



Pulse Secure Services Director Splunk Deployment Guide

Supporting Pulse Secure Services Director 19.1r1

| | |
|-------------------|--------------------------|
| Product Release | 19.1r1 |
| Published Date | 30 September 2019 |
| Document Revision | 1.0 |

Pulse Secure, LLC

2700 Zanker Road, Suite 200 San Jose, CA 95134

The Pulse Secure Logo, the Pulse logo, and PulseE are trademarks of Pulse Secure, LLC. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Pulse Secure assumes no responsibility for any inaccuracies in this document. Pulse Secure reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Pulse Secure or components thereof might be covered by one or more of the following patents that are owned by or licensed to Pulse Secure: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899,

6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Services Director Splunk Deployment Guide

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

END USER LICENSE AGREEMENT

The Pulse Secure product that is the subject of this technical documentation consists of (or is intended for use with) Pulse Secure software. Use of such software is subject to the terms and conditions of the End User License Agreement (“EULA”) posted at <https://www.pulsesecure.net/support>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

| | |
|-----------------------------------------------------------------------|----|
| About This Document | 4 |
| Introduction..... | 5 |
| Key Concepts..... | 5 |
| Required Configuration..... | 6 |
| Network Diagram..... | 6 |
| Prerequisites | 7 |
| Overview..... | 8 |
| Accessing the CLI and the OS Shell of the Services Director VA | 8 |
| Variations..... | 8 |
| Configuring the Splunk System..... | 9 |
| Creating the Indexes..... | 9 |
| Creating a Source Type for Transactions | 11 |
| Creating a Data Input for Transactions | 13 |
| Creating a Data Input For Logs..... | 15 |
| Setting Up Event Log Entry Parsing..... | 19 |
| Setting Up Audit Log Entry Parsing | 22 |
| Setting Up Case-Insensitive Searches for HTTP Headers | 25 |
| Configuring Services Director..... | 27 |
| Creating a Collection Endpoint..... | 27 |
| Creating and Applying an Analytics Profile | 29 |
| Creating a Search Endpoint | 31 |
| Adding Certificate Verification and Transaction Export over TLS | 34 |
| Setting TLS Server Certificates on the Splunk System Endpoints | 34 |
| Setting Up Transaction Data Export over TLS and TLS Verification..... | 36 |
| Setting Up Log Data Export TLS Verification | 37 |
| Setting Up Search Endpoint TLS Verification | 39 |
| Diagnosing Problems | 41 |

About This Document

The aim of this document is to provide a definitive set of steps to create a working Splunk¹ / Services Director Analytics App deployment.

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

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

¹ Splunk is a registered trademark of Splunk Inc. in the USA and other countries.

Introduction

When used in conjunction with Services Director version 18.1, clusters of Pulse Secure Virtual Traffic Manager (vTM) versions 17.2 and above can be configured to export live analytics data to an external Splunk analytics system. Virtual Traffic manager clusters 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. Please note that 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

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:

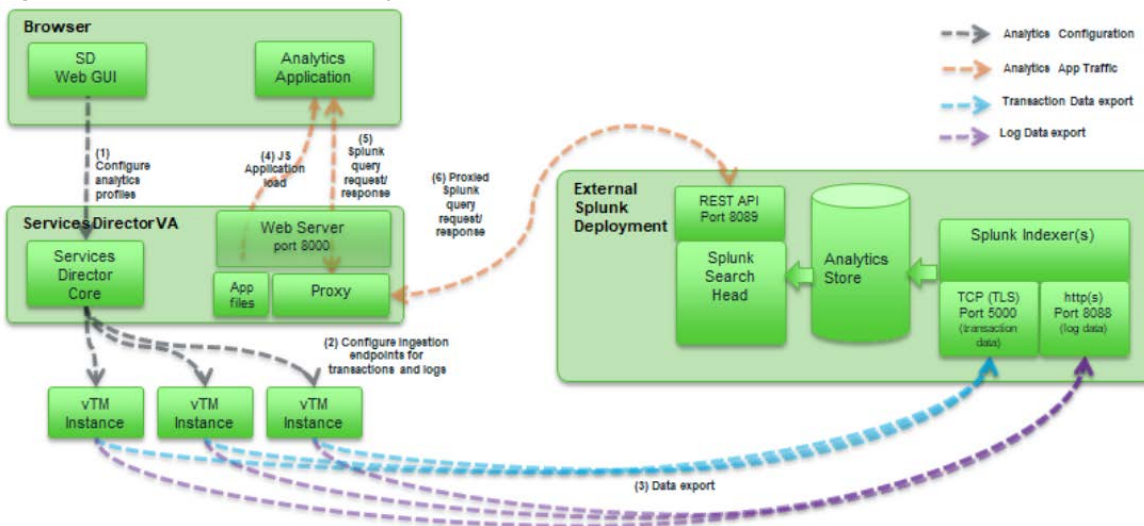
1. The Splunk system must be configured with appropriate collection endpoints to accept and store these records in the correct indices for later analysis.
Specifically, the Splunk system must be configured with a raw TCP input and an HTTP Event Collector input. The raw TCP input will collect the transaction 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.

Figure 1: Services Director and Splunk



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

You must already have the following deployed:

- A machine running a Splunk system.



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

- 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.

Overview

This document is focused on the correct configuration of the Splunk system and Services Director so that analytics data can be seen in the Services Director "vADC Analytics" application. The processes of setting up and licensing Services Director and vTM are covered in detail elsewhere. The first section of this guide covers configuring Splunk. Once Splunk is configured correctly and running, the instructions for configuring Services Director can be followed.

Accessing the CLI and the OS Shell of the Services Director VA

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**

Variations

For simplicity and clarity, 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.

Configuring the Splunk System

The configuration process can be split into a number of stages, most of which are short.

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 the Indexes

Services Director requires log data and transaction data to be in separate indexes. This keeps events data organised, 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 fields as follows:
 - **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`

The screenshot shows the 'New Index' configuration window. The 'Index Name' field is filled with 'zxtm_transactions'. The 'Home Path', 'Cold Path', and 'Thawed Path' fields are all set to 'optional'. The 'Data Integrity Check' is a toggle set to 'Enable'. The 'Max Size of Entire Index' is '500 GB'. The 'Max Size of Hot/Warm/Cold Bucket' is 'auto_high_volume GB'. The 'Frozen Path' is 'optional'. The 'App' dropdown is set to 'Search & Reporting'. Under 'Storage Optimization', the 'Tsidx Retention Policy' is set to 'Enable Reduction'. A warning message is visible below this setting: 'Warning: Do not enable reduction without understanding the full implications. It is extremely difficult to rebuild reduced buckets. [Learn More](#)'. The 'Cancel' button is on the left and the 'Save' button is on the right.

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

The screenshot shows the 'New Index' configuration window. The 'Index Name' is set to 'zxtm_logs'. The 'Home Path', 'Cold Path', and 'Thawed Path' are all set to 'optional'. The 'Data Integrity Check' is set to 'Enable'. The 'Max Size of Entire Index' is set to '500 GB'. The 'Max Size of Hot/Warm/Cold Bucket' is set to 'auto_high_volume GB'. The 'Frozen Path' is set to 'optional'. The 'App' is set to 'Search & Reporting'. The 'Tsidx Retention Policy' is set to 'Enable Reduction'. There is a warning message below the retention policy: 'Warning: Do not enable reduction without understanding the full implications. It is extremely difficult to rebuild reduced buckets. [Learn More](#)'. The 'Save' button is highlighted in green.

8. Click **Save**.

Creating a Source Type for Transactions

Creating a source type allows the Splunk system to interpret transaction data correctly. Unfortunately, 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

The screenshot shows the 'Create Source Type' dialog box with the following configuration:

- Name:** zxtm_transactions
- Description:** optional
- Destination app:** Search & Reporting
- Category:** Network & Security
- Indexed Extractions?** none
- Event Breaks:** Break Type: Auto, Every Line, Regex...
- Timestamp:** Extraction: Auto, Current time, Advanced...
- Advanced:**

| Name | Value |
|------------------|--------------------|
| CHARSET | |
| SHOULD_LINEMERGE | false |
| NO_BINARY_CHECK | true |
| category | Network & Security |
| KV_MODE | json |

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 fields:
 - **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

Add Data Progress: Select Source | **Input Settings** | Review | Done < **Review >**

Input Settings
Optionally set additional input parameters for this data input as follows:

Source type
The source type is one of the default fields that Splunk assigns to all incoming data. It tells Splunk what kind of data you've got, so that Splunk can format the data intelligently during indexing. And it's a way to categorize your data, so that you can search it easily.
Buttons: Select, New
Dropdown: zxtm_transactions

App context
Application contexts are folders within a Splunk instance that contain configurations for a specific use case or domain of data. App contexts improve manageability of input and source type definitions. Splunk loads all app contexts based on precedence rules. [Learn More](#)
App Context: Search & Reporting

Host
When Splunk indexes data, each event receives a "host" value. The host value should be the name of the machine from which the event originates. The type of input you choose determines the available configuration options. [Learn More](#)
Method?: IP, DNS, Custo...

Index
Splunk stores incoming data as events in the selected index. Consider using a "sandbox" index as a destination if you have problems determining a source type for your data. A sandbox index lets you troubleshoot your configuration without impacting production indexes. You can always change this setting later. [Learn More](#)
Index: zxtm_transactions [Create a new index](#)

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

Figure 6: Add Data Wizard: Review

Add Data Progress: Select Source | Input Settings | **Review** | Done < **Submit >**

Review

| | |
|----------------------|----------------------------------|
| Input Type | TCP Port |
| Port Number | 5001 |
| Source name override | N/A |
| Restrict to Host | N/A |
| Source Type | zxtm_transactions |
| App Context | launcher |
| Host | (DNS entry of the remote server) |
| Index | zxtm_transactions |

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

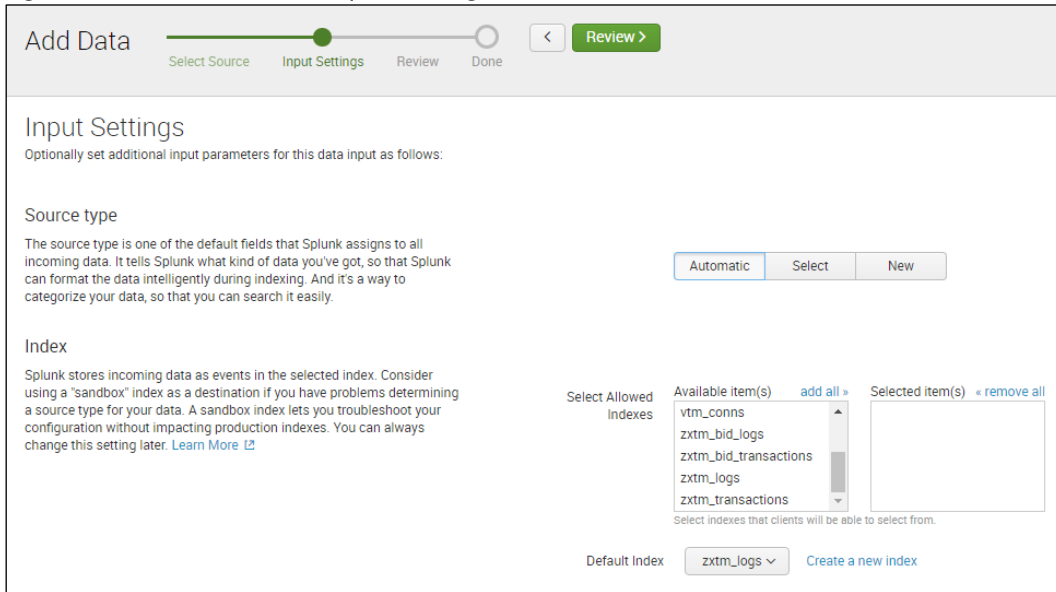
1. Navigate to the Splunk system's web interface and login.
2. From the menu bar, select **Settings > Data > Data inputs**.
3. On the **Data inputs** page, under **Local Inputs > HTTP Event Collector**, click the **Add New** action.

The **Add Data** wizard starts.

4. In the **Select Source** pane of the wizard, complete the following fields:
 - **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`.

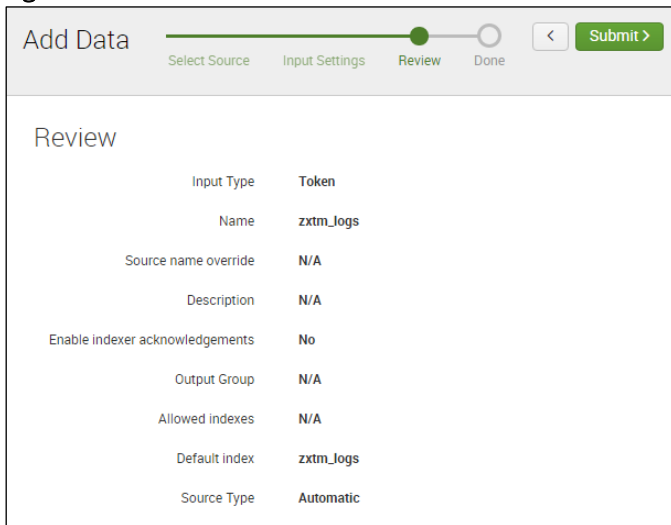
For example:

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:

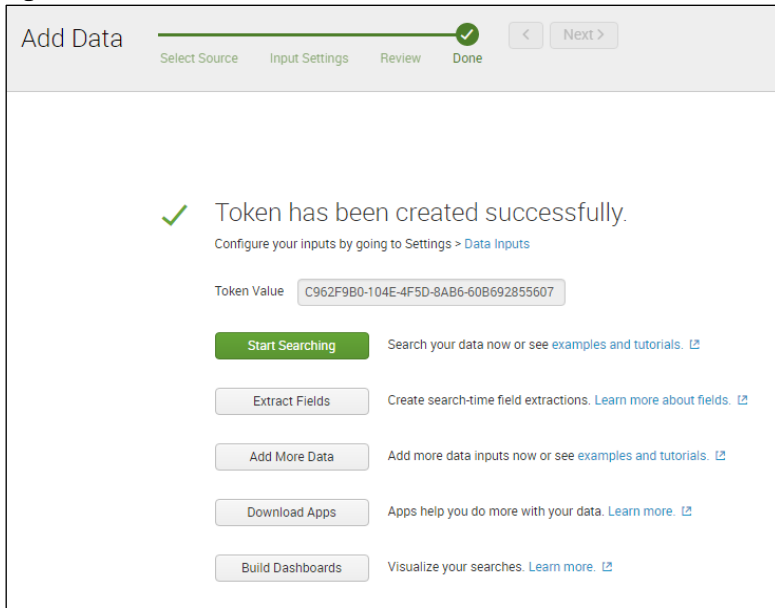
Figure 8: Add Data Wizard: Review



9. Click **Submit** >.

A confirmation screen is displayed. For example:

Figure 9: Add Data Wizard: Token



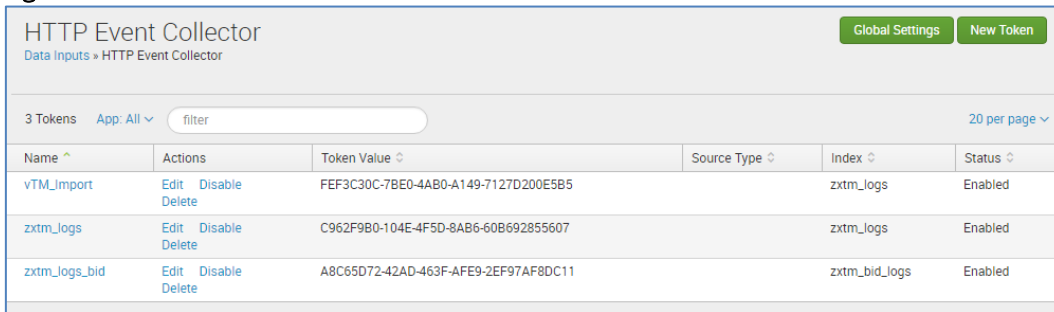
10. Make a note of the **Token Value**.

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



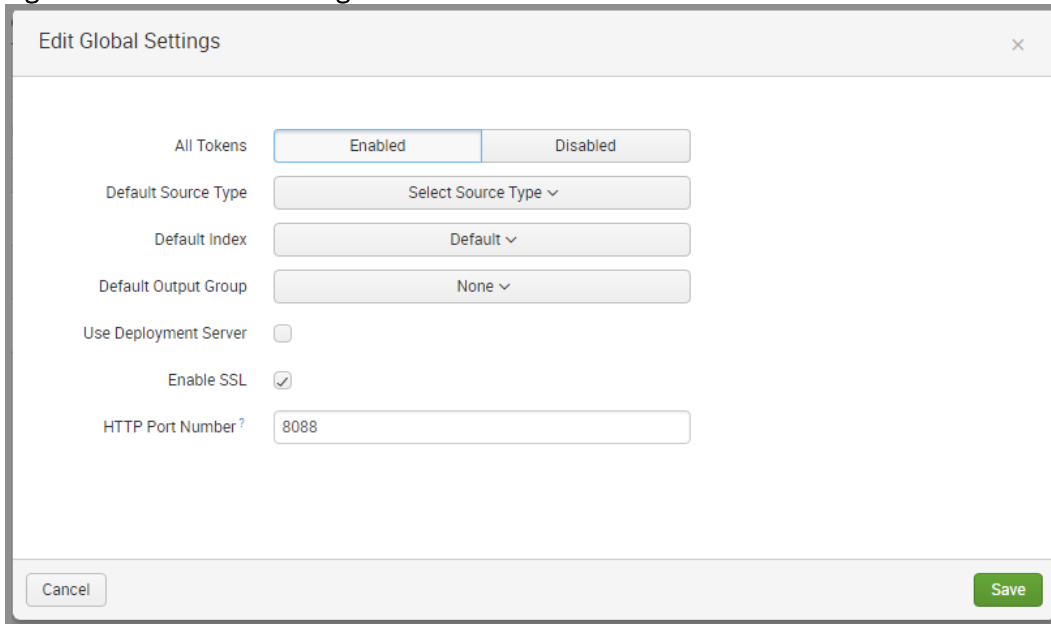
13. Click **Global Settings**.

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

- **All Tokens:** *Enabled*
- Leave all other fields with their default values.

For example:

Figure 11: Edit Global Settings



The screenshot shows a dialog box titled "Edit Global Settings" with a close button in the top right corner. The settings are as follows:

- All Tokens: Enabled (selected), Disabled
- Default Source Type: Select Source Type (dropdown)
- Default Index: Default (dropdown)
- Default Output Group: None (dropdown)
- Use Deployment Server:
- Enable SSL:
- HTTP Port Number?: 8088 (text input)

At the bottom of the dialog, there are two buttons: "Cancel" on the left and "Save" on the right.

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

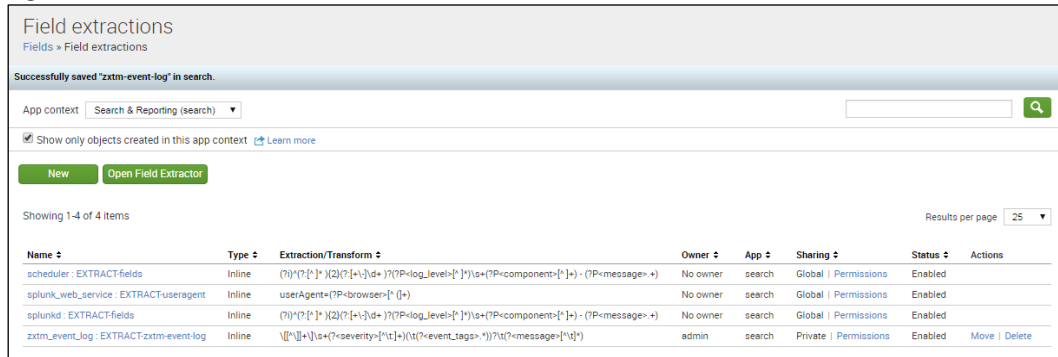
The screenshot shows the 'Add new' configuration page in Splunk. The breadcrumb is 'Fields > Field extractions > Add new'. The form contains the following fields:

- Destination app:** A dropdown menu with 'search' selected.
- Name *:** A text input field containing 'zxtm-event-log'.
- Apply to:** A dropdown menu with 'sourcetype' selected.
- named *:** A text input field containing 'zxtm_event_log'.
- Type *:** A dropdown menu with 'Inline' selected.
- Extraction/Transform *:** A text input field containing the regular expression: `\\[[^\\]]+\\]s+(?<severity>[^\\t:]+)(\\t(?<event_tags>.*))?\\t(?<message>[^\\t]*)`. Below this field is a small note: 'If the field extraction is inline, provide the regular expression. If the field extraction uses a transform, specify the transform name.'

At the bottom of the form, there are 'Cancel' and 'Save' buttons.

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

Figure 13: 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

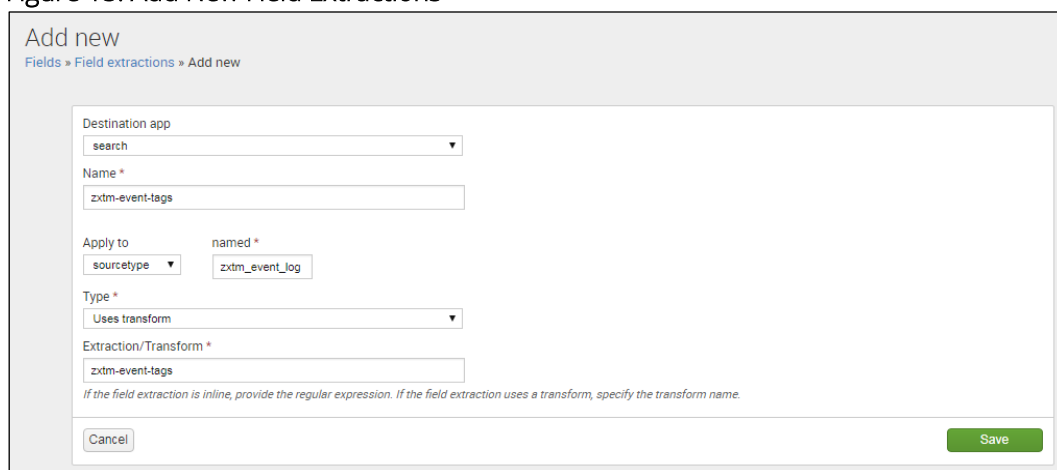
The screenshot shows the 'Add new' configuration page for field transformations. The breadcrumb trail is 'Fields > Field transformations > Add new'. The form contains the following fields and options:

- Destination app:** A dropdown menu with 'search' selected.
- Name *:** A text input field containing 'zxtm-event-tags'.
- Type:** A dropdown menu with 'delimiter-based' selected.
- Delimiters:** A text input field containing '|'.
- Field list:** A text input field containing 'tag, tag, tag, tag, tag'. Below the field is the instruction: 'Specify the comma separated list of field names.'
- Source Key:** A text input field containing 'events_tag'. Below the field is the instruction: 'Specify the key the transforms extraction applies to. Default is _raw.'
- Options:**
 - Create multivalued fields**
If checked the extractor will create multivalued fields if the field is already extracted.
 - Automatically clean field names**
If checked the field names will be cleaned such that they only contain: a-zA-Z0-9_.
- Buttons:** 'Cancel' and 'Save' (green).

- 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

The screenshot shows the 'Add new' configuration page for field transformations. The breadcrumb trail is 'Fields > Field transformations > Add new'. The form contains the following fields and options:

- Destination app:** A dropdown menu with 'search' selected.
- Name *:** A text input field containing 'zxtm-audit-fields'.
- Type:** A dropdown menu with 'delimiter-based' selected.
- Delimiters:** A text input field containing '"\t", "="'.
- Field list:** An empty text input field.
- Source Key:** A text input field containing '_raw'.
- Options:**
 - Create multivalued fields. Below it, the text reads: 'If checked the extractor will create multivalued fields if the field is already extracted.'
 - Automatically clean field names. Below it, the text reads: 'If checked the field names will be cleaned such that they only contain: a-zA-Z0-9_.'
- Buttons:** 'Cancel' and 'Save' buttons are located at the bottom of the form.

- 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

The screenshot shows a web form titled "Add new" with a breadcrumb "Fields » Field extractions » Add new". The form contains the following fields and values:

- Destination app: search
- Name *: zxtm-audit-logs
- Apply to: sourcetype
- named *: zxtm_audit_log
- Type *: Uses transform
- Extraction/Transform *: zxtm-audit-fields


Below the fields is a note: "If the field extraction is inline, provide the regular expression. If the field extraction uses a transform, specify the transform name." At the bottom of the form are "Cancel" and "Save" buttons.

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 extractions**, 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 field**).
 - 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

| Name | Field aliases | Owner | App | Sharing | Status | Actions |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|--------|-----------------------|---------|-----------------------|
| zxtm_transactions : FIELDALIAS-User-Agent | "http.request.USER-AGENT" AS "http.request.User-Agent" "http.request.User-agent" AS "http.requestUser-Agent" "http.request.user-agent" AS "http.request.User-Agent" | admin | search | Private Permissions | Enabled | Clone Move Delete |

Configuring Services Director

Services Director is relatively simple to configure.

- 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.

For clarity and simplicity, the CLI commands are given where possible. The equivalent steps in the GUI should be similar, with fields being left with their default settings unless mentioned in the CLI command.

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 Creating a data input for logs procedure earlier in this document.

CLI

1. Start the CLI in configuration mode (not the shell), as described in Accessing the CLI and the OS shell.
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

| Analytics Endpoints | | | |
|----------------------|----|--------------------|------------|
| Collection Endpoints | | | |
| Name | ID | Transaction Export | Log Export |
| No Data | | | |
| Search Endpoints | | | |
| Name | ID | Address | |
| No Data | | | |

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

Add Collection Endpoint

Name:

Transaction Export Collector Settings

Address (<IP address/hostname>:<port>):

Export over TLS:

Verify TLS:

Certificate:

From file

[Choose File](#)

From text

Log Export Collector Settings

HTTP(S) URL:

Verify TLS:

Authentication Method:

HEC Token:

Certificate:

From file

[Choose File](#)

From text

6. Click **Apply**.

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

Figure 22: New Collection Endpoint

The screenshot shows the 'Analytics Endpoints' configuration page. It has two sections: 'Collection Endpoints' and 'Search Endpoints'. The 'Collection Endpoints' section has an 'Add' button and a table with the following data:

| Name | ID | Transaction Export | Log Export |
|------|-----------------------------------------|--------------------|--------------------------------------------------|
| Test | Collection-Endpoint-2MWH-Y35Y-RWIU-9UBV | example.com:5000 | http://example.com:8088/services/collector/event |

The 'Search Endpoints' section also has an 'Add' button and a table that is currently empty, displaying 'No Data'.

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 CLI and the OS shell](#).
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

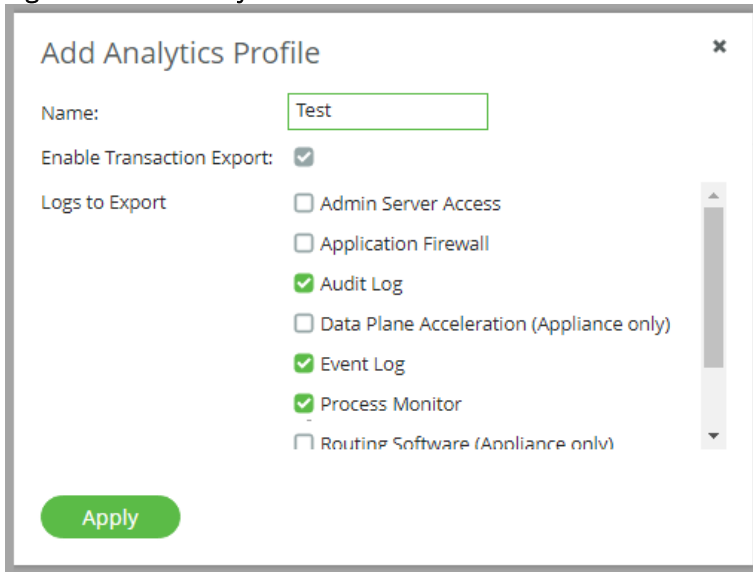
The screenshot shows the 'Analytics Profiles' configuration page. It has an 'Add' button and a table with the following headers:

| 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 some 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

| Analytics Profiles | | | |
|--------------------|---------------------------------------|---------------------------------------|-------------------------|
| Name | ID | Logs to export | Transaction Data Export |
| Test | Analytics-Profile-M17V-9N35-7KA8-0G3W | Audit Log, Event Log, Process Monitor | Enabled |

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

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

Figure 26: vTM Clusters

| vTM Clusters | | | | | | | |
|-----------------------------|------------|--------|-------------------|-----------------|------------------|------------|-------------|
| Cluster Name | Type | In Use | Analytics Profile | Backup Schedule | Next Backup Time | Action | Last Action |
| Cluster-1QFP-Y1AC-UBN3-3SR0 | Discovered | ✓ | N/A | N/A | | Backup Now | |

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

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

Figure 27: vTM Clusters Expanded

The screenshot shows the 'vTM Clusters' interface. At the top, there is an 'Add' button. Below it is a table with the following columns: Cluster Name, Type, In Use, Analytics Profile, Backup Schedule, Next Backup Time, Action, Last Action, and Last Action Status. A single cluster is listed: Cluster-1QFP-Y1AC-UBN3-3SR0, with Type 'Discovered', In Use checked, Analytics Profile 'N/A', and Backup Schedule 'N/A'. An action button 'Backup Now' is visible. Below the table is a configuration form with fields for Cluster Name, Owner (JK), Analytics Profile (Test), Backup Schedule (N/A), and Number of Backups (5). There are 'Apply' and 'Revert' buttons at the bottom of the form.

- 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

- Start the CLI in configuration mode (not the shell), as described in [Accessing the CLI and the OS shell](#).
- 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

- Access your Services Director VA from a browser, using its Service Endpoint IP Address.
- Log in as the administration user. The **Home** page appears.
- Click the **Catalogs** menu, and then click **Analytics > Analytics Endpoints**.

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

Figure 28: Search Endpoints

The screenshot shows the 'Analytics Endpoints' page. It has two main sections: 'Collection Endpoints' and 'Search Endpoints'. The 'Collection Endpoints' section has an 'Add' button and a table with one row. The 'Search Endpoints' section also has an 'Add' button and a table that is currently empty, showing 'No Data'.

| Collection Endpoints | | | |
|----------------------|-----------------------------------------|--------------------|--------------------------------------------------|
| 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 |
| No Data | | |

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

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

- 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.
- 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

Add Search Endpoint

Name:

Address:

Transactions index:

Logs index:

Query using TLS:

Verify TLS:

Certificate: From file
 [Choose File](#)

From text

Username:

Password:

Connection succeeded

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 <input type="button" value="Test Connection"/> |

Adding Certificate Verification and Transaction Export over TLS

It's simplest to get Services Director, vTM and the Splunk system successfully communicating with a minimal amount of security, as this reduces the number of things that can go wrong and makes debugging easier. Once you're ready to move on, you can add transaction export over TLS and certificate verification at various points. If you're confident enough, you can integrate these instructions with those above to set up your system securely in the first place.

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.
- Configuring SSL on the search endpoint.
- Configuring SSL on the collection endpoint.

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:** There is at present no way to 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/system.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.

- 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)
```

i Note: 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.

Once this change has been made, the Splunk system can be restarted using `$SPLUNK_HOME/bin/splunk restart`, or you can continue to the next section to also configure the collection endpoint.

Configuring SSL on the Collection Endpoint

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

i Note: There is at present no way to perform this configuration via the Splunk system's GUI or CLI.

- Log into the Splunk server
- Edit the `$SPLUNK_HOME/etc/system/local/inputs.conf` file using a text editor, creating the file if necessary.

i 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.

- 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 = $1$S9aGnIDlcA==
<<---- (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.

Once this change has been made, restart the Splunk server using `$SPLUNK_HOME/bin/splunk restart`.

Setting Up Transaction Data Export over TLS and TLS Verification

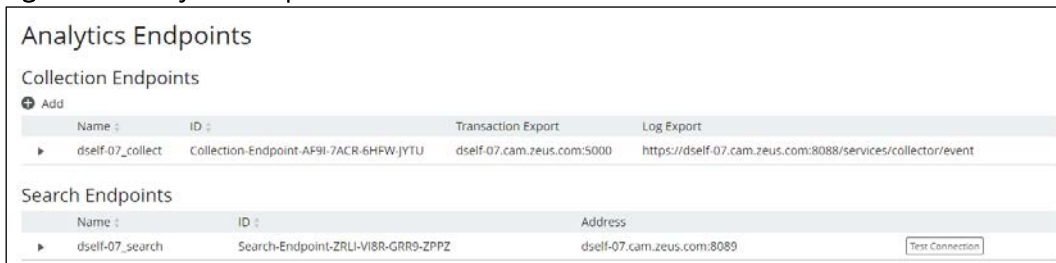
This procedure assumes you already have a collection endpoint with operational transaction data import settings, and which uses unverified TLS.

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**.

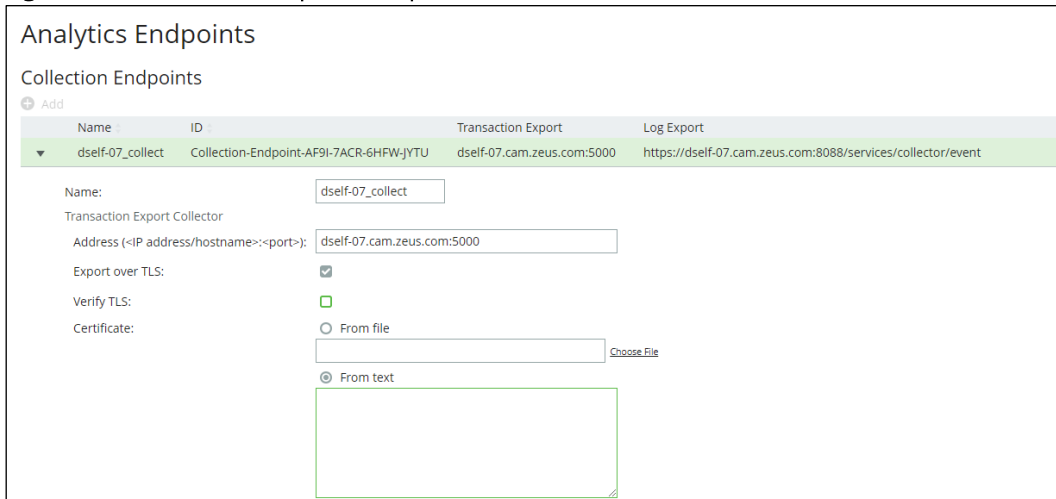
The **Analytics Endpoints** page appears. For example:

Figure 31: Analytics Endpoints



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

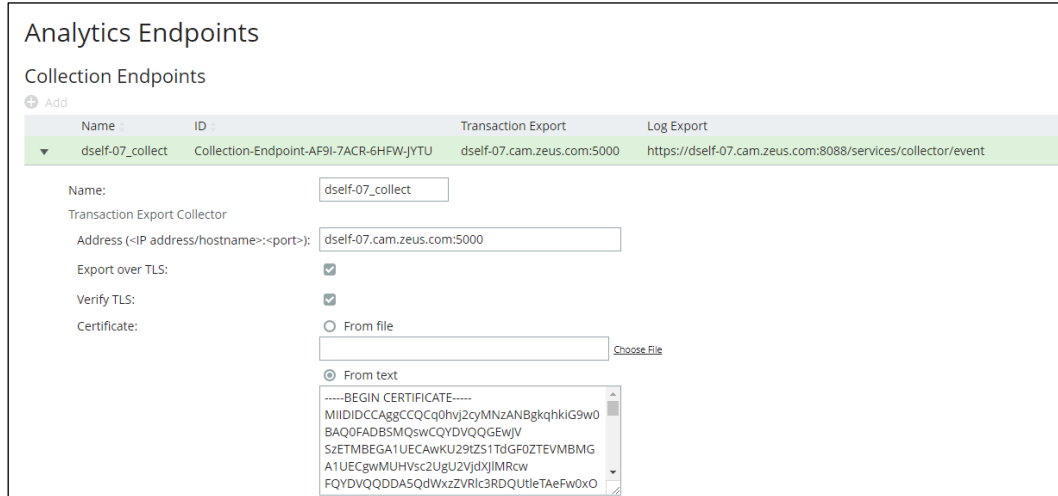
Figure 32: Collection Endpoints Expanded



- 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 above:

Figure 33: Collection Endpoints Complete



- Select the **Verify TLS** checkbox.
- Press the **Apply** button at the end of the form.

Setting Up Log Data Export TLS Verification

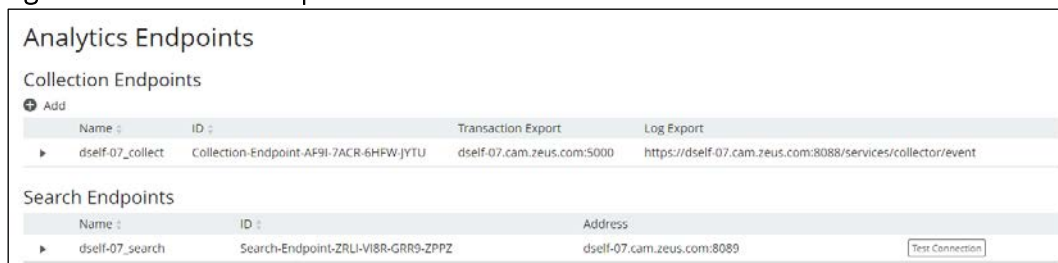
This procedure assumes you already have a collection endpoint with operational transaction data import settings, and which uses unverified TLS.

GUI

- Access your Services Director VA from a browser, using its Service Endpoint IP Address.
- Log in as the administration user. The **Home** page appears.
- Click the **Catalogs** menu, and then click **Analytics > Analytics Endpoints**.

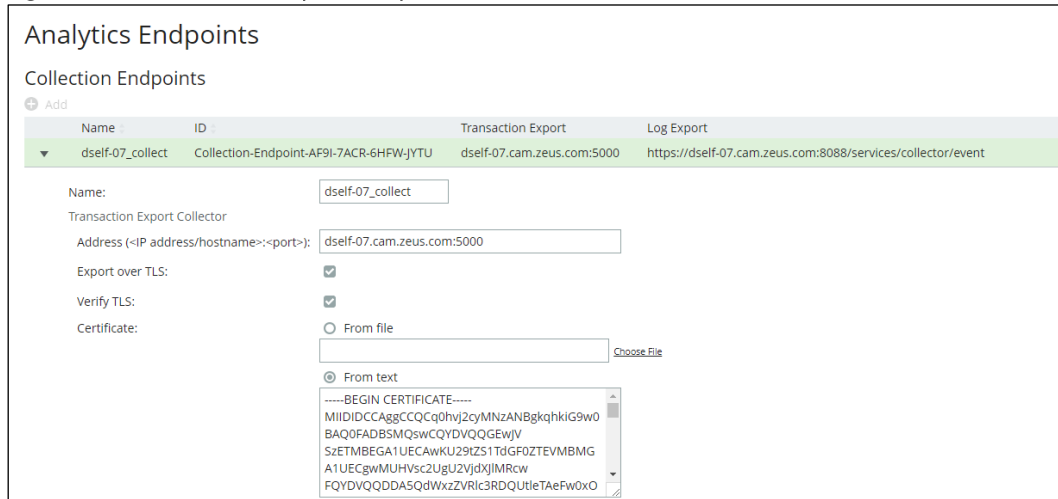
The **Analytics Endpoints** page appears. For example:

Figure 34: Collection Endpoints



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

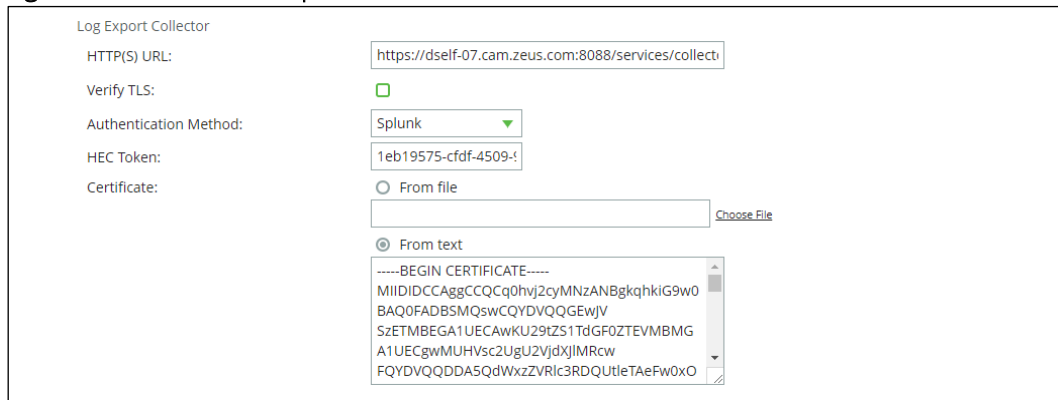
Figure 35: Collection Endpoint Expanded



- 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:

Figure 36: Collection Endpoint New Certificate



- Check the **Verify TLS** checkbox
- Press the **Apply** button at the end of the form.

Setting Up Search Endpoint TLS Verification

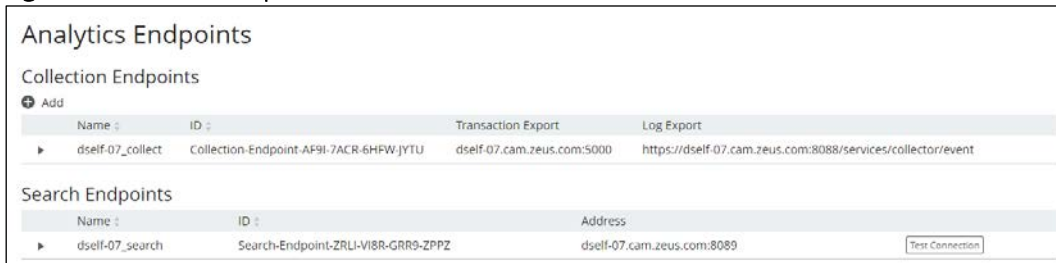
This procedure assumes you already have a search endpoint with operational transaction data import settings, and which uses unverified TLS.

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**.

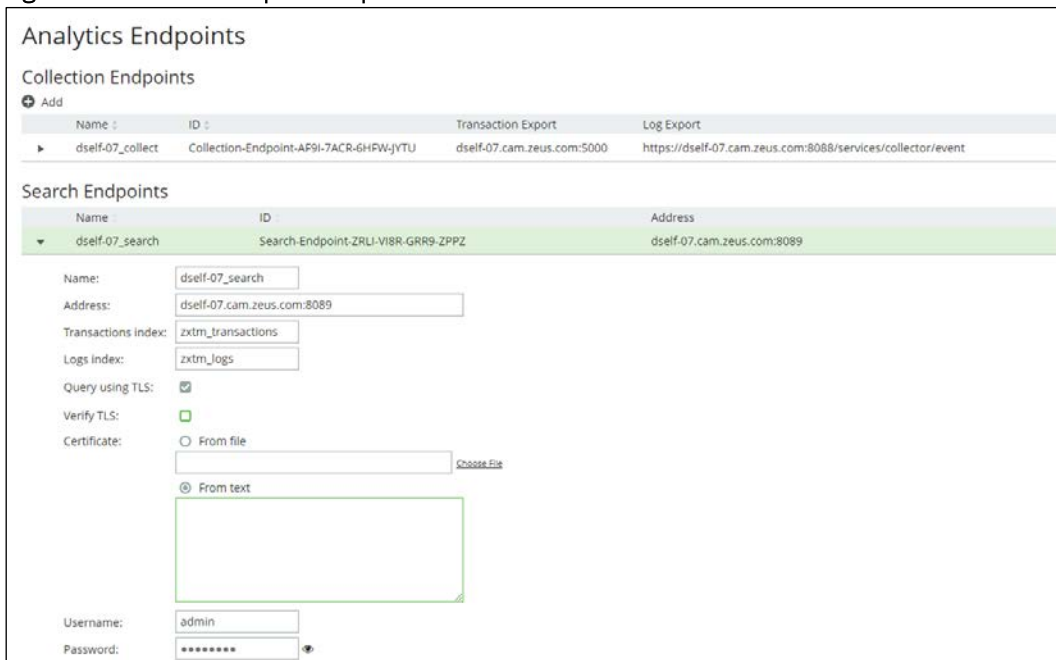
The **Analytics Endpoints** page appears. For example:

Figure 37: Search Endpoints



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

Figure 38: Search Endpoint Expanded



- 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 above:

Figure 39: Search Endpoint Complete

The screenshot shows the 'Analytics Endpoints' configuration page. Under 'Collection Endpoints', there is a table with one entry: 'dself-07_collect' with ID 'Collection-Endpoint-AF9I-7ACR-6HFW-JYTU', Transaction Export 'dself-07.cam.zeus.com:5000', and Log Export 'https://dself-07.cam.zeus.com:8088/services/collector/event'. Under 'Search Endpoints', there is a table with one entry: 'dself-07_search' with ID 'Search-Endpoint-ZRLI-VI8R-GRR9-ZPPZ' and Address 'dself-07.cam.zeus.com:8089'. Below the table, the configuration for 'dself-07_search' is shown: Name: dself-07_search, Address: dself-07.cam.zeus.com:8089, Transactions index: zxtm_transactions, Logs index: zxtm_logs, Query using TLS: checked, Verify TLS: unchecked, Certificate: From text (selected), Certificate content: -----BEGIN CERTIFICATE-----\nMIIDIDCCAggCCQcQ0hmj2oyMNzANBgkqhkiG9w0\nBAQ9FADBSMQswCQYDVQQGEwjV\nSZETMBEGA1UECAwKU29tZS1TdGF0ZTEVMBMG\nA1UECgwMUHVsc2UgU2VjdXJlMRcw\nFOYDVQODDASQdWxzZVRlc3RDQUtleTaeFw0xO\n, Username: admin, Password: *****.

- Check the **Verify TLS** checkbox.
- Press the **Apply** button at the end of the form.
- (Optional) Click the **Test Connection** button to ensure the search endpoint is responsive as configured.

Diagnosing 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.