Pulse Secure

Pulse Secure Services Director Splunk Deployment Guide

Supporting Pulse Secure Services Director 20.1

Product Release	**20.1**
Published	**15 April 2020**
Document Version	**1.0**

Pulse Secure, LLC
2700 Zanker Road,
Suite 200 San Jose
CA 95134

https://www.pulsesecure.net

Pulse Secure Services Director Splunk Deployment Guide

The information in this document is current as of the date on the title page.

END USER LICENSE AGREEMENT

Contents

Preface

Document Conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Pulse Secure technical documentation.

Text Formatting Conventions

Text formatting conventions such as boldface, italic, or Courier font may be used in the flow of the text to highlight specific words or phrases.

Format	Description
bold text	Identifies command names
	Identifies keywords and operands
	Identifies the names of user-manipulated GUI elements
	Identifies text to enter at the GUI
italic text	Identifies emphasis
	Identifies variables
	Identifies document titles
`Courier Font`	Identifies command output
	Identifies command syntax examples

Command Syntax Conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
italic text	Identifies a variable.
[ ]	Syntax components displayed within square brackets are optional.
	Default responses to system prompts are enclosed in square brackets.

Convention	Description
{ **x** | **y** | **z** }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x | **y**	A vertical bar separates mutually exclusive elements.
< >	Non-printing characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, member[member...].
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.
bold text	Identifies command names, keywords, and command options.

Notes and Warnings

Note, Attention, and Caution statements might be used in this document.

Note: A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

ATTENTION
An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.

CAUTION
A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

Requesting Technical Support

Technical product support is available through the Pulse Secure Global Support Center (PSGSC). If you have a support contract, file a ticket with PSGSC.

- Product warranties—For product warranty information, visit https://support.pulsesecure.net/product-service-policies/

Self-Help Online Tools and Resources

For quick and easy problem resolution, Pulse Secure provides an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://support.pulsesecure.net
- Search for known bugs: https://support.pulsesecure.net
- Find product documentation: https://www.pulsesecure.net/techpubs
- Download the latest versions of software and review release notes: https://support.pulsesecure.net

- Open a case online in the CSC Case Management tool: https://support.pulsesecure.net
- To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://support.pulsesecure.net

For important product notices, technical articles, and to ask advice:

- Search the Pulse Secure Knowledge Center for technical bulletins and security advisories: https://kb.pulsesecure.net
- Ask questions and find solutions at the Pulse Community online forum: http://kb.pulsesecure.net

Opening a Case with PSGSC

You can open a case with PSGSC on the Web or by telephone.

- Use the Case Management tool in the PSGSC at https://support.pulsesecure.net.
- Call 1-844 751 7629 (Toll Free, US).

For international or direct-dial options in countries without toll-free numbers, see https://support.pulsesecure.net/support/support-contacts/

Reporting Documentation Issues

To report any errors or inaccuracies in Pulse Secure technical documentation, or to make suggestions for future improvement, please send your comments to: techpubs-comments@pulsesecure.net. Include a full description of your issue or suggestion and the document(s) to which it relates.

Introduction

About This Document

The aim of this document is to provide a definitive set of steps to create a working Splunk / Services Director Analytics App deployment.

This document includes some Services Director configuration processes that are also included in the *Services Director Getting Started Guide*.

Note: This guide uses fixed names where possible (for example, the Splunk system's indexes called *zxtm_logs* and *zxtm_transactions*) and default ports. These can be changed (most likely in customer deployments), but if you're unfamiliar with configuration of a Splunk system or have problems with this process, we recommend using the names and ports shown in this document.

Instructions for the installation and initial configuration of Splunk are not included. Refer to the relevant Splunk documentation for these processes.

Overview of vTM/Splunk

Clusters of Pulse Secure Virtual Traffic Manager (vTM) versions 17.2 (or later) can be configured to export live analytics data to an external Splunk analytics system. This requires the vTMs to be in the estate of Services Director version 18.1 (or later). Each vTM cluster can export data about all the traffic they process, and can also export entries from log files present on the individual cluster members.

Splunk is a business intelligence tool that allows you to collect, store, search, analyze and visualize data. The deployment guide assumes that you already have a Splunk deployment running; see https://www.splunk.com for details on how to get started with Splunk if you are not already familiar with it.

Note: There are a variety of deployment models for Splunk analytics software, so the instructions in this guide may need to be adapted to suit your own Splunk deployment.

This guide leads users through a set of steps after which they should have a Services Director VA capable of:

- Configuring some or all of its estate of vTMs to export log and transaction analytics data to an external deployment of Splunk analytics software.

- Querying that external deployment of Splunk analytics software in order to visualise and navigate the collected analytics data.

Key Concepts

Working with Splunk involves the following key concepts:

- **Analytics records**: A vTM (as of version 17.2 and onwards) is capable of exporting two types of analytics records, specifically transaction records and log records. Each transaction record deals with a connection or higher level request that has passed through a vTM. Each log record is a single log line from a vTM log. These different record types are stored in two separate indices within the Splunk system. The traffic manager exports transaction records over a TCP connection, which can optionally be secured with TLS. Log records are exported with HTTP POST requests, and can also be sent over a secure connection if necessary. Both types are exported as JSON objects, with individual records separated by newline characters.

- **Indexes**: Indexes are where a Splunk system stores the data it receives. Transaction records and log records from Virtual Traffic Manager are stored in separate indexes to allow them to be queried independently (and independently of other unrelated data stored in the Splunk system) by the Services Director analytics application.

- **Source Types**: All records stored by a Splunk system have a 'sourcetype' field, which is assigned to the record by the Splunk system input that is configured to receive the raw analytics data. The sourcetype field references a 'source type' configuration object in the Splunk system's configuration, which controls how the raw data is parsed into separate records and how information such as the timestamp of the event can be determined. Transaction records and log records from Virtual Traffic Manager have separate source types to allow transaction record and log record specific processing to occur.

Required Configuration

In order for the system to operate:

1. The Splunk system must be configured with appropriate collection endpoints to accept and store these records in the correct indices for later analysis. Specifically, the Splunk system must be configured with a raw TCP input and an HTTP Event Collector input. The raw TCP input will collect the transaction records, and the HTTP Event Collector will collect log records. Each has a different way of processing and storing the data it receives.

2. Services Director must be configured with details of these collection endpoints in order that it can configure the vTM estate to export analytics data to the Splunk system.

3. Services Director must also be configured with details of the Splunk system's search endpoint in order to allow the analytics application to direct queries at the Splunk system.

Note: Services Director must be appropriately licensed with Enterprise Management resource pack licenses. Licensing is outside the scope of this document.

Network Diagram

Architecturally speaking, the system can be considered to look as follows.
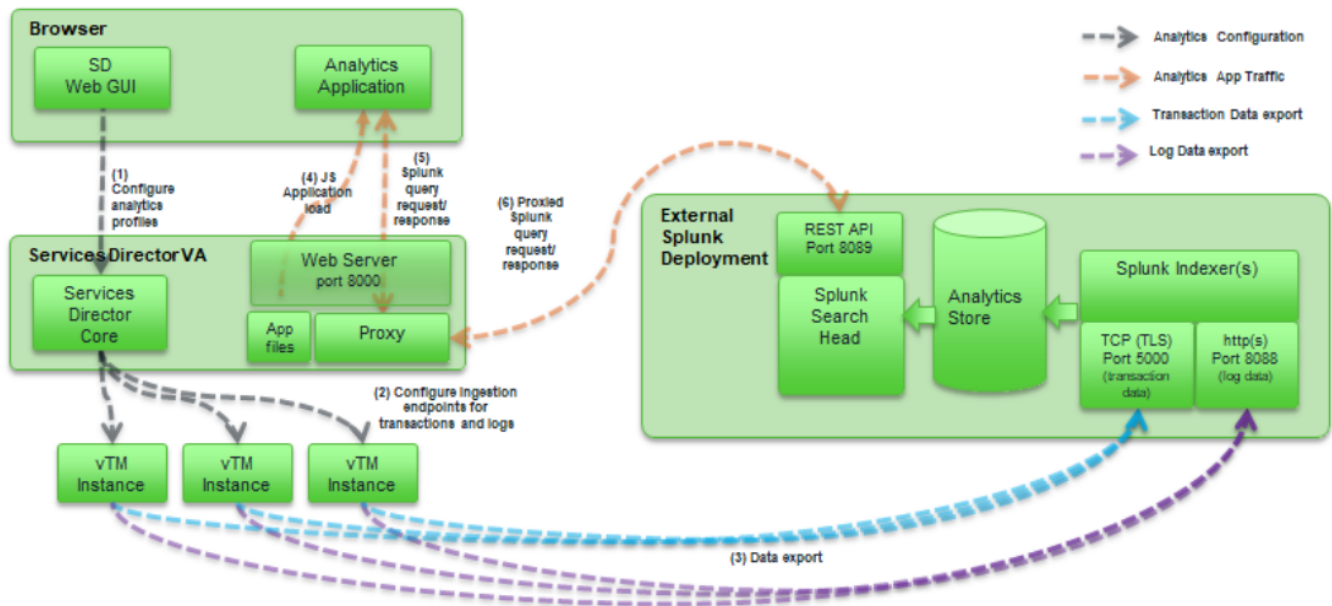
Please note the ports on which the browser, Services Director, vTM and the Splunk system communicate:

- Port 8000 (on Services Director) - Web server (serving the Analytics application itself, and acting as a proxy to Splunk.

- Port 8089 (on Splunk system) - REST API offering query capability.

- Port 5000 (on Splunk system) - As configured by instructions in this guide, a TCP port for collection of analytics transaction records.

- Port 8088 (on Splunk system) - As configured by instructions in this guide, a HTTP(s) port for collection of analytics log records.

Prerequisites

Note: The following instructions have been verified with Splunk versions 6.5.0 and 7.0.1.

You must already have the following deployed:

- A machine running a Splunk system.

- A Services Director 18.1 VA or later.

- A vTM that supports Analytics Export (version 17.3 or later).

You must also have done the following:

- Configured and started the vTM.

- Completed the Services Director Setup Wizard.

- Configured the Services Director to license the vTM so that it supports analytics export. Services Director should also show the vTM as healthy.

Accessing the Services Director VA CLI and OS Shell

The command-line instructions in this book intended for use on the Services Director VA require you to access the CLI and OS shell. For example (using the default *amnesiac* hostname):

1. Connect to the CLI. The login sequence appears. For example:

```
login as: admin
Pulse Services Director
admin@<host>'s password:
Last login: <timestamp> from <IP_address>
amnesiac >
```

2. Enable configuration mode:

```
amnesiac > enable
amnesiac # configure terminal
amnesiac (config) #
```

3. (Optionally) Start the OS shell:

```
amnesiac (config) # _shell
[admin@amnesiac ~]#
```

4. You are now in the operating system shell, and can enter shell commands.

5. To exit the OS shell, type *<ctrl> + D*

Configuring the Splunk System

Introduction

The configuration process for your Splunk system requires a number of separate stages.

Many stages can be performed via both the Splunk system's CLI or GUI, but in some cases only the GUI is supported.

In rare cases, the only way to change a setting is by editing one of Splunk system's configuration files.

Creating Indexes

Services Director requires log data and transaction data to be in separate indexes. This keeps events data organized, allows different retention policies to be set, and can speed up searches.

CLI

1. Log into the Splunk server's command line using SSH.

2. Enter the */opt/splunk* directory.

3. Run the following shell commands:

```
sudo bin/splunk add index zxtm_transactions -maxDataSize auto_high_volume
sudo bin/splunk add index zxtm_logs -maxDataSize auto_high_volume
```

GUI

1. Navigate to the Splunk system's web interface and login.

2. From the menu bar, select **Settings > Data > Indexes**.

3. On the **Indexes** page, click the **New Index** button.

4. In the **New Index** dialog, complete the following fields:

    - **Index Name**: *zxtm_transactions.*

    - **Max Size of Hot/Warm/Cold Bucket**: *auto_high_volume.*

    Leave all other fields with their default values. For example:

    FIGURE 2   New Index - zxtm_transactions

5. Click **Save**.

6. Click the **New Index** button.

7. In the **New Index** dialog, complete the fields as follows:

    - **Index Name**: *zxtm_logs*

    - **Max Size of Hot/Warm/Cold Bucket**: *auto_high_volume*

    Leave all other fields with their default values. For example:

FIGURE 3   New Index - zxtm_logs

8. Click **Save**.

Creating a Source Type for Transactions

Creating a source type allows the Splunk system to interpret transaction data correctly. This cannot be done via the CLI.

GUI

Navigate to the Splunk system's web interface and login.

1. From the menu bar, select **Settings > Data > Source types**.

2. On the **Source Types** page, click the **New Source Type** button.

3. In the **Create Source Type** dialog, complete the following fields:

   - **Name**: *zxtm_transactions*

   - **Category**: *Network & Security*

   - **Indexed Extractions**: *none*

4. In the **Event Breaks** section:

   - Set **Break Type** to *Every Line*.

5. In the **Timestamp** section:

   - Set **Extraction** to *Auto*.

6. In the **Advanced** section:

   - Click **New setting**. A new entry row appears.

   - In **Name**, type *KV_MODE*.

   - In **Value**, select *json*.

   Leave all other fields with their default values. For example:

FIGURE 4   Create Source Type

7. Click **Save**.

Creating a Data Input for Transactions

You must now configure the Splunk system to listen for transaction data on port 5000.

This can be performed using either the Splunk system's CLI or GUI.

CLI

1. Log into the Splunk server's command line using SSH.

2. Enter the */opt/splunk* directory.

3. Run the following shell command:

   ```
   sudo bin/splunk add tcp 5000 -sourcetype zxtm_transactions -index zxtm_transactions
   -resolvehost true
   ```

GUI

1. Navigate to the Splunk system's web interface and login.

2. From the menu bar, select **Settings > Data > Data inputs**.

3. On the **Data inputs** page, under **Local Inputs > TCP**, click the **Add New** action.

   The **Add Data** wizard starts.

4. In the Select Source pane of the wizard, complete the following field:

   • **Port**: *5000*.

   Leave all other fields with their default values.

5. Click the **Next >** button.

6. In the **Input Settings** pane of the wizard, complete the following fields:

   • Under Source type:

      • Click **Select**.

      • Click **Select source type** and select **Network & Security > zxtm_transactions**.

   • Under **Host**, select *DNS*.

   • Under **Index**, select *zxtm_transactions*.

   Leave all other fields with their default values. For example:

FIGURE 5   Add Data Wizard: Input Settings

7. Click the **Review >** button.

8. In the **Review** pane of the wizard, review all selections. For example:

FIGURE 6   Add Data Wizard: Review

9. Click **Submit >**.

Creating a Data Input For Logs

The following steps will configure the Splunk system to listen for log data on the default HTTP Event Collector port (8088). This can be done via the Splunk system's CLI or GUI.

CLI

1. Log into the Splunk server's command line using SSH.

2. Enter the */opt/splunk* directory.

3. Run the following shell command:

   ```
   sudo bin/splunk http-event-collector create zxtm_logs -uri https://localhost:8089 -
   index zxtm_logs
   ```

4. Make a note of the token in the output of the above command. This is referred to as *<auth-token>* in later procedures.

5. Run the following command:

   ```
   sudo bin/splunk http-event-collector enable -uri https://localhost:8089
   ```

GUI

1. Navigate to the Splunk system's web interface and login.

2. From the menu bar, select **Settings > Data > Data inputs**.

3. On the **Data inputs** page, under **Local Inputs > HTTP Event Collector**, click the **Add New** action.

   The **Add Data** wizard starts.

4. In the **Select Source** pane of the wizard, complete the following field:

   • **Name**: *zxtm_logs*

   Leave all other fields with their default values.

5. Click the **Next >** button.

6. In the **Input Settings** pane of the wizard, complete the following fields:

   • Under **Source type**, click *Automatic*.

   • Under **Index**:

     • **Select Allowed Indexes**: Ensure there are no selections.

     • **Default Index**: *zxtm_logs*.

For example:

FIGURE 7   Add Data Wizard: Input Settings

7. Click the **Review >** button.

8. In the **Review** pane of the wizard, review all selections. For example:

FIGURE 8   Add Data Wizard: Review

9. Click **Submit >**.

    A confirmation screen is displayed. For example:

    FIGURE 9   Add Data Wizard: Token

10. Make a note of the **Token Value**.

11. From the menu bar, select **Settings > Data > Data inputs**.

12. On the **Data inputs** page, click **Local Inputs > HTTP Event Collector**.

    The **HTTP Event Collector** page appears. For example:

    FIGURE 10   HTTP Event Collector

13. Click **Global Settings**.

    The **Edit Global Settings** dialog appears.

14. In the **Edit Global Settings** dialog, complete the following field:

- **All Tokens**: *Enabled*

Leave all other fields with their default values. For example:

FIGURE 11   Edit Global Settings

15. Click **Save**.

Setting Up Event Log Entry Parsing

You can configure the Splunk system to extract additional data from the event logs. Unfortunately, this cannot be done via the CLI. Additionally, for Splunk versions below 7, only some parts can be done via the GUI. Alternative instructions are provided where required in the GUI section below.

GUI

1. Navigate to the Splunk system's web interface and login.

2. From the menu bar, select **Settings > Knowledge > Fields**.

3. On the **Fields** page, under **Field extractions**, click the **Add New** action.

4. On the **Add new** page, complete the following fields:

   - **Destination app**: *search*

   - **Name**: *zxtm-event-log*

   - **Apply to**: *sourcetype*

   - **Named**: *zxtm_event_log*

   - **Type**: *Inline*

   - **Extraction/Transform**: Cut/paste the following code extract.

     ```
     \[[^\]]+\]\s+(?<severity>[^\t:]+)(\t(?<event_tags>.*))?\t(?<message>[^\t]*)
     ```

   For example:

FIGURE 12   Add New Field Extractions

5. Click **Save**. The addition is confirmed. For example:

Field Extractions

6. If your Splunk version is before 7, perform the following steps:

- From a shell prompt, edit the etc/users/admin/search/local/transforms.conf file, creating the directory and file if necessary.

- Cut/paste the following content into the file:

```
[zxtm-event-tags]
CLEAN_KEYS = 0
DELIMS = "\t"
FIELDS = tag, tag, tag, tag, tag, tag
MV_ADD = 1
SOURCE_KEY = event_tags
```

- Save and close the file.

- Run the following command:

```
sudo bin/splunk restart
```

7. If your Splunk version is 7 or above, perform the following steps:

- From the menu bar, select **Settings > Knowledge > Fields**.

- On the **Fields** page, under **Field transformations**, click the **Add New** action.

- On the **Add new** page, complete the following fields:

  - **Destination app**: *search*

  - **Name**: *zxtm-event-tags*

  - **Type**: *delimiter-based*

  - **Delimiters**: *"\t"*

  - **Field list**: *tag, tag, tag, tag, tag, tag*

  - **Source key**: *event_tags*

- **Create multivalued fields**: Select this check box.

- **Automatically clean field names**: Clear this check box.

For example:

FIGURE 14   Add New Field Transformations

- Click **Save**.

8. From the menu bar, select **Settings > Knowledge > Fields**.

9. On the **Fields** page, under **Field extractions**, click the **Add New** action.

10. On the **Add new** page, complete the following fields:

- **Destination app**: *search*

- **Name**: *zxtm-event-tags*

- **Apply to**: *sourcetype*

- **Named**: *zxtm_event_log*

- **Type**: *Uses transform*

- **Extraction/Transform**: *zxtm-event-tags*

For example:

FIGURE 15   Add New Field Extractions

11. Click **Save**.

Setting Up Audit Log Entry Parsing

You can configure the Splunk system to extract additional data from the audit logs. Unfortunately, this cannot be done via the CLI. Additionally, for Splunk versions below 7, only some parts can be done via the GUI. Alternative instructions are provided where required in the GUI section below.

GUI

1. Navigate to the Splunk system's web interface and login.

2. If your Splunk version is before 7, perform the following steps:

    -   From a shell prompt, edit the *etc/users/admin/search/local/transforms.conf* file, creating the directory and file if necessary.

    -   Cut/paste the following content into the file:

        ```
        [zxtm-audit-fields]
        CLEAN_KEYS = 0
        DELIMS = "\t", "="
        FIELDS =
        ```

    -   Save and close the file.

    -   Run the following command:

        ```
        sudo bin/splunk restart
        ```

3. If your Splunk version is 7 or above, perform the following steps:

    -   From the menu bar, select **Settings > Knowledge > Fields**.

    -   On the **Fields** page, under **Field transformations**, click the **Add New** action.

    -   On the **Add new** page, complete the following fields:

        -   **Destination app**: *search*

        -   **Name**: *zxtm-audit-fields*

        -   **Type**: *delimiter-based*

        -   **Delimiters**: *"\t", "="*

        -   **Field list**: Enter a single space.

        -   **Source key**: *_raw*.

        -   **Create multivalued fields**: Clear this check box.

        -   **Automatically clean field names**: Clear this check box.

    For example:

FIGURE 16   Add New Field Transformations

- Click **Save**.

4. On the **Fields** page, under **Field extractions**, click the **Add New** action.

5. On the **Add new** page, complete the following fields:

  - **Destination app**: *search*

  - **Name**: *zxtm-audit-logs*

  - **Apply to**: *sourcetype*

  - **Named**: *zxtm_audit_log*

  - **Type**: *Uses transform*

  - **Extraction/Transform**: *zxtm-audit-fields*

  For example:

FIGURE 17   Add New Field Extractions

6.   Click **Save**.

Setting Up Case-Insensitive Searches for HTTP Headers

HTTP header field names are treated as case-sensitive in the Splunk system. This is opposite to the general HTTP specification of HTTP header field names, which are case-insensitive.

As a result, you may want your Splunk system to search for multiple variants of the field name. For example, User-Agent, User-agent and USER-AGENT. To do this, field aliases must be added. Unfortunately, this cannot be performed using the CLI.

Note: The built-in header filters always search using title case. For example, HTTP Request Header User-Agent. If you require case insensitivity for other data exported by vTM (for example, *http.request.cookies* or *http.response.cookies*), repeat the steps below for the affected headers.

GUI

1. Navigate to the Splunk system's web interface and login.

2. On the **Fields** page, under **Field aliases**, click the **Add New** action.

3. On the **Add new** page, complete the following fields:

   - **Destination app**: *search*

   - **Name**: Enter your own choice of name for the alias. For example, *User-Agent*.

   - **Apply to**: *sourcetype*

   - **Named**: *zxtm_transactions*

4. Decide on a consistent name for all of the variants. This name will be used in searches. For example, *Consistent_Name*.

5. For each case-sensitive variant of *Consistent_Name*, create a **Field alias** entry (adding extra entries as required by clicking **Add another**).

   - For request header variants, each **Field alias** entry should take the form:

     ```
     http.request.<variant> = http.request.<Consistent_Name>
     ```

   - For response header variants, each **Field alias** entry should take the form:

     ```
     http.response.<variant> = http.response.<Consistent_Name>
     ```

   Note: You do not need to create a **Field alias** entry that exactly matches *Consistent_Name*.

   For example:

FIGURE 18   Add New Field Aliases

6. Click **Save**.

    A summary of the new alias appears. For example:

    FIGURE 19   Field Aliases

Configuring Services Director

Introduction

You configure Services Director as follows:

- First, you create a collection endpoint. This represents where data for the Splunk system should be sent.

- Then, an analytics profile needs to be creating and applied to a vTM cluster, so the cluster knows what to export and what collection endpoint to export to.

- Finally, a search endpoint needs to be created so that Services Director can get data from the Splunk system to make graphs.

This chapter describes both CLI and GUI procedures.

Creating a Collection Endpoint

This stage can be completed using the Services Director CLI or GUI. The following properties are required:

- *<collection-endpoint>* is the FQDN or IP address of your Splunk machine.

- *<auth-token>* is the token displayed at the end of the procedure for creating a data input for logs, see "Creating a Data Input For Logs" on page 16.

CLI

1. Start the CLI in configuration mode (not the shell), as described in "Accessing the Services Director VA CLI and OS Shell" on page 8.

2. Enter the following CLI command:

```
ssc collection-endpoint create name Test txn-export-address <collection-
endpoint>:5000 txn-tls false log-export-address https://<collection-
endpoint>:8088/services/collector/event log-tls-verify false auth-type splunk
auth-token <auth-token>
```

GUI

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.

2. Log in as the administration user.

   The **Home** page appears.

3. Click the **Catalogs** menu, and then click **Analytics > Analytics Endpoints**.

   The **Analytics Endpoints** page appears. For example:

   FIGURE 20   Collection Endpoints

4. Click the **Add** button above the **Collection Endpoints** table.

   The **Add Collection Endpoint** dialog box appears.

5. Complete the following fields:

   - **Name**: *Test*. This name will appear in the **Collection Endpoints** table after you apply the endpoint.

   - Under **Transaction Export Collector Settings**:

     - **Address**: *<collection-endpoint>:5000*

   - Under **Log Export Collector Settings**:

     - **Address**: *https://<collection-endpoint>:8088/services/collector/event*

     - **Authentication Method**: *Splunk*

     - **HEC Token**: *<auth-token>*

   Leave all other fields with their default values. For example:

FIGURE 21   Add Collection Endpoint

6. Click **Apply**.

    The collection endpoint is added to the **Collection Endpoint** table. For example:

    FIGURE 22   New Collection Endpoint

Creating and Applying an Analytics Profile

This stage can be completed using the Services Director CLI or GUI.

The following properties are required:

- *<logs_export_list>* is a comma-separated list of log IDs. For example:

```
"Audit Log","Event Log","System - authentication log"
```

- *<cluster-name>* is the name or ID of the target vTM cluster.

CLI

1. Start the CLI in configuration mode (not the shell), as described in "Accessing the Services Director VA CLI and OS Shell" on page 8.

2. Create an analytics profile:

```
ssc analytics-profile create logs-to-export <logs_export_list> tag Test
```

3. Apply the analytics profile to an analytics cluster:

```
ssc cluster update cluster-name <cluster_name> analytics-profile Test
```

GUI

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.

2. Log in as the administration user. The **Home** page appears.

3. Click the **Catalogs** menu, and then click **Analytics > Analytics Profiles**.

   The **Analytics Profiles** page appears. For example:

   FIGURE 23   Analytics Profiles

   ## Analytics Profiles

   ⊕ Add

Name :	ID :	Logs to export	Transaction Data Export
		No Data	

4. Click the **Add** button above the table.

5. In the **Add Analytics Profile** dialog, complete the following fields:

   - **Name**: *Test*

   - **Enable Transaction Export**: Select this check box.

   - **Logs to Export**: Check one or more of the listed logs.

For example:

FIGURE 24   Add Analytics Profile

6. Click **Apply**.

    The profile is added to the **Analytics Profile** page. For example:

    FIGURE 25   New Analytics Profile

7. Click the **Services** menu, and then click **Services Director > vTM Clusters**.

    The **vTM Clusters** page appears. For example:

    FIGURE 26   vTM Clusters

8. Expand the cluster entry corresponding to your registered vTM.

9. In the expanded view, set **Analytics Profile** to *Test*. For example:

FIGURE 27   vTM Clusters Expanded

10. Click **Apply**.

Creating a Search Endpoint

This stage can be completed using the Services Director CLI or GUI. The following properties are required:

- *<search-endpoint>* is the FQDN or IP address of your Splunk machine.

- *<auth-password>* is the password required to log into the Splunk system.

After completing this stage, you should be able to log into the Services Director GUI and access exported analytics data in the vADC Analytics application.

CLI

1. Start the CLI in configuration mode (not the shell), as described in "Accessing the Services Director VA CLI and OS Shell" on page 8.

2. Create a search endpoint:

```
ssc search-endpoint create search-endpoint address <search-endpoint>:8089 use-tls
true name Test auth-username admin auth-password <auth-password> logs-index
zxtm_logs transactions-index zxtm_transactions
```

GUI

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.

2. Log in as the administration user. The **Home** page appears.

3. Click the **Catalogs** menu, and then click **Analytics > Analytics Endpoints**.

   The **Analytics Endpoints** page appears. For example:

   FIGURE 28   Search Endpoints

4. Click the **Add** button above the **Search Endpoints** table.

   The **Add Search Endpoint** dialog box appears.

5. Complete the following fields:

- **Name**: *Test*. This name will appear in the **Search Endpoints** table after you apply the endpoint.

- **Address**: *<collection-endpoint>:8089*

- **Transactions index**: *zxtm_transactions*

- **Logs index**: *zxtm_logs*

- **Query using TLS**: Select this check box.

- **Username**: *admin*

- **Password**: *<auth-password>*

Leave all other fields with their default values.

6. Click **Test Connection**. You should see a "Connection succeeded" message.

Note: If you see a warning about no data being found, check that the Splunk system's indexes exist and have received data.

For example:

FIGURE 29   Add Search Endpoint

7. Click **Apply**.

    The search endpoint is added to the **Search Endpoint** table. For example:

    FIGURE 30   New Search Endpoint

Analytics Endpoints

Collection Endpoints

⊕ Add

	Name	ID	Transaction Export	Log Export
▶	Test	Collection-Endpoint-TWTB-OH3F-1FMJ-FK1Z	example.com:5000	http://example.com:8088/services/collector/event

Search Endpoints

	Name	ID	Address		
▶	Test	Search-Endpoint-T4NE-VB03-ICSY-TMCS	example.com:8089	Test Connection	✕

Configuring Certificate Verification and Transaction Export over TLS

Introduction

Once you have Services Director, vTM and the Splunk system successfully communicating, you can add transaction export over TLS and certificate verification at various points.

Setting TLS Server Certificates on the Splunk System Endpoints

Establishing TLS server certificates on the Splunk system's endpoints requires three tasks:

- "Obtaining a Signed Server Certificate" on page 39.

- "Configuring SSL on the Search Endpoint" on page 40.

- "Configuring SSL on the Collection Endpoint" on page 40.

Obtaining a Signed Server Certificate

You must obtain a set of keys and certificates signed by a CA. These replace the generic certificates installed by default when deploying a Splunk system.

Guidance on how such a set of keys and certificates might be prepared (without using a commercial Certificate Authority) is provided by Splunk in the following pages:

- http://docs.splunk.com/Documentation/Splunk/7.0.2/Security/Howtoself-signcertificates

  This document includes a step-by-step guide to generating a root CA key/certificate pair, then a server key and a server certificate signed by that CA.

- http://docs.splunk.com/Documentation/Splunk/7.0.2/Security/HowtoprepareyoursignedcertificatesforSplunk

  This document includes a step-by-step guide on how to chain together the server certificate/key and the CA certificate created earlier, to create a server certificate chain.

Once you have a set of certificates and keys, these need to be referenced in the configuration files for the Splunk system, and (in the case of the CA certificate) in Services Director.

Configuring SSL on the Search Endpoint

To secure the search endpoint, a Splunk configuration file needs to be amended to reference the generated keys/certificates.

Note: You currently cannot perform this configuration via the Splunk system's GUI or CLI.

1. Log into the Splunk server.

2. Edit the *$SPLUNK_HOME/etc/system/local/server.conf* file using a text editor, creating the file if necessary.

   Note: On Linux, *$SPLUNK_HOME* will normally be equivalent to */opt/splunk*. It may not be defined as an environment variable, so the fully-qualified filepath may be necessary.

3. Cut/paste the following content into the file, replacing the referenced file paths/names and certificate password as required:

   ```
   [sslConfig]
   sslPassword = $1$D5PA3wWpcA==
   <<---- (encoded version of certificate password - see note below on passwords)

   serverCert = $SPLUNK_HOME/etc/auth/mycerts/myNewServerCertificate.pem
   <<---- (server certificate chain file)

   caCertFile = $SPLUNK_HOME/etc/auth/mycerts/myCACertificate.pem
   <<---- (CA certificate)
   ```

   When editing *server.conf*, the *sslPassword* can be entered in plain text. When the Splunk server next restarts, it will encode the password into the format shown in the example above.

4. Once this change has been made, the Splunk system can be restarted using:

   ```
   $SPLUNK_HOME/bin/splunk restart
   ```

   Alternatively, you can continue to the next section to also configure the collection endpoint, see .

Configuring SSL on the Collection Endpoint

To secure collection endpoints, a Splunk configuration file needs to be amended to reference the generated certificates.

Note: You currently cannot perform this configuration via the Splunk system's GUI or CLI.

1. Log into the Splunk server.

2. Edit the *$SPLUNK_HOME/etc/system/local/inputs.conf* file using a text editor, creating the file if necessary.

   Note: On Linux, *$SPLUNK_HOME* will normally be equivalent to */opt/splunk*. It may not be defined as an environment variable, so the fully-qualified filepath may be necessary.

3. Cut/paste the following content into the file, replacing the referenced file paths/names and certificate password as required:

```
[splunktcp-ssl:5000]
<<---- (the number is the port used for the transaction export collector endpoint)
disabled = 0

[SSL]
serverCert = $SPLUNK_HOME/etc/auth/mycerts/myNewServerCertificate.pem
<<---- (server certificate chain file)

sslPassword = $1$S9aGnlDlcA==
<<---- (encoded version of certificate password - see note below on passwords)
```

Note: When editing *inputs.conf*, the *sslPassword* can be entered in plain text. When the Splunk system next restarts, it will encode the password into the format shown in the example above.

4. Restart the Splunk server using:

```
$SPLUNK_HOME/bin/splunk restart
```

Setting Up Transaction Data Export over TLS and TLS Verification

This procedure assumes you already have a collection endpoint with operational transaction data import settings, and which uses unverified TLS.

GUI

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.

2. Log in as the administration user. The **Home** page appears.

3. Click the **Catalogs** menu, and then click **Analytics > Analytics Endpoints**.

   The **Analytics Endpoints** page appears. For example:

   FIGURE 31   Analytics Endpoints

4. Expand the collection endpoint that you wish to secure (in this example, *dself-07_collect*):

FIGURE 32   Collection Endpoints Expanded

5. In the **Transaction Export Collector** section, paste the PEM contents of the CA certificate file into the **Certificate > From** text field.

For example, the contents of the *myCACertificate.pem* file from the Configuring SSL on the search endpoint section, see "Configuring SSL on the Search Endpoint" on page 40.

FIGURE 33   Collection Endpoints Complete

6. Select the **Verify TLS** checkbox.

7. Click **Apply**.

Setting Up Log Data Export TLS Verification

This procedure assumes you already have a collection endpoint with operational transaction data import settings, and which uses unverified TLS.

GUI

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.

2. Log in as the administration user. The **Home** page appears.

3. Click the **Catalogs** menu, and then click **Analytics > Analytics Endpoints**.

    The **Analytics Endpoints** page appears. For example:

    FIGURE 34   Collection Endpoints

4. Expand the collection endpoint that you wish to secure (in this example, *dself-07_collect*):

    FIGURE 35   Collection Endpoint Expanded

5. In the **Log Export Collector** section, paste the PEM contents of the CA certificate file into the **Certificate > From** text field.

For example, the contents of the *myCACertificate.pem* file from the Configuring SSL on the Collection Endpoint section above, see "Configuring SSL on the Collection Endpoint" on page 40.

FIGURE 36   Collection Endpoint New Certificate

6. Check the **Verify TLS** checkbox.

7. Click **Apply**.

Setting Up Search Endpoint TLS Verification

This procedure assumes you already have a search endpoint with operational transaction data import settings, and which uses unverified TLS.

GUI

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.

2. Log in as the administration user. The **Home** page appears.

3. Click the **Catalogs** menu, and then click **Analytics > Analytics Endpoints**.

The **Analytics Endpoints** page appears. For example:

FIGURE 37   Search Endpoints

4. Expand the search endpoint you wish to secure (in this example, *dself-07_search*):

FIGURE 38   Search Endpoint Expanded

5. Paste the PEM contents of the CA certificate file into the **Certificate > From** text field.

For example, the contents of the *myCACertificate.pem* file from "Configuring SSL on the Search Endpoint" on page 40.

FIGURE 39   Search Endpoint Complete

6. Check the **Verify TLS** checkbox.

7. Click **Apply**.

8. (Optional) Click the **Test Connection** button to ensure the search endpoint is responsive as configured.

Diagnosing and Resolving Problems

Diagnosing and Resolving Problems

If the Splunk system is not behaving as expected, it is possibly a problem caused by an edit to one of its configuration files. Perform following the steps:

1. Log into the Splunk server's command line using SSH.

2. Enter the *opt/splunk* directory.

3. Perform a check of Splunk configuration files:

   ```
   sudo bin/splunk btool check
   ```

4. If there is a problem with a file, try to fix it using the appropriate Splunk documentation until the command above exits with no complaints.

5. Restart the Splunk system to make sure it has picked up the fixes.