COHESITY

Version 1.0

December 2019

The Spotlight on Cohesity Auditing

Leverage the Configuration and Protocol Audit Capabilities of Cohesity DataPlatform

ABSTRACT

The need to comply with an ever-increasing number of regulations is requiring customers to assess their IT environment more thoroughly and more often. This white paper elucidates the audit capabilities of Cohesity DataPlatform, including the built-in Cohesity Spotlight application. In addition, the paper explains how to use syslog forwarding to send audit logs to other applications, such as Splunk, for data analysis. Ultimately, Cohesity DataPlatform's audit capabilities allow customers to monitor their Cohesity clusters for access, changes to the documents in their files and folders structure, shares and permissions, as well as configuration changes to the clusters themselves.

Table of Contents

Figures

Tables

Introduction to Auditing Unstructured Data

The rapid growth of unstructured data, which is estimated to be 80 percent of worldwide data by 2025, coupled with an ever-increasing number of compliance and security requirements, compels auditors to be constantly aware of the data being stored in their data centers.

The unstructured data often includes sensitive information such as intellectual property, confidential customer or employee data, and proprietary company records. The need to audit unstructured data to keep confidential company information secure, as well as the need to comply with governmental regulations, drives the need for advanced audit capabilities.

Cohesity DataPlatform's built-in auditing makes it invaluable to private and public companies, as well as government agencies, that need to adhere to an ever-changing set of standards and regulatory mandates.

Table 1: Cohesity DataPlatform Auditing Addresses Key Business Drivers in Many Segments

SEGMENT	KEY BUSINESS DRIVERS
Financial Services	Compliance requirements for: ▪ Sarbanes-Oxley Act (SOX) ▪ Payment Card Industry Data Security Standard (PCI DSS) ▪ Federal Financial Institutions Examination Council (FFIEC)
Health Care	Compliance requirements for the Health Insurance Portability and Accountability Act (HIPAA)
Federal Agencies	Security requirements for the Security Technical Information Guide (STIG)/Federal Information Security Management Act (FISMA)
Law Enforcement	Criminal Justice Information Services (CJIS)
European Union	▪ Compliance with the General Data Protection Regulation (GDPR) ▪ Protecting Personal Identifiable Information (PII)

Audit Events in Your Data Protection Infrastructure

Cohesity DataPlatform provides the ability to audit events across the system, as well as SMB and NFS protocol access events.

For every audit event, the following information is logged:

- Timestamp
- Protocol
- Client Ip
- Username/User ID
- Domain name
- Share name
- Entity Name(s)/Entity ID(s)
- Entity attributes

Table 2 below describes the cluster and protocol events that are logged.

Table 2: Cluster and Protocol Events Captured by Cohesity Auditing

EVENT	DESCRIPTION	EVENT TYPE
Log On	A user logged in to Cohesity DataPlatform.	Cluster
Log Off	A user has logged out of Cohesity DataPlatform.	Cluster
Mount	Call to mount a view using NFS.	Protocol
Open	Call to open a SMB file.	Protocol
Close	Call to close an open SMB file.	Protocol
Create	Call to create an entity (such as a Protection Policy).	Cluster
Delete	Call to delete an entity (such as a view).	Cluster
Rename	Call to rename an entity (such as a Storage Domain).	Cluster
Set Attributes	Call to set permissions for files and folders.	Protocol

Enable File Services Audit Logs

Before you can collect and analyze information about protocol events using either NFS and/or SMB, you need to enable audit logs for file services.

To enable file services audit logs:

1. Start the Cohesity DataPlatform command-line interface (CLI), remotely or locally, as described in the *Cohesity DataPlatform CLI Reference Guide*.
2. Enter the IP address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use:

```
[cohesity@node-1 ~]$ ./iris_cli -server 192.168.100.100 -username=admin -
password=admin
```

3. Enter the following command:

```
[cohesity@node-1 ~]$ admin>cluster edit filer-audit-enabled=true filer-
audit-retention-days=<NumberOfDays>
```

   Note that `filer-audit-retention-days=`***<NumberOfDays>*** is optional. The default file audit retention value is 90 days.

4. To set the time zone for logs, enter:

```
[cohesity@node-1 ~]$ admin>cluster edit time-zone=<Timezone>
```

   To set the time zone in the Cohesity DataPlatform browser interface, see Appendix B: Time Zones.

5. By default, audit logs are disabled for all views. You can enable logs for each view individually by entering:

```
[cohesity@node-1 ~]$ admin>view edit name=<ViewName> filer-audit-
enabled=true
```

Audit log entries are created every time you create and delete a file. Audit logs are saved in the "`filesystem_audit folder/directory`" location and are written in JSON to the "`auditlog-`*YYYY-MM-DD*`.txt`" file.

Access Cohesity Audit Logs

To review cluster and protocol events, use SMB or NFS to connect to a hidden internal view called '`COHESITY_AUDIT_VIEW`'.

For SMB, use:

```
\\<cluster_name>\COHESITY_AUDIT_VIEW\filesystem_audit
```

For NFS, use:

```
linux_client# mount <cluster_name>:/COHESITY_AUDIT_VIEW/filesystem_audit
/mnt/<path>
```

Use Syslog Forwarding to Access Data in Third-party Tools

Syslog forwarding allows you to use your favorite tools to access and analyze your data and gain invaluable insights. Cohesity DataPlatform supports forwarding audit logs to a centralized syslog server or to third-party applications, such as Splunk.

Figure 1: Syslog Forwarding to Third-party Applications

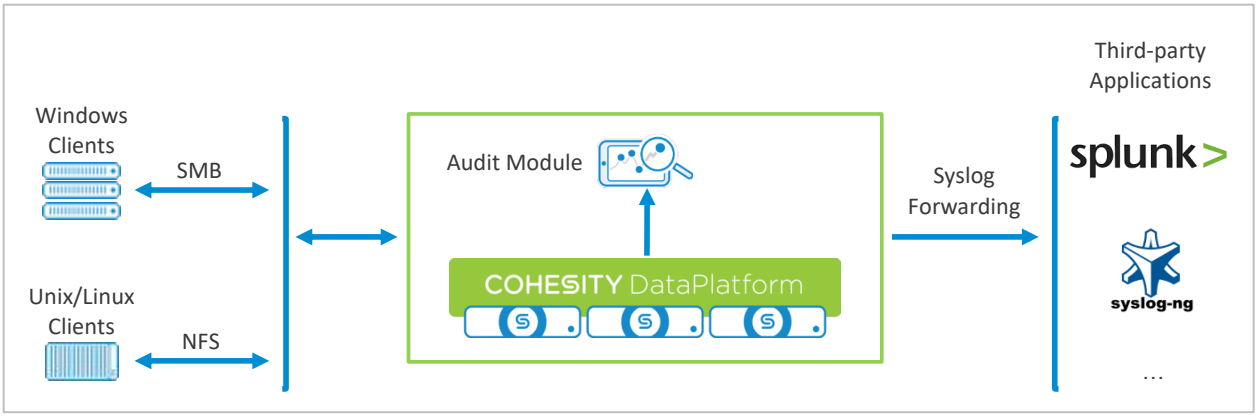

Configure Syslog Forwarding

To forward your syslogs to third-party applications, you'll need their IP address. To configure syslog forwarding with the third-party IP address, you'll need to use the Cohesity DataPlatform CLI.

To configure syslog forwarding:

1. Start the Cohesity DataPlatform CLI, remotely or locally, as described in the *Cohesity DataPlatform CLI Reference Guide*.
2. Enter:

```
[cohesity@node-1 ~]$ admin> syslog-server add address=<IP_address> filer-
audit=true port=514 protocol=tcp
```

At this point, you have successfully configured Cohesity DataPlatform to send audit events via syslog forwarding. In a later section, you'll find an example of configuring Splunk Enterprise to accept syslog events from Cohesity DataPlatform.

Use the Cohesity Spotlight App to Track File Events

Cohesity Spotlight is one of many useful apps in the Cohesity MarketPlace that extend the capabilities of Cohesity DataPlatform. With the Spotlight app, you can search audit logs to determine who is creating, modifying, accessing, or deleting files. Spotlight also flags anomalous users and entities. In doing all this, Spotlight helps customers understand their data in three important ways:

1. **Awareness**. See what's happening with your data.
2. **Business and Compliance Needs**. Gain insights into the data access.
3. **Monitoring**. Discover anomalous file access events.

Figure 2: The Cohesity Spotlight App Gives You Insights on Your Audit Log

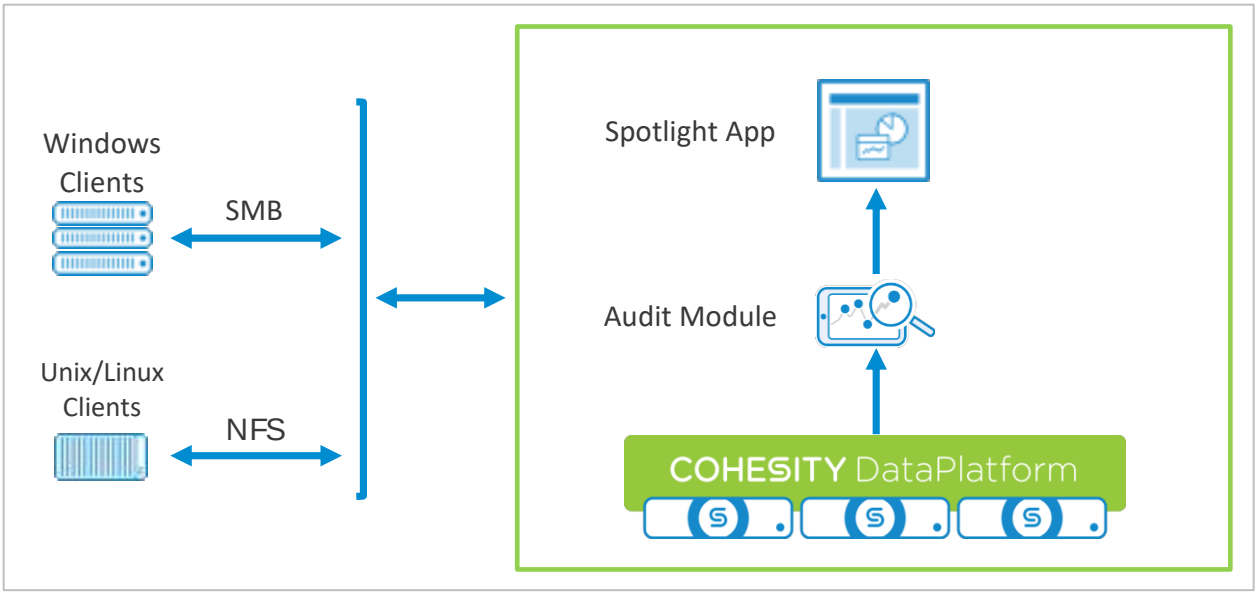

Customers can use Cohesity Spotlight to monitor any modifications to file data, as Spotlight is able to read the audit logs generated on Cohesity DataPlatform directly. For example, this application can determine, through anomaly detection, if there might have been a potential internal or external security breach, like a ransomware attack. As a result, Cohesity Spotlight gives customers much greater visibility and control over their enterprise data.

Figure 3: The Cohesity Spotlight App Shows What's Happening with Your Files

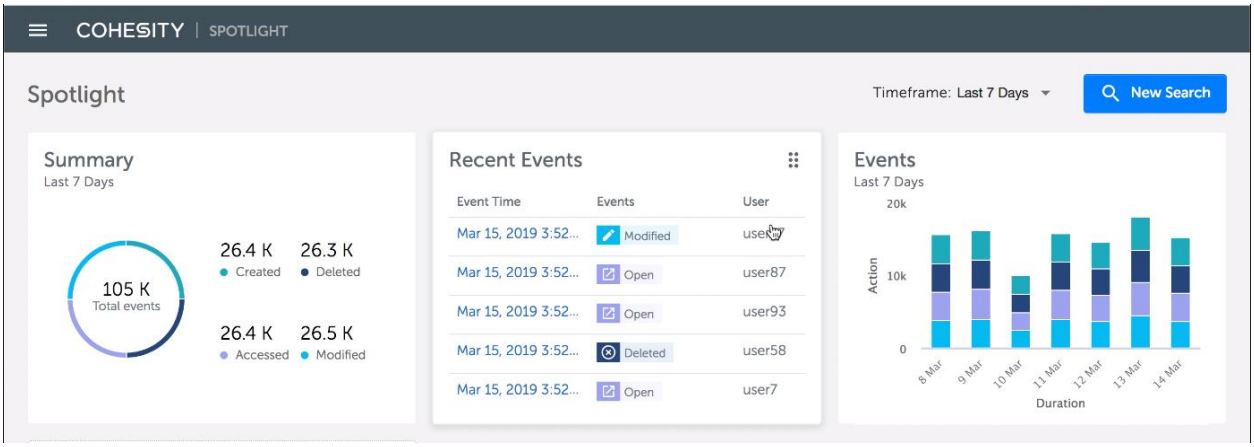

Use Splunk Enterprise to Visualize Your Data

Splunk is a software platform that helps organizations turn their data into the answers they need to solve their toughest IT, security, and business challenges. With Splunk Enterprise, you can collect, index, search, analyze, and visualize the machine-generated data gathered from the websites, applications, sensors, and devices that make up your IT infrastructure and business.

Configure Splunk

Splunk Enterprise can index remote data from syslog-ng or any other application that transmits via TCP. As the recommended protocol for sending data from a remote host to your Splunk Enterprise server, TCP is the protocol that underlies the Splunk Enterprise data distribution scheme. You can configure Splunk Enterprise to consume any data that arrives on TCP ports, and use this method to capture data from network services such as syslog.

After you download and install the software, you must start Splunk Enterprise and launch Splunk Web. Splunk Web is the browser interface to Splunk Enterprise.

Use Splunk Web to Add Network Input

You can use Splunk Web to add Cohesity DataPlatform as an input.

To add inputs from network ports using Splunk Web:

1. Log in to your Splunk Enterprise account.

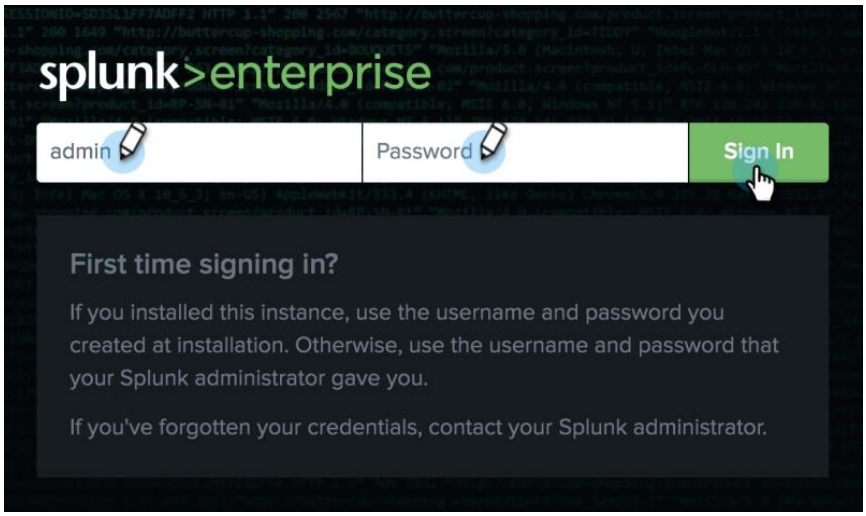

NOTE: By default, Splunk Web runs on port 8000 of the host on which it is installed. If you use Splunk Enterprise on your local machine, the URL to access Splunk Web is http://localhost:8000.

2. Navigate to **Settings** > **Data** > **Data inputs**.

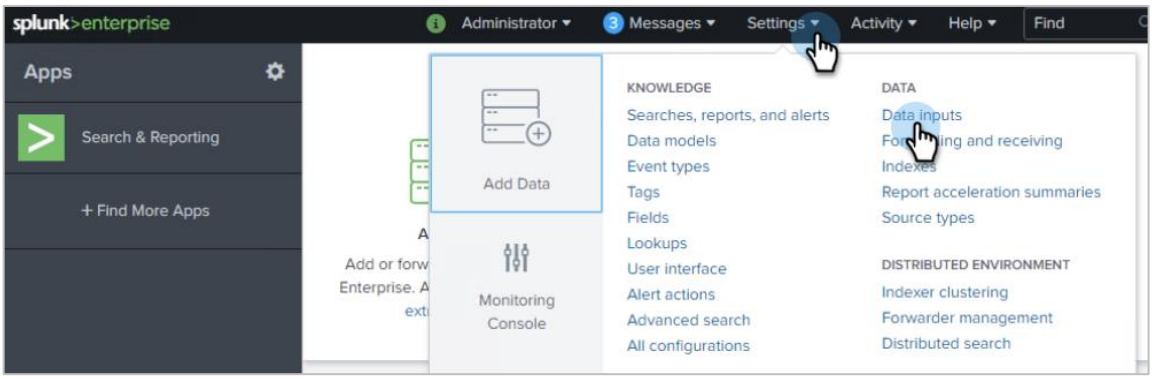

3. In the row for TCP, click **+ Add new**.

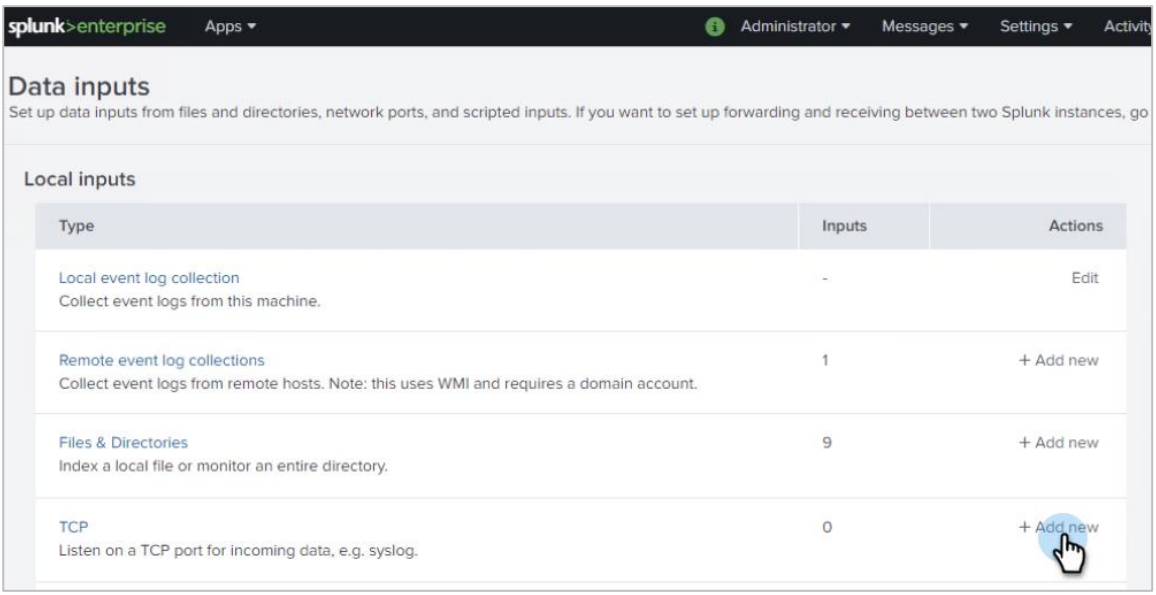

4. In the **Select Source** page, enter a **Port** number (the default port for syslog is 514) and, if necessary, enter a new **Source name override** value to override the default source value. If this is a TCP input, specify whether this port should accept connections from all hosts or only one host in the **Only accept connections from** field. If you only want the input to accept connections from one host, enter the hostname or IP address of the host. You can use wildcards to specify hosts. Click **Next**.

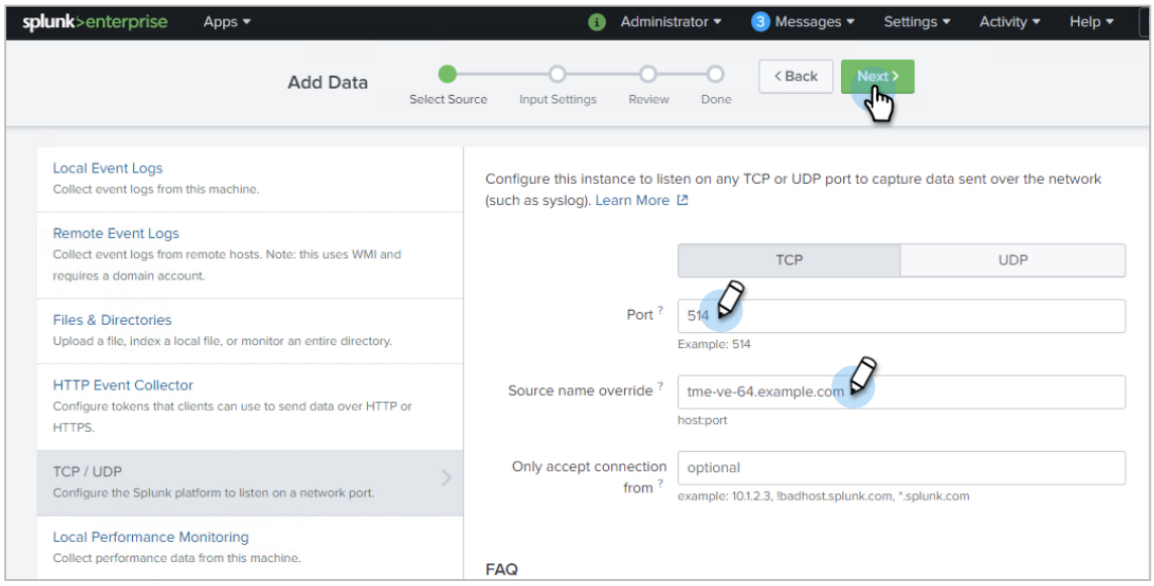

5. All of the parameters in the **Input Settings** page are optional, but here you can set the **Source type** to **syslog** and choose a **Host Method**:

   a. **IP**. Sets the input processor to rewrite the host with the IP address of the remote server.
   b. **DNS**. Sets the host to the DNS entry of the remote server.
   c. **Custom**. Sets the host to a user-defined label.

Once completed, click **Review**.

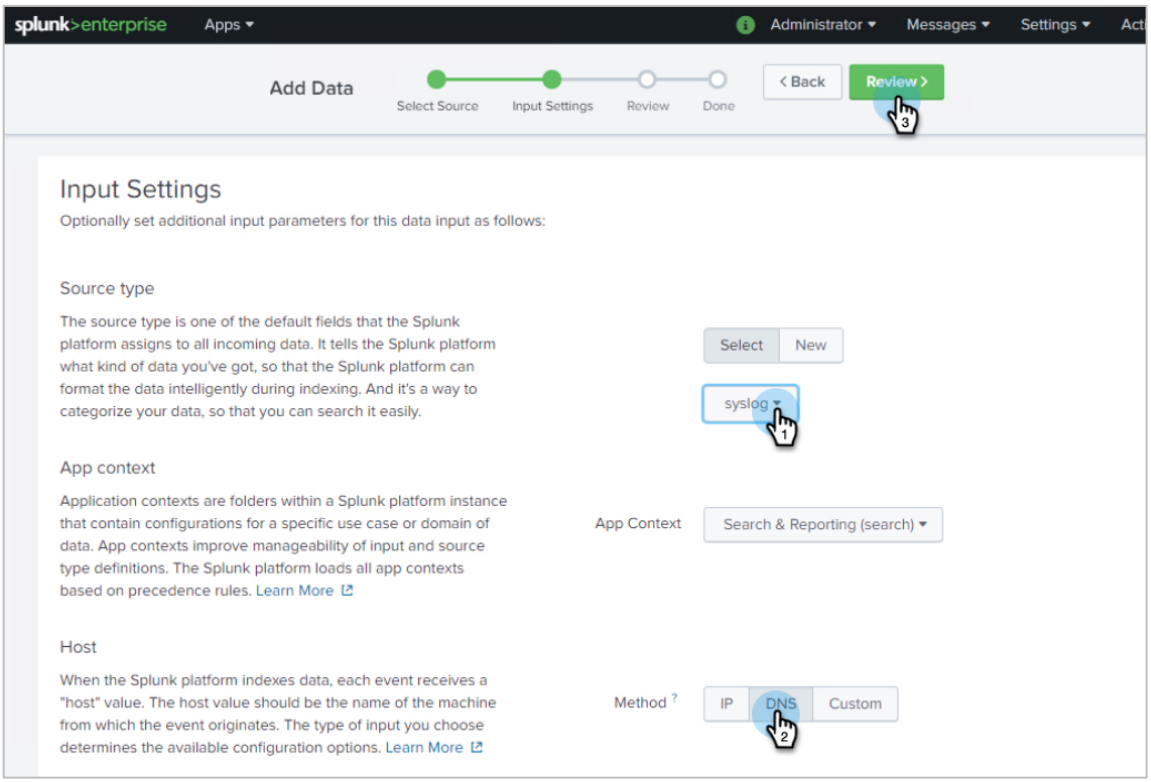

6. Review the settings to confirm they're correct and click **Submit**.

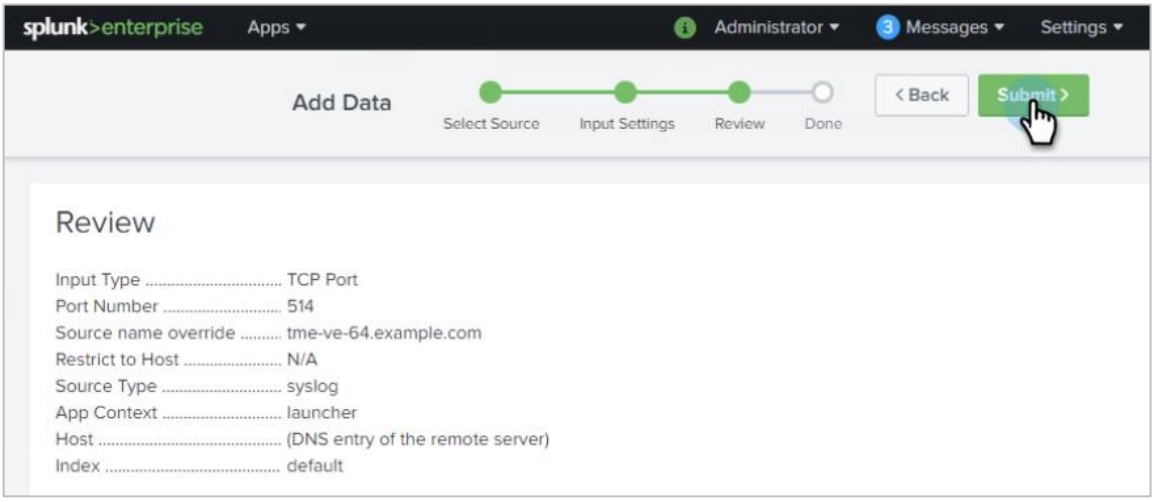

At this point, Splunk is configured to accept audit events from Cohesity DataPlatform. As events are generated, the events will be searchable via Splunk

Figure 4: Search Cohesity DataPlatform Events in Splunk Enterprise

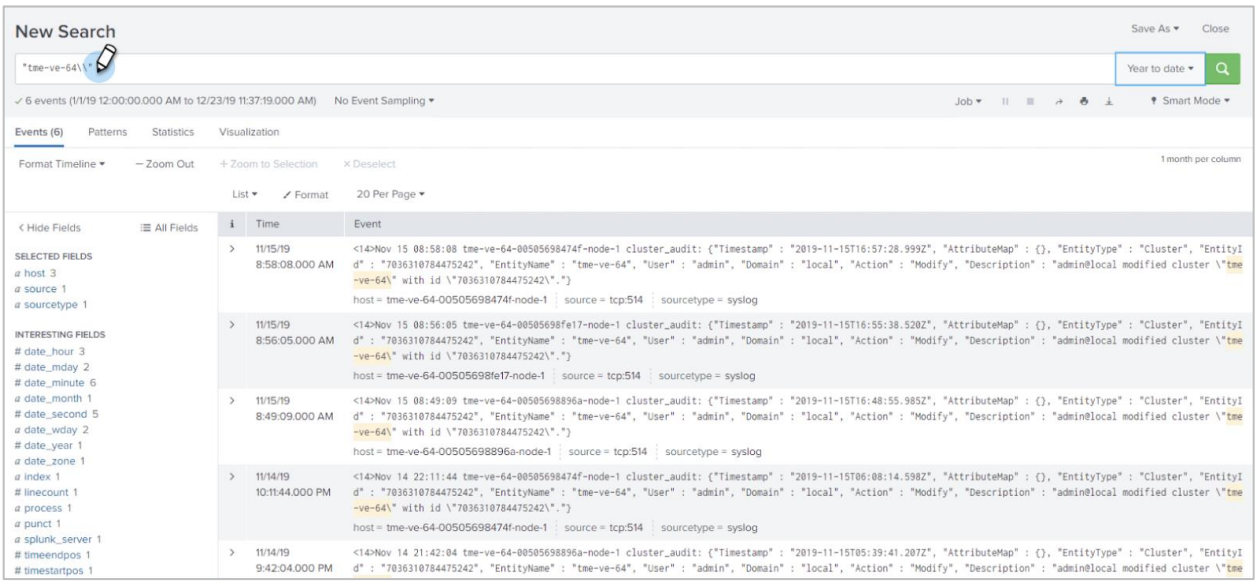

Appendix A: Example Audit Logs

The following is a sample of an SMB protocol audit log capturing the connection to an SMB share.

```
<14>Nov 17 07:15:12 tme-ve-64-00505698474f-node-1 filesystem_audit: {
"Timestamp" : "2019-11-17T15:13:36.802Z",
"RecordID" : "eccc:3",
"Protocol" : "SMB",
"ClientIP" : "10.2.165.230",
"UserName" : "Administrator",
"UserID" : 0,
"DomainName" : "EXAMPLE",
"RequestType" : "Mount",
"EntityPath" : "\\\\10.2.171.79\\nas"}
```

The above audit log shows an example of an administrator user using the SMB protocol to mount a share named '**nas**' by using one of the cluster's IP addresses.

Appendix B: Time Zones

To ensure that your audit logs include accurate timestamps for the captured events, you need to set your time zone in Cohesity DataPlatform explicitly.

To select your time zone:

1. Log in to Cohesity DataPlatform and select **Settings** > **Cluster** > **Summary**.

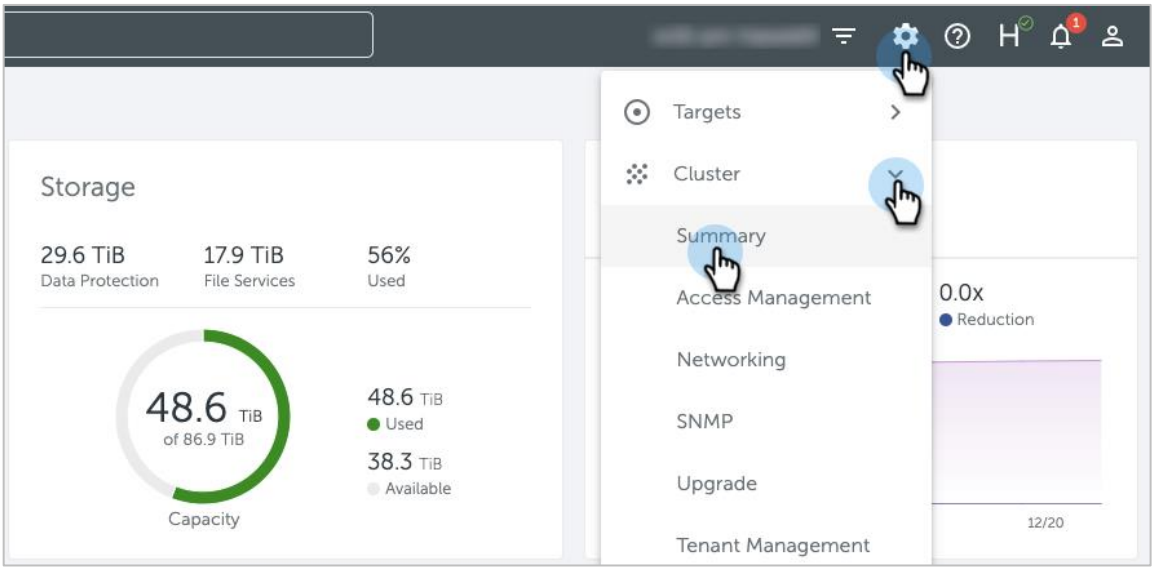

2. On the **Cluster Summary** page, click **Configure**.

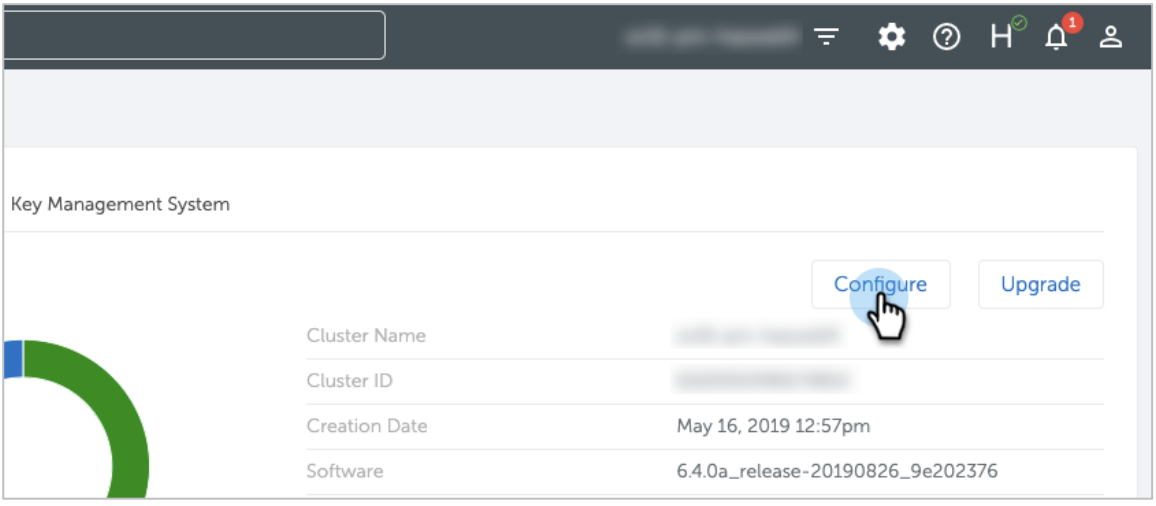

3. In the **Edit Cluster Settings** page, scroll to the bottom of the right column and click **Change Time Zone**.

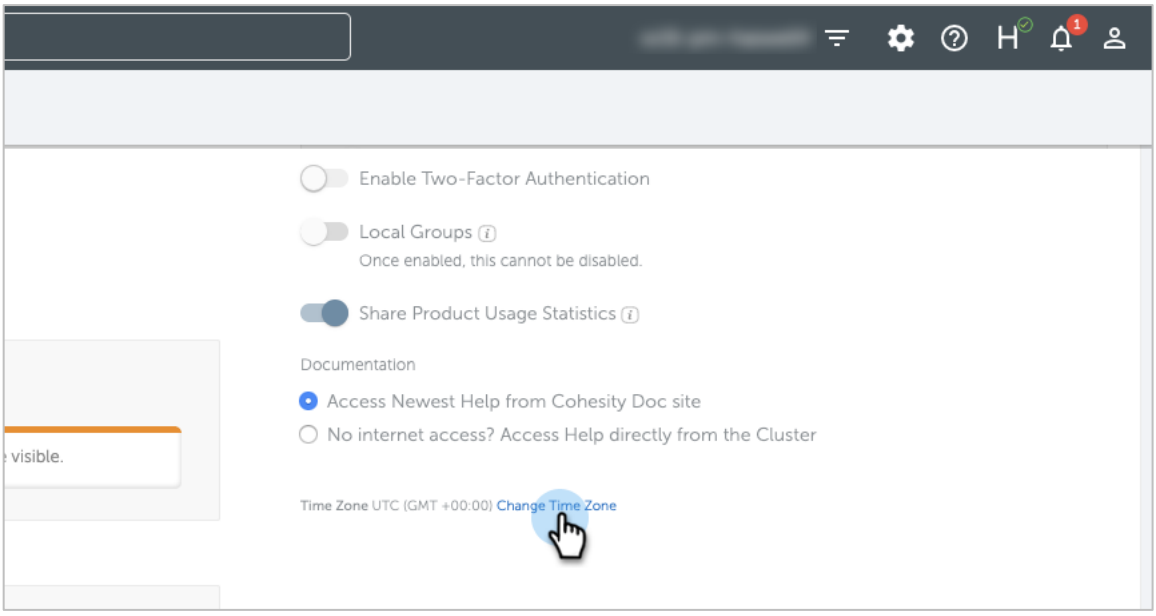

4. Start typing to see the list of available time zones and pick your time zone from the list.

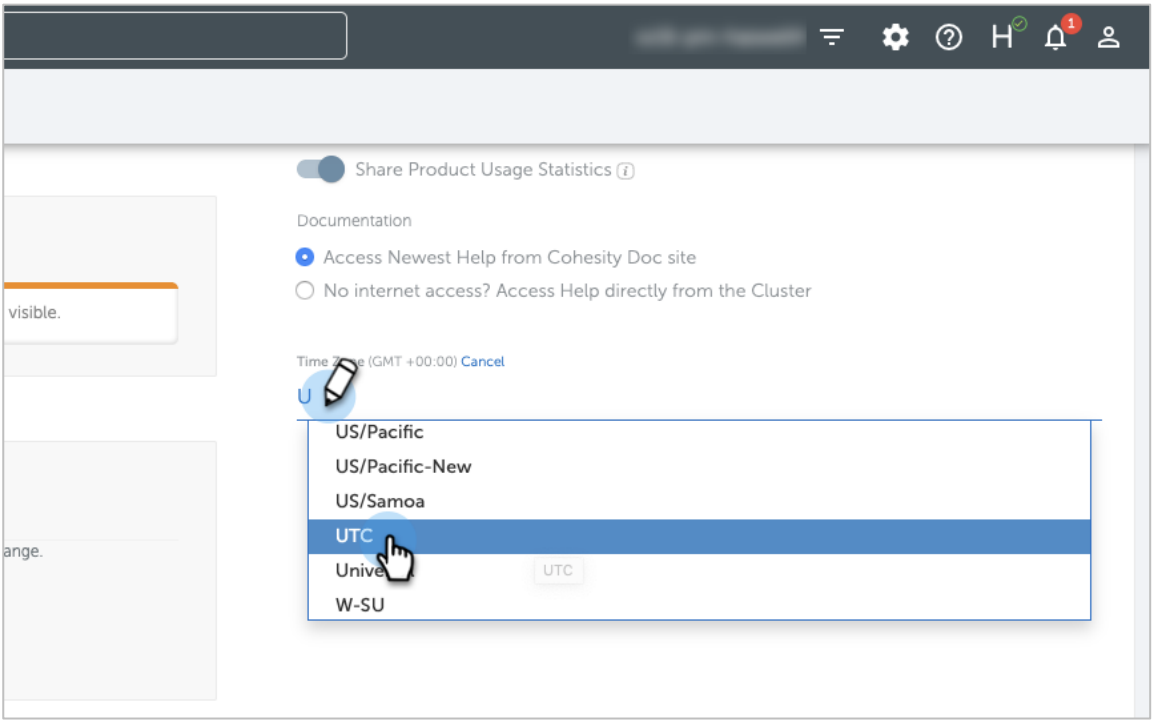

5. Scroll to the bottom of the **Edit Cluster Settings** page and click **Save**.

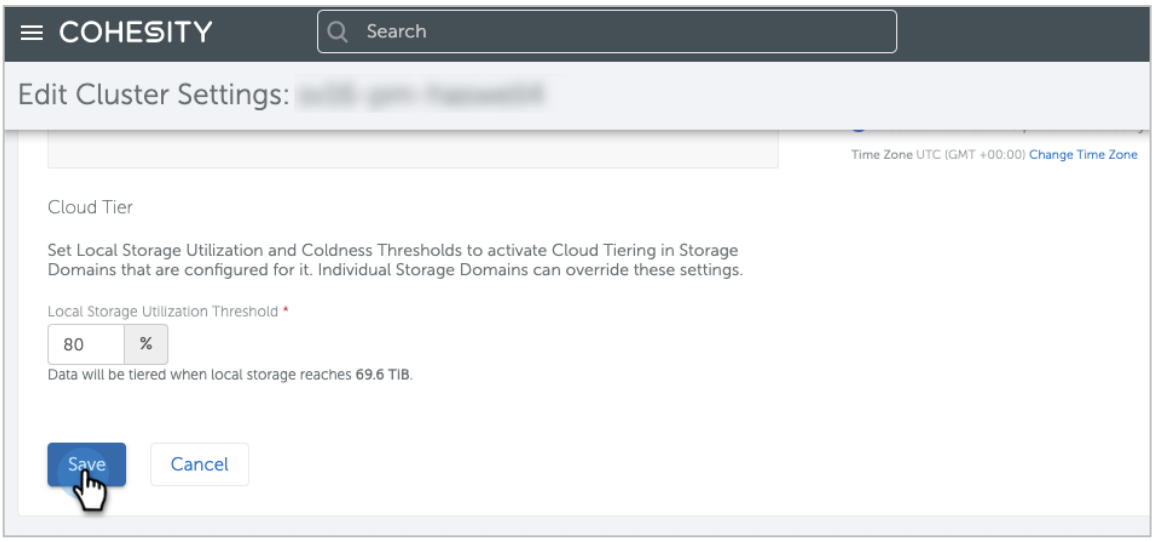

Your Cohesity cluster is now set to the time zone you selected.

Your Feedback

Was this document helpful? Send us your feedback!

About the Author

Scott Owens is a Technical Marketing Engineer at Cohesity. In his role, Scott focuses on file services.

Other essential contributors included:

- Vibhor Gupta, Product Management File Services

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.0	Dec 2019	First full release

ABOUT COHESITY

Cohesity ushers in a new era in data management that solves a critical challenge facing businesses today: mass data fragmentation. The vast majority of enterprise data — backups, archives, file shares, object stores, and data used for test/dev and analytics — sits in fragmented infrastructure silos that makes it hard to protect, expensive to manage, and difficult to analyze. Cohesity consolidates silos onto one web-scale platform, spanning on-premises, cloud, and the edge, and uniquely empowers organizations to run apps on that platform — making it easier than ever to back up and extract insights from data. Cohesity is a 2019 CNBC Disruptor and was named a Technology Pioneer by the World Economic Forum.

Visit our website and blog, follow us on Twitter and LinkedIn and like us on Facebook.

© 2019. Cohesity, Inc.

Cohesity, the Cohesity logo, SnapFS, SnapTree, SpanFS, and SpanOS, are registered trademarks, and DataPlatform, DataProtect, and Helios are trademarks of Cohesity, Inc. All rights reserved.

2000026-001-EN