      10.11.12.13
          cifs
svm1      default     4      cifs,      0.0.0.0/0      any
          nfs
svm2      test        1      nfs3,      cloudsecure_rule,  never
          nfs4,      10.11.12.13
          cifs
svm3      test        3      cifs,      0.0.0.0/0      any
          nfs,
          flexcache
4 entries were displayed.
```

2. Delete the rules across all policies on the SVM which have “cloudsecure_rule” as Client Match by specifying its respective RuleIndex. Workload Security rule will usually be at 1.

```
contrail-qa-fas8020:*> export-policy rule delete -vserver <svm name>
-policyname * -ruleindex 1
```

3. Ensure Workload Security rule is deleted (optional step to confirm).

```
contrail-qa-fas8020:*> export-policy rule show -vserver <svm name>
      Policy          Rule  Access   Client          RO
Vserver   Name        Index  Protocol Match          Rule
-----
-----
svm0      default      4      cifs,      0.0.0.0/0      any
          nfs
svm2      test         3      cifs,      0.0.0.0/0      any
          nfs,
          flexcache
2 entries were displayed.
```

Manually Restore Users for SMB

Use the following steps to manually restore any users from ONTAP if your Workload Security trial expires, or if the agent/collector is down.

You can get the list of users blocked in Workload Security from the users list page.

1. Login to the ONTAP cluster (where you want to unblock users) with cluster *admin* credentials. (For Amazon FSx, login with FSx credentials).
2. Run the following command to list all users blocked by Workload Security for SMB in all SVMs:

```
vserver name-mapping show -direction win-unix -replacement " "
```

```
Vserver:  <vservername>
Direction: win-unix
Position Hostname          IP Address/Mask
-----
1          -                  -          Pattern: CSLAB\\US040
                                         Replacement:
2          -                  -          Pattern: CSLAB\\US030
                                         Replacement:
2 entries were displayed.
```

In the above output, 2 users were blocked (US030, US040) with domain CSLAB.

1. Once we identify the position from the above output, run the following command to unblock the user:

```
vserver name-mapping delete -direction win-unix -position <position>
```

2. Confirm the users are unblocked by running the command:

```
vserver name-mapping show -direction win-unix -replacement " "
```

No entries should be displayed for the users previously blocked.

Troubleshooting

Problem	Try This
Some of the users are not getting restricted, though there is an attack.	<ol style="list-style-type: none">1. Make sure that the Data Collector and Agent for the SVMs are in <i>Running</i> state. Workload Security won't be able to send commands if the Data Collector and Agent are stopped.2. This is because the user may have accessed the storage from a machine with a new IP which has not been used before. Restricting happens via IP address of the host through which the user is accessing the storage. Check in the UI (Alert Details > Access Limitation History for This User > Affected IPs) for the list of IP addresses which are restricted. If the user is accessing storage from a host which has an IP different from the restricted IPs, then the user will still be able to access the storage through the non-restricted IP. If the user is trying to access from the hosts whose IPs are restricted, then the storage won't be accessible.
Manually clicking on Restrict Access gives "IP addresses of this user have already been restricted".	The IP to be restricted is already being restricted from another user.
Policy could not be modified. Reason: not authorized for that command.	Check if using csuser, permissions are given to the user as mentioned above.

Problem	Try This
<p>User (IP Address) blocking for NFS works, but for SMB / CIFS, I see an error message: "SID to DomainName transformation failed. Reason timeout: socket is not established"</p>	<p>This can happen if <i>csuser</i> does not have permission to perform ssh. (Ensure connection at cluster level, then ensure user can perform ssh). <i>csuser</i> role requires these permissions.</p> <p>https://docs.netapp.com/us-en/cloudinsights/cs_restrict_user_access.html#prerequisites-for-user-access-blocking</p> <p>For <i>csuser</i> with cluster credentials, do the following from the ONTAP command line:</p> <pre>security login role create -role csrole -cmddirname "vserver export-policy rule" -access all security login role create -role csrole -cmddirname set -access all security login role create -role csrole -cmddirname "vserver cifs session" -access all security login role create -role csrole -cmddirname "vserver services access-check authentication translate" -access all security login role create -role csrole -cmddirname "vserver name-mapping" -access all</pre> <p>If <i>csuser</i> is not used and if admin user at cluster level is used, make sure that the admin user has ssh permission to ONTAP.</p>
<p>I'm getting the Error Message <i>SID translate failed. Reason:255:Error: command failed: not authorized for that commandError: "access-check" is not a recognized command</i>, when a user should have been blocked.</p>	<p>This can happen when <i>csuser</i> does not have correct permissions. See Prerequisites for User Access Blocking for more information.</p> <p>After applying the permissions, it is recommended to restart the ONTAP data collector and User Directory data collector. The required permission commands are listed below.</p> <pre>----- security login role create -role csrole -cmddirname "vserver export-policy rule" -access all security login role create -role csrole -cmddirname set -access all security login role create -role csrole -cmddirname "vserver cifs session" -access all security login role create -role csrole -cmddirname "vserver services access-check authentication translate" -access all security login role create -role csrole -cmddirname "vserver name-mapping" -access all -----</pre>

Workload Security: Simulating File Tampering

You can use the instructions on this page to simulate file tampering for testing or demonstrating Workload Security using the included file tampering Simulation script.

Things to note before you begin

- The file tampering simulation script works on Linux only. The simulation script should also generate High Confidence Alerts in the event that the user has integrated ONTAP ARP with Workload Security.
- Workload Security will detect events and alerts generated with NFS 4.1 only if ONTAP version is 9.15 or higher.
- The script is provided with the Workload Security agent installation files. It is available on any machine that has a Workload Security agent installed.
- You can run the script on the Workload Security agent machine itself; there is no need to prepare another Linux machine. However, if you prefer to run the script on another system, simply copy the script and run it there.
- Users can opt for either the Python or shell script based on their preferences and system requirements.
- The Python script has pre-requisite installations. If you don't want to use python, use the shell script.

Guidelines:

This script should be executed on an SVM containing a folder with a substantial number of files for encryption, ideally 100 or more, including files in sub-folders. Ensure that the files are not empty.

To generate the alert, temporarily pause the collector before test data creation. Once the sample files are generated, resume the collector and initiate the encryption process.

Steps:

Prepare the system:

First, mount the target volume to the machine. You can mount either an NFS or CIFS export.

To mount NFS export in Linux:

```
mount -t nfs -o vers=4.0 10.193.177.158:/svmvol1 /mntpt
mount -t nfs -o vers=4.0 Vserver data IP>:/nfsvol /destinationlinuxfolder
```

Do not mount NFS version 4.1; it is not supported by Fpolicy.

To mount CIFS in Linux:

```
mount -t cifs //10.193.77.91/sharedfolderincluster
/root/destinationfolder/ -o username=raisa
```

Enable ONTAP Autonomous Ransomware Protection (Optional):

If your ONTAP cluster version is 9.11.1 or higher you may enable the ONTAP Ransomware Protection service by executing the following command on the ONTAP command console.

```
security anti-ransomware volume enable -volume [volume_name] -vserver [svm_name]
```

Next, set up a Data Collector:

1. Configure the Workload Security agent if not already done.
2. Configure an SVM data collector if not already done.
3. Make sure the mount protocol is selected while configuring the data collector.

Generate the sample files programmatically:

Before creating the files, you must first stop or [pause the data collector](#) processing.

Before running the simulation, you must first add files to be encrypted. You can either manually copy the files to be encrypted into the target folder, or use one of the included scripts to programmatically create the files. Whichever method you use, make sure at least 100 files are present to encrypt.

If you choose to programmatically create the files, you can use the Shell or Python:

Shell:

1. Log into the Agent box.
2. Mount an NFS or CIFS share from the SVM of the filer to the Agent machine. Cd to that folder.
3. Copy the script from Agent installation directory (%AGENT_INSTALL_DIR/agent/install/ransomware_simulation/shell/create_dataset.sh) to the target mount location.
4. Execute the following command using the scripts within the mounted directory (e.g. /root/demo) to create the test dataset folder and files:

```
'./create_dataset.sh'
```

5. This will create 100 non empty files with various extensions inside the mount folder under a directory called "test_dataset".

Python:

Python script Prerequisite:

- Install Python (if not already installed.)
 - Download Python 3.5.2 or above from <https://www.python.org/>.
 - To check Python installation, run `python --version`.
 - The Python script has been tested on versions as early as 3.5.2.

- Install pip if not already installed:
 - Download the get-pip.py script from <https://bootstrap.pypa.io/>.
 - Install pip using `python get-pip.py`.
 - Verify pip installation with `pip --version`.
- PyCryptodome Library:
 - The script uses the PyCryptodome library.
 - Install PyCryptodome with `pip install pycryptodome`.
 - Confirm PyCryptodome installation by running `pip show pycryptodome`.

Python create file script:

1. Log into the Agent box.
2. Mount an NFS or CIFS share from the SVM of the filer to the Agent machine. Cd to that folder.
3. Copy the script from Agent installation directory
(%AGENT_INSTALL_DIR/agent/install/ransomware_simulation/python/create_dataset.py) to the target mount location.
4. Execute the following command using the scripts within the mounted directory (for e.g. /root/demo) to create the test dataset folder and files:

```
'python create_dataset.py'
```

5. This will create 100 non empty files with various extensions inside the mount folder under a directory called "test_dataset"

Resume the collector

If you paused the collector before following these steps, please be sure to resume the collector once the sample files are created.

Generate the sample files programmatically:

Before creating the files, you must first stop or [pause the data collector](#) processing.

To generate a file tampering alert, you can execute the included script which will simulate a file tampering alert in Workload Security.

Shell:

1. Copy the script from Agent installation directory
(%AGENT_INSTALL_DIR/agent/install/ransomware_simulation/shell/simulate_attack.sh) to the target mount location.
2. Execute the following command using the scripts within the mounted directory (for e.g. /root/demo) to encrypt the test dataset:

```
'./simulate_attack.sh'
```

3. This will encrypt the sample files created under the "test_dataset" directory.

Python:

1. Copy the script from Agent installation directory (%AGENT_INSTALL_DIR/agent/install/ransomware_simulation/python/simulate_attack.py) to the target mount location.
2. Please note that python prerequisites are installed as per Python script Prerequisite section
3. Execute the following command using the scripts within the mounted directory (for e.g. /root/demo) to encrypt the test dataset:

```
'python simulate_attack.py'
```

4. This will encrypt the sample files created under the "test_dataset" directory.

Generate an Alert in Workload Security

Once the simulator script execution has finished, an alert will be seen on the Web UI within a few minutes.

Note: In the event that all of the following conditions are met, a High Confidence Alert will be generated.

1. Monitored SVM's ONTAP version higher than 9.11.1
2. ONTAP Autonomous Ransomware Protection configured
3. Workload Security Data collector is added in Cluster mode.

Workload Security detects file tampering patterns based on user behaviour while ONTAP ARP detects file tampering activity based on encryption activities in files.

If the conditions are met, Workload Security marks the alerts as High Confidence Alert.

Example of High Confidence Alert on the Alerts list page:

Example of High Confidence Alert detail:

Triggering alert multiple times

Workload Security learns user behavior and will not generate alerts on repeated file tampering attacks within 24 hours for the same user.

To generate a new alert with a different user, please follow the same steps again (creating test data and then encrypting the test data).

Configuring Email Notifications for Alerts, Warnings, and Agent/Data Source Collector health

Email notifications enable you to stay informed about potential attacks, security warnings, and infrastructure health issues as they occur. Configure recipient email addresses in the Admin > Notifications settings to receive real-time alerts tailored to each recipient's responsibilities.

Potential Attack Alerts and Warnings

To send *Potential Attack* alert notifications, enter the recipients' email addresses in the *Send Potential Attack Alerts* section.

Email notifications are sent to the alert recipient list for every action on the alert.

To send *Warning* notifications, enter the recipients' email addresses in the *Send Warning Alerts* section.

Agent and Data Collector Health monitoring

You can monitor the health of Agents and Data Sources through notifications.

In order to receive notifications in the event that an Agent or Data Source collector is not functioning, enter the email addresses of the recipients in the *Data Collection Health Alerts* section.

Keep the following in mind:

- Health alerts will be sent only after the agent/collector stops reporting for at least one hour.
- Only one email notification is sent to the intended recipients in a given 24 hour period, even if the Agent or Data collector is disconnected for a longer duration.
- In case of an Agent failure, one alert will be sent (not one per collector). The email will include a list of all impacted SVMs.
- Active directory collection failure is reported as a warning; it does not impact threat detection.
- The Getting Started setup list now includes a new *Configure email notifications* phase.

Receiving Agent And Data Collector Upgrade Notifications

- Enter the email ID(s) in the “Data Collection Health Alerts”.
- The “Enable upgrade notifications” check box becomes enabled.
- Agent and Data Collector upgrade email notifications are sent to the email IDs one day in advance of the planned upgrade.

Troubleshooting

Problem:	Try this:
Email IDs are present in the “Data Collector Health Alerts”, however I am not receiving notifications.	Notification emails are sent from the NetApp Data Infrastructure Insights domain, i.e from accounts@service.cloudinsights.netapp.com . Some companies block incoming emails if they are from an external domain. Ensure that external notifications from NetApp Data Infrastructure Insights domains are whitelisted.

Webhook Notifications

Workload Security notifications using webhooks

Webhooks allow users to send critical or warning alert notifications to various applications using a customized webhook channel.

Many commercial applications support webhooks as a standard input interface, for example: Slack, PagerDuty, Teams, and Discord. By supporting a generic, customizable webhook channel, Workload Security can support many of these delivery channels. Information about configuring the webhooks can be found on the respective application’s websites. For example, Slack provides [this useful guide](#).

You can create multiple webhook channels, each channel targeted for a different purpose, separate applications, different recipients, etc.

The webhook channel instance is comprised of the following elements

Name	Description
URL	Webhook target URL, including the http:// or https:// prefix along with the url params
Method	GET/POST - Default is POST
Custom Header	Specify any custom headers here
Message Body	Put the body of your message here
Default Alert Parameters	Lists the default parameters for the webhook
Custom Parameters and Secrets	Custom parameters and secrets allows you to add unique parameters and secure elements such as passwords

Creating a webhook

To create a Workload Security Webhook, go to Admin > Notifications and select “Workload Security Webhooks” tab. The following image shows a sample slack webhook creation screen.

Note: User must be a Workload Security *Admin* in order to create and manage Workload Security Webhooks.

Add a Webhook

Name
Test-Webhook-1

Template Type
Slack

URL ?
https://hooks.slack.com/services/<id>

Validate SSL Certificate for secure communication

Method
POST

Custom Header
Content-type: application/json
Accept: application/json

Message Body
{
 "blocks": [
 {
 "type": "section",
 "text": {
 "type": "mrkdwn",
 "text": "*%severity% Alert: %%synopsis%*"
 }
 },
 {
 "type": "divider"
 }
]
}

Buttons
Cancel Test Webhook Create Webhook

- Enter appropriate information for each of the fields, and click "Save".
- You can also click the "Test Webhook" button to test the connection. Note that this will send the "Message Body" (without substitutions) to the defined URL according to the selected Method.
- SWS webhooks comprise a number of default parameters. Additionally, you can create your own custom parameters or secrets.

Parameters: What are they and how to use them?

Alert Parameters are dynamic values populated per alert. For example, the `%%severity%%` parameter will be replaced with the severity type of the alert.

Note that substitutions are not performed when clicking the "Test Webhook" button; the test sends a payload that shows the parameter's placeholders (`%%<param-name>%%`) but does not replace them with data.

Custom Parameters and Secrets

In this section you can add any custom parameters and/or secrets you wish. A custom parameter or secret can be in the URL or message body. Secrets allow user to configure a secure custom parameter like password, apiKey etc.

The following sample image shows how custom parameters are used in webhook creation.

The screenshot shows the 'Add Webhook' interface. On the left, there are configuration fields: 'Template Type' (Slack), 'URL' (https://hooks.slack.com/services/%%slack-id%%), 'Method' (POST), 'Custom Header' (Content-type: application/json, Accept: application/json), and 'Message Body' (JSON template). The 'Message Body' field contains the following JSON:

```

    {
      "text": "Status: %%status%%",
      "type": "mrkdwn",
      "text": "Configured by: %%webhookConfiguredBy%%"
    }
  }
  {
  }
  
```

On the right, a table lists custom parameters and secrets:

%%alertDetailsPageUrl%%	https://%%cloudInsightsHostname%%/%%alertDetailsPageUrl%%
%%alertTimestamp%%	Alert timestamp in Epoch format (milliseconds)
%%changePercentage%%	Change Percentage
%%detected%%	Alert timestamp in GMT (Tue, 27 Oct 2020 01:20:30 GMT)
%%id%%	Alert ID
%%note%%	Note
%%severity%%	Alert severity
%%status%%	Alert status
%%synopsis%%	Alert Synopsis
%%type%%	Alert type
%%userId%%	User id
%%userName%%	User name
%%filesDeleted%%	Files deleted
%%encryptedFilesSuffix%%	Encrypted files suffix
%%filesEncrypted%%	Files encrypted

The 'Custom Parameters and Secrets' section is highlighted with a red box, showing the following entries:

Name	Value	Description
%%webhookConfiguredBy	system_admin_1	...
%%slack-id%%	*****	*****

Custom Parameters and Secrets

+ Parameter

Workload Security Webhooks List Page

On the Webhooks list page, displayed are the Name, Created By, Created On, Status, Secure, and Last Reported fields.

Note: The value of 'status' column will keep changing based on the result of last webhook trigger result. The following are examples of status results.

Status	Description
OK	Successfully sent notification.
403	Forbidden.
404	URL not found.

400	Bad Request. You might see this status if there is any error in the message body, for example: <ul style="list-style-type: none"> • Badly formatted json. • Providing invalid value for reserved keys. For example, PagerDuty accepts only critical/warning/error/info for “Severity”. Any other result may yield a 400 status. • Application specific validation errors. For example, Slack allows a maximum of 10 fields inside a section. Including more than 10 may result in a 400 status.
410	Resource is no longer available

“Last Reported” column indicates the time when the webhook was last triggered.

From the webhooks listing page users can also Edit/Duplicate/Delete webhooks.

Configure Webhook notification in alert policy

To add a webhook notification to an alert policy, go to -Workload Security > Policies- and select an existing policy or add a new policy. In the *Actions* section > *Webhook Notifications* dropdown, select the required webhooks.

Edit Attack Policy

X

Policy Name*

For Attack Type(s) *

Ransomware Attack
 Data Destruction - File Deletion

On Device

▼

+ Another Device

Actions

Take Snapshot ?
 Block User File Access ?

Time Period

▼

Webhooks Notifications

▼

Test-Webhook-1

Cancel
Save

Webhook notifications are tied to policies. When the attack (RW/DD/WARN) happens, the action configured (Take snapshot / user blocking) will be taken and then the associated webhook notification will be triggered.

Note: Email notifications are independent of policies, they will be triggered as usual.

- If a policy is paused, webhook notifications will not be triggered.
- Multiple webhooks can be attached to a single policy but it is recommended to attach no more than 5 webhooks to a policy.

Workload Security Webhook Examples

Webhooks for [Slack](#)

Webhooks for [PagerDuty](#)

Webhooks for [Teams](#)

Webhooks for [Discord](#)

Workload Security Webhook Example for Discord

Webhooks allow users to send alert notifications to various applications using a customized webhook channel. This page provides an example for setting up webhooks for Discord.



This page refers to third-party instructions, which are subject to change. Refer to the [Discord documentation](#) for the most up-to-date information.

Discord Setup:

- In Discord, select the Server, under Text Channels, select Edit Channel (gear icon)
- Select **Integrations > View Webhooks** and click **New Webhook**
- Copy the Webhook URL. You will need to paste this into the Workload Security webhook configuration.

Create Workload Security Webhook:

1. Navigate to Admin > Notifications and select the *Workload Security Webhooks* tab. Click '+ Webhook' to create a new webhook.
2. Give the webhook a meaningful Name.
3. In the *Template Type* drop-down, select **Discord**.
4. Paste the Discord URL from above into the *URL* field.

Add a Webhook

Name

Template Type

URL

Validate SSL Certificate for secure communication

Method

Custom Header

```
Content-Type: application/json
Accept: application/json
```

Message Body

```
{
  "content": null,
  "embeds": [
    {
      "title": "%%severity%% | %%id%%",
      "description": "%%synopsis%%",
      "url": "https://%%cloudInsightsHostname%%/%%alertDetailsPageUrl%%",
      "color": 3244733,
      "fields": [
        {
          "name": "%%",
          "value": "%%"
        }
      ]
    }
  ]
}
```

In order to test the webhook, temporarily replace the URL value in the message body with any valid URL (such as <https://netapp.com>) then click the *Test Webhook* button. Discord requires that a valid URL be supplied in order for Test Webhook functionality to work.

Be sure to set the message body back once the test completes.

Notifications via Webhook

To notify on events via webhook, navigate to *Workload Security > Policies*. Click on *+Attack Policy* or *+Warning Policy*.

- Enter a meaningful policy name.
- Select the required Attack Type(s), Devices to which policy should be attached, and required Actions.
- Under the *Webhooks Notifications* dropdown, select the required Discord webhooks and save.

Note: Webhooks can also be attached to existing policies by editing them.

Add Attack Policy



Policy Name*

For Attack Type(s) *

- Ransomware Attack
- Data Destruction - File Deletion

On Device

[+ Another Device](#)

Actions

- Take Snapshot [?](#)
- Block User File Access [?](#)

Time Period

Webhooks Notifications

Test-Webhook-1

[Cancel](#)[Save](#)

Workload Security Webhook Example for PagerDuty

Webhooks allow users to send alert notifications to various applications using a customized webhook channel. This page provides an example for setting up webhooks

for PagerDuty.



This page refers to third-party instructions, which are subject to change. Refer to the [PagerDuty documentation](#) for the most up-to-date information.

PagerDuty Setup:

1. In PagerDuty, navigate to **Services > Service Directory** and click on the **+New Service** button.
2. Enter a *Name* and select *Use our API directly*. Select **Add Service**.

Add a Service

A service may represent an application, component or team you wish to open incidents against.

General Settings

Name

Description

Add a description for this service (optional)

Integration Settings

Connect with one of PagerDuty's supported integrations, or create a custom integration through email or API. Alerts from a service from a supported integration or through the Events V2 API.

You can add more than one integration to a service, for example, one for monitoring alerts and one for [change events](#).

Integration Type



Select a tool

PagerDuty integrates with hundreds of tools, including monitoring tools, ticketing systems, code repositories, and deploy pipelines. This may involve configuration steps in the tool you are integrating with PagerDuty.

Integrate via email

If your monitoring tool can send email, it can integrate with PagerDuty using a custom email address.

Use our API directly

If you're writing your own integration, use our Events API. More information is in our [developer documentation](#).

Events API v2

Don't use an integration

If you only want incidents to be manually created, You can always add additional integrations later.

3. Select the **Integrations** tab to see the **Integration Key**. You will need this key when you create the Workload Security webhook below.
1. Go to **Incidents** or **Services** to view Alerts.

Open Incidents (5)

					All statuses	Go to incident #	25 per page	1 - 5 of 5
Status	Priority	Urgency	Alerts	Title	Assigned To	Created		
<input type="checkbox"/> Acknowledged	High	1	Critical Alert: Ransomware attack from user account #403982	+ SHOW DETAILS (1 triggered alert)	Chandan SS	Today at 4:11 AM		
<input type="checkbox"/> Acknowledged	High	1	Critical Alert: Data Destruction - File Deletion attack from user account #403996	+ SHOW DETAILS (1 triggered alert)	Chandan SS	Today at 5:41 AM		

Create Workload Security PagerDuty Webhook:

- Navigate to Admin > Notifications and select the *Workload Security Webhooks* tab. Select '+ Webhook' to create a new webhook.
- Give the webhook a meaningful name.
- In the *Template Type* dropdown, select *PagerDuty Trigger*.
- Create a custom parameter secret named *routingKey* and set the value to the PagerDuty *Integration Key* created above.

Custom Parameters and Secrets 

Name	Value ↑	Description
%%routingKey%%	*****	

 [Parameter](#)

Name 	Value
routingKey	*****
Type	Description
Secret	

[Cancel](#) [Save Parameter](#)

Add a Webhook

Name
Test PagerDuty

Template Type
PagerDuty Trigger

URL 
https://events.pagerduty.com/%%pagerDutyId%%

Validate SSL Certificate for secure communication

Method
POST

Custom Header
Content-Type: application/json
Accept: application/json

Message Body

```
{
  "dedup_key": "%%id%%",
  "event_action": "trigger",
  "links": [
    {
      "href": "https://%%cloudInsightsHostname%%/%%alertDetailsPageUrl%%",
      "text": "%%severity%% | %%id%% | %%detected%%"
    }
  ],
  "payload": {
    "severity": "%%severity%%"
  }
}
```

Buttons

[Cancel](#) [Test Webhook](#) [Create Webhook](#)

Notifications via Webhook

- To notify on events via webhook, navigate to *Workload Security > Policies*. Select *+Attack Policy* or *+Warning Policy*.
- Enter a meaningful policy name.
- Select required Attack Type(s), Devices to which the policy should be attached, and the required Actions.
- Under *Webhooks Notifications* dropdown, select the required PagerDuty webhooks. Save the policy.

Note: Webhooks can also be attached to existing policies by editing them.

Add Attack Policy

Policy Name*
Test policy 1

For Attack Type(s) *

Ransomware Attack
 Data Destruction - File Deletion

On Device

All Devices

+ Another Device

Actions

Take Snapshot ?
 Block User File Access ?

Time Period

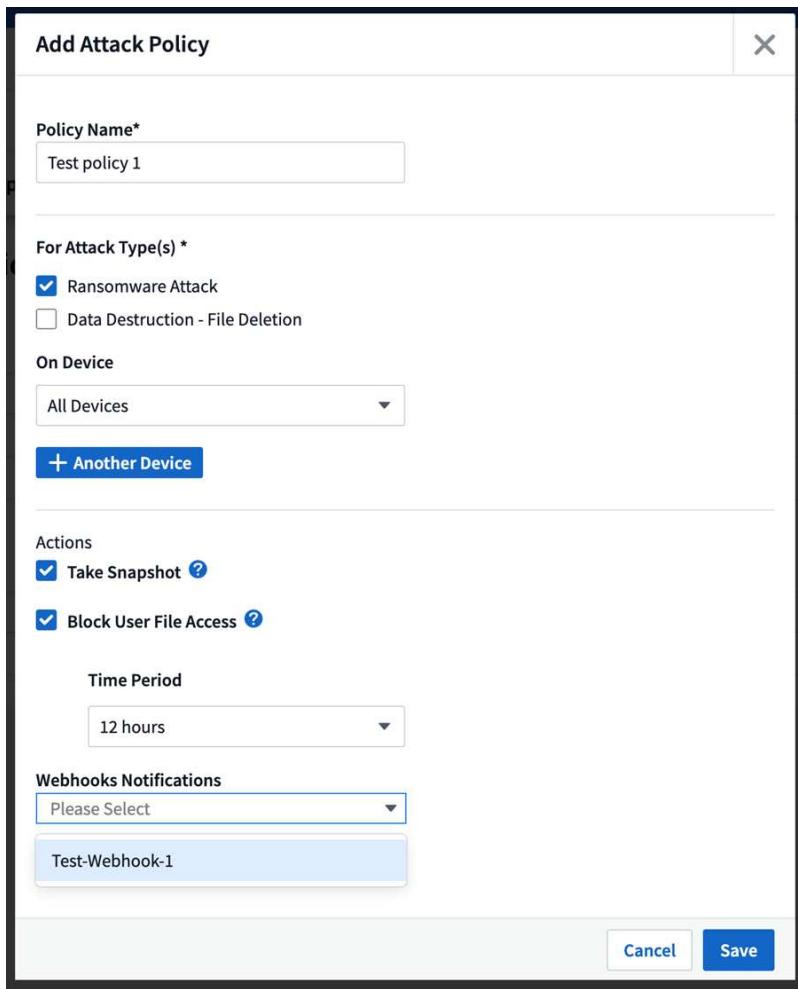
12 hours

Webhooks Notifications

Please Select

Test-Webhook-1

Cancel **Save**



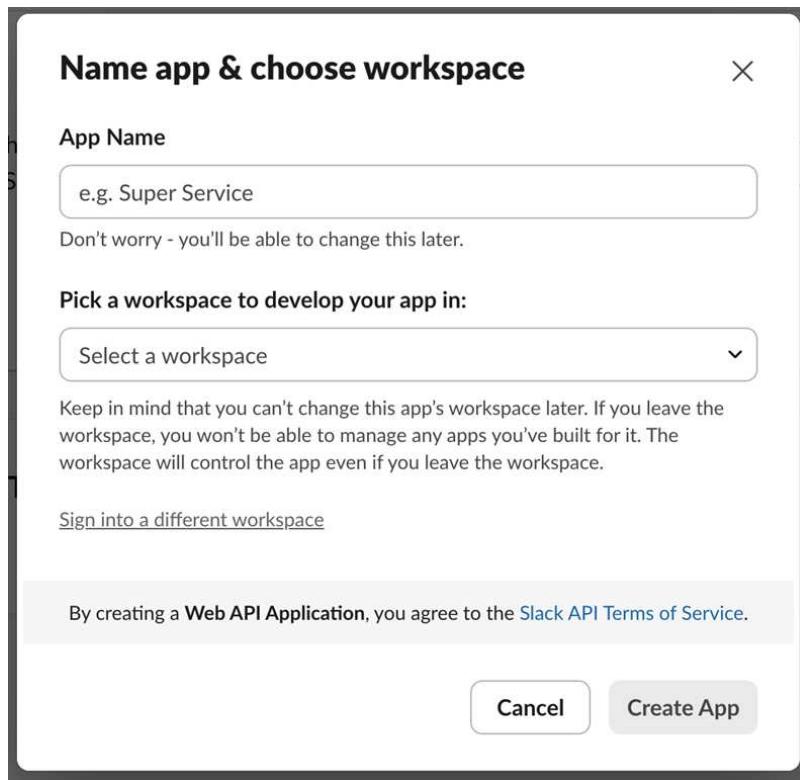
Workload Security Webhook Example for Slack

Webhooks allow users to send alert notifications to various applications using a customized webhook channel. This page provides an example for setting up webhooks for Slack.

This page refers to third-party instructions, which are subject to change. Refer to the Slack documentation for the most up-to-date information.

Slack Example

- Go to <https://api.slack.com/apps> and Create a new App. Give it a meaningful name and select a Workspace.



- Go to Incoming Webhooks, click on *Activate Incoming Webhooks*, select *Add New Webhook*, and select the Channel on which to Post.
- Copy the Webhook URL. This URL will be given when creating a Workload Security webhook.

Create Workload Security Slack Webhook

1. Navigate to Admin > Notifications and select the *Workload Security Webhooks* tab. Select *+ Webhook* to create a new webhook.
2. Give the webhook a meaningful name.
3. In the *Template Type* dropdown, select *Slack*.
4. Paste the URL copied from above.

Add a Webhook

Name
Test-Webhook-1

Template Type
Slack

URL 
https://hooks.slack.com/services/<id>

Validate SSL Certificate for secure communication

Method
POST

Custom Header
Content-type: application/json
Accept: application/json

Message Body
{
 "blocks": [
 {
 "type": "section",
 "text": {
 "type": "mrkdwn",
 "text": "*%severity%% Alert: %%synopsis%%*"
 }
 },
 {
 "type": "divider"
 }
]
}

Buttons
Cancel [Test Webhook](#) [Create Webhook](#)

Notifications via webhook

- To notify on events via webhook, navigate to *Workload Security > Policies*. Click on *+Attack Policy* or *+Warning Policy*.
- Enter a meaningful policy name.
- Select required Attack Type(s), Devices to which the policy should be attached, and required Actions.
- Under the *Webhooks Notifications* dropdown, select the required webhooks. Save the policy.

Note: Webhooks can also be attached to existing policies by editing them.

Add Attack Policy

Policy Name*
Test policy 1

For Attack Type(s) *

Ransomware Attack
 Data Destruction - File Deletion

On Device

All Devices

+ Another Device

Actions

Take Snapshot [?](#)
 Block User File Access [?](#)

Time Period

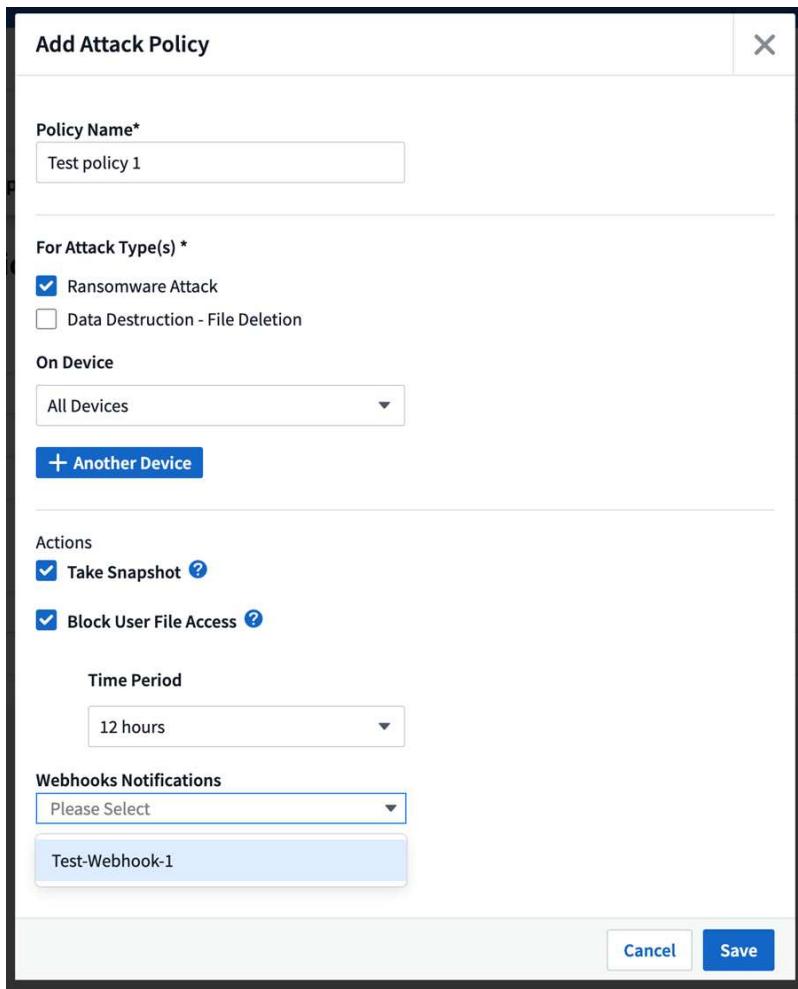
12 hours

Webhooks Notifications

Please Select

Test-Webhook-1

Cancel **Save**



Workload Security Webhook Example for Microsoft Teams

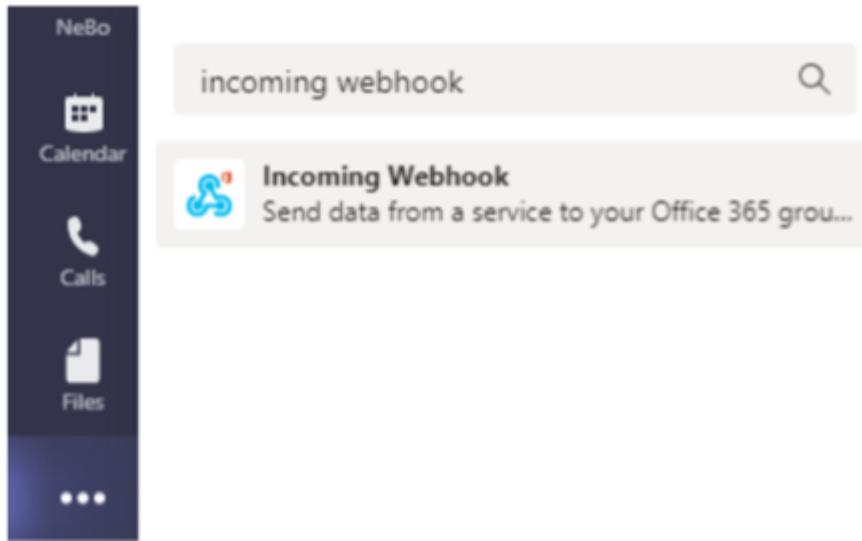
Webhooks allow users to send alert notifications to various applications using a customized webhook channel. This page provides an example for setting up webhooks for Teams.



This page refers to third-party instructions, which are subject to change. Refer to the [Teams documentation](#) for the most up-to-date information.

Teams Setup:

1. In Teams, select the kebab, and search for Incoming Webhook.



2. Select **Add to a Team > Select a Team > Setup a Connector**.
3. Copy the Webhook URL. You will need to paste this into the Workload Security webhook configuration.

Create Workload Security Teams Webhook:

1. Navigate to Admin > Notifications and select the “Workload Security Webhooks” tab. Select **+ Webhook** to create a new webhook.
2. Give the webhook a meaningful Name.
3. In the *Template Type* drop-down, select **Teams**.

Add a Webhook

Name

Template Type

URL

Validate SSL Certificate for secure communication

Method

Custom Header

```
Content-Type: application/json
Accept: application/json
```

Message Body

```
{
  "@type": "MessageCard",
  "@context": "http://schema.org/extensions",
  "themeColor": "0076D7",
  "summary": "%%severity%% Alert: %%synopsis%%",
  "sections": [
    {
      "activityTitle": "%%severity%% Alert: %%synopsis%%",
      "activitySubtitle": "%%detected%%",
      "markdown": false,
      "facts": [
        {
          "name": "Severity",
          "value": "%%severity%%"
        },
        {
          "name": "Detected At",
          "value": "%%detected%%"
        }
      ]
    }
  ]
}
```

4. Paste the URL from above into the *URL* field.

Steps for creating Teams notification with Adaptive Card template

1. Replace the message body with the following template:

```
{
  "type": "message",
  "attachments": [
    {
      "contentType": "application/vnd.microsoft.card.adaptive",
      "content": {
        "type": "AdaptiveCard",
        "version": "1.0",
        "body": [
          {
            "type": "Text",
            "text": "%%severity%% Alert: %%synopsis%%"
          }
        ],
        "actions": [
          {
            "type": "ActionCard",
            "actions": [
              {
                "type": "OpenUrlAction",
                "url": "%%detected%%"
              }
            ]
          }
        ]
      }
    }
  ]
}
```

```
{
  "contentType": "application/vnd.microsoft.card.adaptive",
  "content": {
    "$schema": "http://adaptivecards.io/schemas/adaptive-card.json",
    "type": "AdaptiveCard",
    "version": "1.2",
    "body": [
      {
        "type": "TextBlock",
        "text": "%%severity%% Alert: %%synopsis%%",
        "wrap": true,
        "weight": "Bolder",
        "size": "Large"
      },
      {
        "type": "TextBlock",
        "text": "%%detected%%",
        "wrap": true,
        "isSubtle": true,
        "spacing": "Small"
      },
      {
        "type": "FactSet",
        "facts": [
          {
            "title": "User",
            "value": "%%userName%%"
          },
          {
            "title": "Attack/Abnormal Behavior",
            "value": "%%type%%"
          },
          {
            "title": "Action taken",
            "value": "%%actionTaken%%"
          },
          {
            "title": "Files encrypted",
            "value": "%%filesEncrypted%%"
          },
          {
            "title": "Encrypted files suffix",
            "value": "%%encryptedFilesSuffix%%"
          },
          {
            "title": "Files deleted",
            "value": "%%filesDeleted%%"
          }
        ]
      }
    ]
  }
}
```

```
        "value": "%%filesDeleted%%"
    },
    {
        "title": "Activity Change Rate",
        "value": "%%changePercentage%%"
    },
    {
        "title": "Severity",
        "value": "%%severity%%"
    },
    {
        "title": "Status",
        "value": "%%status%%"
    },
    {
        "title": "Notes",
        "value": "%%note%%"
    }
]
}
],
"actions": [
{
    "type": "Action.OpenUrl",
    "title": "View Details",
    "url": "https://%%cloudInsightsHostname%%/%%alertDetailsPageUrl%%"
}
]
}
]
```

2. If you are using Power Automate Flows, query params in the URL are in encoded format. You must decode the URL before entering.
3. Click "Test Webhook" to ensure there are no errors.
4. Save the webhook.

Notifications via Webhook

To notify on events via webhook, navigate to *Workload Security > Policies*. Select *+Attack Policy* or *+Warning Policy*.

- Enter a meaningful policy name.
- Select required Attack Type(s), Devices to which policy should be attached, and required Actions.

- Under the *Webhooks Notifications* dropdown, select the required Teams webhooks. Save the policy.

Note: Webhooks can also be attached to existing policies by editing them.

Add Attack Policy

Policy Name*
Test policy 1

For Attack Type(s) *
 Ransomware Attack
 Data Destruction - File Deletion

On Device
All Devices

Actions
 Take Snapshot ?
 Block User File Access ?

Time Period
12 hours

Webhooks Notifications
Please Select
Test-Webhook-1

Cancel **Save**

Workload Security API

Integrate Workload Security with your enterprise ecosystem using a REST API protected by secure token-based authentication. Retrieve forensic activity data, manage API access tokens, and develop custom integrations with CMDBs, ticketing systems, and other applications. Interactive Swagger documentation provides complete API specifications and enables you to test endpoints directly.

Requirements for API Access:

- An API Access Token model is used to grant access.
- API Token management is performed by Workload Security users with the Administrator role.

API Documentation (Swagger)

The latest API information is found by logging in to Workload Security and navigating to **Admin > API Access**. Click the **API Documentation** link.

The API Documentation is Swagger-based, which provides a brief description and usage information for the API and allows you to try it out on your tenant.

 If calling the Forensics Activity API, use the `cloudsecure_forensics.activities.v2` API. If you are making multiple calls to this API, ensure that the calls occur sequentially, not in parallel. Multiple parallel calls may cause the API to time out.

API Access Tokens

Before using the Workload Security API, you must create one or more **API Access Tokens**. Access tokens grant read permissions. You can also set the expiration for each access token.

To create an Access Token:

- Click **Admin > API Access**
- Click **+API Access Token**
- Enter **Token Name**
- Specify **Token Expiration**

 Your token will only be available for copying to the clipboard and saving during the creation process. Tokens can not be retrieved after they are created, so it is highly recommended to copy the token and save it in a secure location. You will be prompted to click the Copy API Access Token button before you can close the token creation screen.

You can disable, enable, and revoke tokens. Tokens that are disabled can be enabled.

Tokens grant general purpose access to APIs from a customer perspective, managing access to APIs in the scope of their own tenant.

The application receives an Access Token after a user successfully authenticates and authorizes access, then passes the Access Token as a credential when it calls the target API. The passed token informs the API that the bearer of the token has been authorized to access the API and perform specific actions based on the scope that was granted during authorization.

The HTTP header where the Access Token is passed is **X-CloudInsights-ApiKey**:

For example, use the following to retrieve storages assets:

```
curl https://<Workload Security tenant>/rest/v1/cloudsecure/activities -H  
'X-CloudInsights-ApiKey: <API_Access_Token>'
```

Where *<API_Access_Token>* is the token you saved during API access key creation and *<Workload Security Tenant>* is the tenant URL of your Workload Security environment.

Detailed information can be found in the *API Documentation* link under **Admin > API Access**.

Script to extract data via the API

Workload Security agents include an export script to facilitate parallel calls to the v2 API by dividing the requested time range into smaller batches.

The script is located at */opt/netapp/cloudsecure/agent/export-script*. A README file in the same directory provides usage instructions.

Here is an example command to invoke the script:

```
python3 data-export.py --tenant_url <Workload Security tenant>  
--access_key %ACCESS_KEY% --path_filter "<dir path>" --user_name "<user>"  
--from_time "01-08-2024 00:00:00" --to_time "31-08-2024 23:59:59"  
--iteration_interval 12 --num_workers 3
```

Key Parameters:

--iteration_interval 12: Splits the requested time range into intervals of 12 hours.

--num_workers 3: Fetches these intervals in parallel using 3 threads.

Troubleshooting the ONTAP SVM Data Collector

Workload Security uses data collectors to collect file and user access data from devices. Here you can find tips for troubleshooting issues with this collector.

See the [Configuring the SVM Collector](#) page for instructions on configuring this collector.

In the case of an error, you can click on *more detail* in the *Status* column of the Installed Data Collectors page for detail about the error.

Installed Data Collectors

Name	Status	Type	Agent
9.8_vs1	! Error more detail	ONTAP SVM	agent-11

Known problems and their resolutions are described below.

Problem: Data Collector runs for some time and stops after a random time, failing with: "Error message: Connector is in error state. Service name: audit. Reason for failure: External fpolicy server overloaded."

Try This:

The event rate from ONTAP was much higher than what the Agent box can handle. Hence the connection got terminated.

Check the peak traffic in CloudSecure when the disconnection happened. This you can check from the **CloudSecure > Activity Forensics > All Activity** page.

If the peak aggregated traffic is higher than what the Agent Box can handle, then please refer to the Event Rate Checker page on how to size for Collector deployment in an Agent Box.

If the Agent was installed in the Agent box prior to 4 March 2021, run the following commands in the Agent box:

```
echo 'net.core.rmem_max=8388608' >> /etc/sysctl.conf
echo 'net.ipv4.tcp_rmem = 4096 2097152 8388608' >> /etc/sysctl.conf
sysctl -p
```

Restart the collector from the UI after resizing.

Problem: Collector reports Error Message: "No local IP address found on the connector that can reach the data interfaces of the SVM".

Try This: This is most likely due to a networking issue on the ONTAP side. Please follow these steps:

1. Ensure that there are no firewalls on the SVM data lif or the management lif which are blocking the connection from the SVM.
2. When adding an SVM via a cluster management IP, please ensure that the data lif and management lif of the SVM are pingable from the Agent VM. In case of issues, check the gateway, netmask and routes for the lif.

You can also try logging in to the cluster via ssh using the cluster management IP, and ping the Agent IP. Make sure that the agent IP is pingable:

```
network ping -vserver <vserver name> -destination <Agent IP> -lif <Lif Name> -show-detail
```

If not pingable, make sure the network settings in ONTAP are correct, so that the Agent machine is pingable.

3. If you have tried connecting via Cluster IP and it is not working, try connecting directly via SVM IP. Please see above for the steps to connect via SVM IP.
4. While adding the collector via SVM IP and vsadmin credentials, check if the SVM Lif has Data plus Mgmt

role enabled. In this case ping to the SVM Lif will work, however SSH to the SVM Lif will not work. If yes, create an SVM Mgmt Only Lif and try connecting via this SVM management only Lif.

5. If it is still not working, create a new SVM Lif and try connecting through that Lif. Make sure that the subnet mask is correctly set.

6. Advanced Debugging:

- a. Start a packet trace in ONTAP.
- b. Try to connect a data collector to the SVM from CloudSecure UI.
- c. Wait till the error appears. Stop the packet trace in ONTAP.
- d. Open the packet trace from ONTAP. It is available at this location

```
https://<cluster_mgmt_ip>/spi/<clustername>/etc/log/packet_traces/
```

e. Make sure there is a SYN from ONTAP to the Agent box.

f. If there is no SYN from ONTAP then it is an issue with firewall in ONTAP.

g. Open the firewall in ONTAP, so that ONTAP is able to connect the agent box.

7. If it is still not working, please consult the networking team to make sure that no external firewall is blocking the connection from ONTAP to the Agent box.

8. If none of the above solves the issue, open a case with [Netapp Support](#) for further assistance.

Problem: Message: "Failed to determine ONTAP type for [hostname: <IP Address>. Reason: Connection error to Storage System <IP Address>: Host is unreachable (Host unreachable)"

Try this:

1. Verify that the correct SVM IP Management address or Cluster Management IP has been provided.
2. SSH to the SVM or the Cluster to which you are intending to connect. Once you are connected ensure that the SVM or the Cluster name is correct.

Problem: Error Message: "Connector is in error state. Service.name: audit. Reason for failure: External fpolicy server terminated."

Try this:

1. It is most likely that a firewall is blocking the necessary ports in the agent machine. Verify the port range 35000-55000/tcp is opened for the agent machine to connect from the SVM. Also ensure that there are no firewalls enabled from the ONTAP side blocking communication to the agent machine.
2. Type the following command in the Agent box and ensure that the port range is open.

```
sudo iptables-save | grep 3500*
```

Sample output should look like:

```
-A IN_public_allow -p tcp -m tcp --dport 35000 -m conntrack -ctstate NEW  
-j ACCEPT
```

3. Login to SVM, enter the following commands and check that no firewall is set to block the communication with ONTAP.

```
system services firewall show  
system services firewall policy show
```

[Check firewall commands](#) on the ONTAP side.

4. SSH to the SVM/Cluster which you want to monitor. Ping the Agent box from the SVM data lif (with CIFS, NFS protocols support) and ensure that ping is working:

```
network ping -vserver <vserver name> -destination <Agent IP> -lif <Lif Name> -show-detail
```

If not pingable, make sure the network settings in ONTAP are correct, so that the Agent machine is pingable.

5. If a single SVM is added twice added to a tenant via 2 data collectors, then this error will be shown. Delete one of the data collectors thru the UI. Then restart the other data collector thru the UI. Then the data collector will show "RUNNING" status and will start receiving events from SVM.

Basically, in a tenant, 1 SVM should be added only once, via 1 data collector. 1 SVM should not added twice via 2 data collectors.

6. In instances where the same SVM was added in two different Workload Security environments (tenants), the last one will always succeed. The second collector will configure fpolicy with its own IP address and kick out the first one. So the collector in the first one will stop receiving events and its "audit" service will enter into error state.

To prevent this, configure each SVM on a single environment.

7. This error may also occur if service policies are not configured correctly. With ONTAP 9.8 or later, in order to connect to the Data Source Collector, the data-fpolicy-client service is required along with the data service data-nfs, and/or data-cifs. Additionally, the data-fpolicy-client service must be associated with the data lif(s) for the monitored SVM.

Problem: No events seen in activity page.

Try this:

1. Check if ONTAP collector is in "RUNNING" state. If yes, then ensure that some cifs events are being generated on the cifs client VMs by opening some files.
2. If no activities are seen, please login to the SVM and enter the following command.

```
<SVM>event log show -source fpolicy
```

Please ensure that there are no errors related to fpolicy.

3. If no activities are seen, please login to the SVM. Enter the following command:

```
<SVM>fpolicy show
```

Check if the fpolicy policy named with prefix "cloudsecure_" has been set and status is "on". If not set, then most likely the Agent is unable to execute the commands in the SVM. Please ensure all the prerequisites as described in the beginning of the page have been followed.

Problem: SVM Data Collector is in error state and Error message is "Agent failed to connect to the collector"
Try this:

1. Most likely the Agent is overloaded and is unable to connect to the Data Source collectors.
2. Check how many Data Source collectors are connected to the Agent.
3. Also check the data flow rate in the "All Activity" page in the UI.
4. If the number of activities per second is significantly high, install another Agent and move some of the Data Source Collectors to the new Agent.

Problem: SVM Data Collector shows error message as "fpolicy.server.connectError: Node failed to establish a connection with the FPolicy server "12.195.15.146" (reason: "Select Timed out")"

Try this: Firewall is enabled in SVM/Cluster. So fpolicy engine is unable to connect to fpolicy server.
CLIs in ONTAP which can be used to get more information are:

```
event log show -source fpolicy which shows the error  
event log show -source fpolicy -fields event,action,description which  
shows more details.
```

[Check firewall commands on the ONTAP side.](#)

Problem: Error Message: "Connector is in error state. Service name:audit. Reason for failure: No valid data interface (role: data,data protocols: NFS or CIFS or both, status: up) found on the SVM."

Try this: Ensure there is an operational interface (having role as data and data protocol as CIFS/NFS).

Problem: The data collector goes into Error state and then goes into RUNNING state after some time, then back to Error again. This cycle repeats.

Try this: This typically happens in the following scenario:

1. There are multiple data collectors added.
2. The data collectors which show this kind of behavior will have 1 SVM added to these data collectors. Meaning 2 or more data collectors are connected to 1 SVM.
3. Ensure 1 data collector connects to only 1 SVM.
4. Delete the other data collectors which are connected to the same SVM.

Problem: Connector is in error state. Service name: audit. Reason for failure: Failed to configure (policy on SVM svmname. Reason: Invalid value specified for 'shares-to-include' element within 'fpolicy.policy.scope-modify: "Federal"

Try this: *The share names need to be given without any quotes. Edit the ONTAP SVM DSC configuration to correct the share names.

Include and exclude shares is not intended for a long list of share names. Use filtering by volume instead if you have a large number of shares to include or exclude.

Problem: There are existing fpolicies in the Cluster which are unused. What should be done with those prior to installation of Workload Security?

Try this: It is recommended to delete all existing unused fpolicy settings even if they are in disconnected state. Workload Security will create fpolicy with the prefix "cloudsecure_". All other unused fpolicy configurations can be deleted.

CLI command to show fpolicy list:

```
fpolicy show
```

Steps to delete fpolicy configurations:

```
fpolicy disable -vserver <svmname> -policy-name <policy_name>
fpolicy policy scope delete -vserver <svmname> -policy-name <policy_name>
fpolicy policy delete -vserver <svmname> -policy-name <policy_name>
fpolicy policy event delete -vserver <svmname> -event-name <event_list>
fpolicy policy external-engine delete -vserver <svmname> -engine-name
<engine_name>
```

Problem: After enabling Workload Security, ONTAP performance is impacted: Latency becomes sporadically high, IOPs become sporadically low.

Try this: While using ONTAP with Workload Security sometimes latency issues can be seen in ONTAP. There are a number of possible reasons for this as noted in the following: [1372994](#), [1415152](#), [1438207](#), [1479704](#), [1354659](#). All of these issues are fixed in ONTAP 9.13.1 and later; it is strongly recommended to use one of these later versions.

Problem: Data Collector shows the error message:

“Error: Failed to determine the health of the collector within 2 retries, try restarting the collector again (Error Code: AGENT008)”.

Try this:

1. On the Data Collectors page, scroll to the right of the data collector giving the error and click on the 3 dots menu. Select *Edit*.
Enter the password of the data collector again.
Save the data collector by pressing on the *Save* button.
Data Collector will restart and the error should be resolved.
2. The Agent machine may not have enough CPU or RAM headroom, that is why the DSCs are failing.
Please check the number of Data Collectors which are added to the Agent in the machine.
If it is more than 20, please increase the CPU and RAM capacity of the Agent machine.
Once the CPU and RAM is increased, the DSCs will get into Initializing and then to Running state automatically.
Look into the sizing guide on [this page](#).

Problem: The Data Collector is erroring out when SVM mode is selected.

Try this: While connecting in SVM mode, If cluster management IP is used to connect instead of SVM management IP, then the connection will error out. Make sure that the correct SVM IP is used.

Problem: Data collector shows an error message when Access Denied feature is enabled:

“Connector is in error state. Service name: audit. Reason for failure: Failed to configure fpolicy on SVM test_svm. Reason: User is not authorized.”

Try this: The user might be missing the REST permissions needed for the Access Denied feature. Please follow the instructions on [this page](#) to set the permissions.

Restart the collector once the permissions are set.

Problem: Collector is in Error state with the message:

Connector is in error state. Service name: audit. Reason for failure: Failed to configure persistent store on SVM <SVM Name>. Reason: Unable to find a suitable aggregate for volume "<volumeName>" in SVM "<SVM Name>". Reason: Performance information for aggregate "<aggregateName>" is currently not available. Wait a few minutes and try the command again.

Try this: Wait a few minutes and then restart the Collector.

If you are still experiencing problems, reach out to the support links mentioned in the **Help > Support** page.

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.