

*IBM Z Common Data Provider Hands-
On Guide for streaming SMF data to
Elasticsearch*



Contents

- Overview.....1**
- Scenario in this book.....2**
- Preparation checklist..... 3**
- Installing..... 4**
- Configuring the Configuration Tool on Liberty..... 5**
 - Configuring user IDs, group IDs, and security product..... 5
 - Setting up a Liberty server directory and a working directory for the Configuration Tool..... 7
 - Creating the started tasks of the Configuration Tool server and its angel server..... 8
 - Starting the Configuration Tool server.....8
 - Accessing the Configuration Tool..... 9
- Creating a policy to stream SMF data..... 9**
- Creating the System Data Engine started task for streaming SMF data..... 11**
- Configuring the Data Streamer..... 11**
- Configuring and starting Logstash..... 12**
- Starting components..... 14**
- Stopping components.....14**

Overview

IBM® Z Common Data Provider provides the infrastructure for accessing IT operational data from z/OS® systems and streaming it to the analytics platform in a consumable format. It is a single data provider for sources of both structured and unstructured data, and it can provide a near real-time data feed of z/OS log data and System Management Facilities (SMF) data to your analytics platform.

Purpose of this book

This book provides step-by-step instructions for setting up a Proof-of-Value environment with minimum configuration and customization. The purpose of this book is to help users who are new to IBM Z Common Data Provider understand its installation and configuration flow, and its basic functions.

Basic concepts used in the book

Ensure that you understand the following concepts that are generally used in the book before you get started:

Data stream

A stream of data is a set of logically associated data in a standard format. It can be a real log, or a virtual log. It is routed to, and transformed by, the Data Streamer in a predictable way, and is delivered to one or more subscribers.

Subscriber

A subscriber is the software that you define to receive operational data. It can be analytics software like Splunk, or intermediary software like Logstash.

Policy

A policy defines how the data is collected, transformed, and streamed. It is created in the Configuration Tool and includes the following information:

- The source from which you want to collect data.
- How to alter the operational data in the data stream so that it is consumable at its target destination.
- The subscribers that receive the data.

Working directory

When you configure some IBM Z Common Data Provider components, you must define a working directory for the component to store its runtime files.

IBM Z Common Data Provider basic components

IBM Z Common Data Provider includes the following basic components:

- A Configuration Tool for defining the sources from which you want to collect operational data
- The data gatherer components (System Data Engine and Log Forwarder) for gathering different types of operational data
- A Data Streamer for streaming all data to its destination

Configuration Tool

The IBM Z Common Data Provider Configuration Tool is a web-based user interface that is provided as an application for IBM WebSphere® Application Server for z/OS Liberty, or as a plug-in for IBM z/OS Management Facility (z/OSMF). In the tool, you specify the configuration information as part of creating a *policy* for streaming operational data to its destination.

Data gatherer components

Each of the following components gathers a different type of data:

System Data Engine

The System Data Engine gathers System Management Facilities (SMF) data and IBM Information Management System (IMS) log data in near real time. It can also gather SMF data and IMS data in batch.

Log Forwarder

The Log Forwarder gathers z/OS log data from the following sources:

- Job log, which is output that is written to a data definition (DD) by a running job
- z/OS UNIX log file, including the UNIX System Services system log (syslogd)
- Entry-sequenced Virtual Storage Access Method (VSAM) cluster
- z/OS system log (SYSLOG)
- IBM Tivoli® NetView® for z/OS messages
- IBM WebSphere Application Server for z/OS High Performance Extensible Logging (HPXL) log
- IBM Information Management System (IMS) log
- IBM Resource Measurement Facility (RMF) Monitor III reports

Data Streamer

The Data Streamer streams operational data to configured subscribers in the appropriate format. It receives the data from the data gatherers, alters the data to make it consumable for the subscriber, and sends the data to the subscriber.

Data Receiver

The Data Receiver acts as a target subscriber if the intended subscriber of a data stream cannot directly ingest the data feed from IBM Z Common Data Provider. The Data Receiver writes any data that it receives to disk files, which can then be ingested into an analytics platform such as Splunk.

Scenario in this book

A Proof-of-Value environment with basic functions will be set up to demonstrate the installation and configuration flow.

This scenario includes the following features:

- A Configuration Tool on Liberty.
- A System Data Engine for collecting SMF data.
- A Data Streamer for sending data to the subscriber.
- A policy that streams SMF 30 and SMF_110_1_KPI to Elastic Stack via Logstash.

The following high-level tasks demonstrate the roadmap for setting up the Proof-of-Value environment.

1. Go through the preparation checklist and make sure that your environment meets the requirements.
2. Install the product image.
3. Configure the Configuration Tool on Liberty.
4. Create a policy in the Configuration Tool.
5. Configure the System Data Engine.
6. Configure the Data Streamer.
7. Configure and start Logstash.
8. Start the components.

Preparation checklist

Before you install and configure the IBM Z Common Data Provider components, ensure that your environment meets the system and security requirements in this checklist.

System requirements

Prepare your system according to the following requirements.

z/OS system requirements

You must run the IBM Z Common Data Provider in each z/OS logical partition (LPAR) from which you want to gather z/OS operational data.

IBM Z Common Data Provider must be run with the following software:

- IBM z/OS V2.2, or later (product number 5655-ZOS)

For z/OS V2.2 only: If you use IBM z/OS V2.2, the following software is also required on the system where configuration is done:

For z/OS V2.2

IBM z/OS Management Facility V2.2 (product number 5650-ZOS), with APAR/PTF PI52426/UI36315

- The following Java™ library:
 - IBM 64-bit SDK for z/OS Java Technology Edition V8 (product number 5655-DGH)
- On an IBM z/OS V2.2 system, to collect System Management Facilities (SMF) data from SMF in-memory resources, you must apply APAR OA49263.

SMF requirements

This scenario uses in-memory resource, therefore, the SMF must be running in log stream recording mode instead of data set recording mode.

Configuration Tool browser requirements

You can access the Liberty-based Configuration Tool by using the following web browsers:

- Apple Safari
- Google Chrome
- Microsoft Internet Explorer 10 or later
- Mozilla Firefox

Working directory requirements

When you configure the Configuration Tool and Data Streamer, you must define a working directory for each component. Prepare the working directories for these components according to the following guidelines:

- The working directory must be readable and writable by the user ID that runs the component.
- To avoid possible conflicts, the working directory for each component must be unique.

Data Streamer port definition

When you configure the Data Streamer, you define the port number on which the Data Streamer listens for data from the data gatherers. Prepare an available port before you configure the Data Streamer. The default port is 51401.

Security requirements

Prepare user IDs with requested authorities before you configure the components.

Configuration Tool

- You must have a user ID with the RACF SPECIAL authority to run a script that creates the users and groups for the Liberty-based Configuration Tool.
- You must have a user ID that has the UID 0 attribute to run the setup script that creates the working directory and the Liberty server directory.

Data Streamer

The user ID that is associated with the Data Streamer started task must have the appropriate authority to access the IBM Z Common Data Provider program files, which include the installation files and the policy file. It must also have read and execute permissions to the Java libraries in the UNIX System Services file system.

System Data Engine

The user ID that is associated with the System Data Engine started task must have authority to read the SMF in-memory resource.

For example, if you are using the RACF® as your System Authorization Facility (SAF) product, you must give the System Data Engine user ID read authority to the profile that you set up to secure your SMF in-memory resource. In the following examples, *IFASMF.resource* represents the name of the SMF in-memory resource that is being used to gather SMF records, and *userid* represents the System Data Engine user ID.

In-memory resource example

```
PERMIT IFA.IFASMF.resource CLASS(FACILITY) ACCESS(READ) ID(userid)
```

Installing Z Common Data Provider

Install IBM Z Common Data Provider by using SMP/E. For installation instructions, see the IBM Z Common Data Provider Program Directory.

About this task

Install all IBM Z Common Data Provider components according to the instructions in the Program Directory.

- [IBM Z Common Data Provider Program Directory](#)

Because this scenario runs the Configuration Tool on IBM WebSphere Application Server for z/OS Liberty, you must install the IBM Z Common Data Provider Embedded Liberty.

Table 1 on page 4 lists the target libraries.

Table 1. Target libraries for IBM Z Common Data Provider components	
Component	Target library
Configuration Tool	<ul style="list-style-type: none">• /usr/lpp/IBM/zcdp/v2r1m0/UI <p>For the following libraries, customize the high-level qualifier (.hlq) according to site requirements.</p> <ul style="list-style-type: none">• hlq.SHBODEFS• hlq.SHBOSAMP

Table 1. Target libraries for IBM Z Common Data Provider components (continued)	
Component	Target library
Data Streamer	<ul style="list-style-type: none"> • /usr/lpp/IBM/zcdp/v2r1m0/DS <p>For the following library, customize the high-level qualifier (.hlq) according to site requirements.</p> <ul style="list-style-type: none"> • hlq.SHBOSAMP
Log Forwarder	<ul style="list-style-type: none"> • /usr/lpp/IBM/zcdp/v2r1m0/LF <p>For the following libraries, customize the high-level qualifier (.hlq) according to site requirements.</p> <ul style="list-style-type: none"> • hlq.SHBOCLST • hlq.SHBOLPA
System Data Engine	<p>For the following libraries, customize the high-level qualifier (.hlq) according to site requirements.</p> <ul style="list-style-type: none"> • hlq.SHBOLLST • hlq.SHBOLOAD
Common dependencies	<ul style="list-style-type: none"> • /usr/lpp/IBM/zcdp/v2r1m0/DEPS

Configuring the Configuration Tool on Liberty

In this scenario, the IBM Z Common Data Provider Configuration Tool will be deployed as an application to IBM WebSphere Application Server for z/OS Liberty. The other scenario where the Configuration Tool is deployed as a plug-in for IBM z/OS Management Facility (z/OSMF) is not covered in this guide.

Configuring user IDs, group IDs, and security product

You must create user IDs and group IDs with necessary permissions to run the IBM Z Common Data Provider Configuration Tool.

About this task

A default properties file /usr/lpp/IBM/zcdp/v2r1m0/UI/LIB/cdpui.properties is provided with default user IDs and group IDs to run the Configuration Tool. You can run the defracf.cmd script to change the default values. The new values are saved in /var/cdp-uiconfig/cdpui.properties for the savingpolicy.sh script to use in the next task. If you are using RACF as your SAF product, you can allow the script to run necessary RACF commands to create the IDs and permissions. If you do not use RACF, you can exit the script after verifying or changing the values and continue with the configuration.

In this scenario, RACF is used as the SAF product, and we will allow the script to create the IDs and permissions.

To run the defracf.cmd script, you must be logged in to the z/OS system with a user ID that has the RACF SPECIAL authority.

Procedure

1. Run the following script under UNIX System Services. Only the default z/OS shell is supported.

```
/usr/lpp/IBM/zcdp/v2r1m0/UI/LIB/defracf.cmd
```

2. Verify the default user IDs and group IDs.

STC_USRID

The user ID that is assigned to the Configuration Tool server started task procedure. The default value is HBOSTCID.

STC_GROUP

The group that contains **STC_USRID**. The default value is HBOSTCGP.

AUTHORIZED_GROUP

The group that is granted the permission of logging in and using the Configuration Tool. The default value is HB0USRGP.

GUEST_USER

The user ID that is used by Liberty for accessing the Configuration Tool login page. The default value is HB0GUEST.

GUEST_GROUP

The group that contains **GUEST_USER**. The default value is HB0UNGRP.

AUTHORIZED_USER

The user ID that is granted the permission of logging in and using the Configuration Tool. The default value is HB0USER.

Remember: You must specify an existing user for this parameter. If you don't specify any value for this parameter, no user is able to access the Configuration Tool.

AUTOID

Determines whether the UID and GID are automatically assigned. The default value is OFF, and you can set values for the following parameters. Make sure that the UIDs and GIDs that you specify meet the requirements of your environment. If the UIDs and GIDs are not accepted by your security product, the Configuration Tool cannot be installed successfully. If you don't specify any value for these parameters, default values will take effect.

STC_USRID_UID

The UID for **STC_USRID**.

STC_GROUP_GID

The GID for **STC_GROUP**.

AUTHORIZED_GROUP_GID

The GID for **AUTHORIZED_GROUP**.

GUEST_USER_UID

The UID for **GUEST_USER**.

GUEST_GROUP_GID

The GID for **GUEST_GROUP**.

If automatic assignment of UID and GID is enabled on your environment, you can change the value of this parameter to ON to have required UIDs and GIDs automatically assigned by the system. In this case, skip the UID and GID parameters that are listed previously.

3. When you are prompted to choose exit or go, enter GO.

Check the output from the RACF commands in the `/var/cdp-uiconfig/defracf.log` file and verify that all commands are successfully issued by the script.

- There should be no RACF error messages from the UNIX System Services issued to the terminal after the script finishes running.
- If you see the messages ICH10006I, ICH06011I, and IRRD175I indicating that RACLISTED PROFILES must be refreshed before they are effective, and a message "All related RACLIST CLASS

are refreshed successfully" after the script finishes running, it means that the RACLISTED PROFILES are refreshed by the script and are effective.

- Message ICH10102I that says BBG.AUTHMOD.BBGZSAFM, and BBG.AUTHMOD.BBGZSAFM.SAFCRED are already defined, can be safely ignored. These profiles are shared with other Liberty Angel Servers, and they might be defined by a Liberty Angel Server that was created before.

Setting up a Liberty server directory and a working directory for the Configuration Tool

You must set up a Liberty server directory to contain the configuration of the Configuration Tool server, and a working directory to store the policy definition files. A setup script (`savingpolicy.sh`) is provided to automate this process. Default values are used in this scenario.

About this task

User ID criteria for running the setup script

To run the setup script, you must be logged in to the z/OS system with a user ID that has the UID 0 attribute.

Procedure

1. Enter the directory that contains the setup script.

```
cd /usr/lpp/IBM/zcdp/v2r1m0/UI/LIB/
```

2. Run the following command to start the setup script. Only the default z/OS shell is supported.

```
savingpolicy.sh
```

3. Follow the prompts of the script to provide necessary values.

To accept the default value that is shown in the parentheses, enter a blank value.

- a) When you are prompted to choose if you are deploying the Configuration Tool on z/OSMF or Liberty, enter 2 to select Liberty.
- b) When you are prompted to specify the full path of directory where the Configuration Tool server is installed, accept the default value.
 - Ensure that the directory is readable by the **STC_USRID** and **STC_GROUP** that are specified in the `/var/cdp-uiconfig/cdpui.properties` file that is created in [“Configuring user IDs, group IDs, and security product”](#) on page 5.
- c) When you are prompted to specify the Java home directory, accept the default value, or specify the Java home directory for your system.
- d) When you are prompted to specify the Configuration Tool Source Script Directory, accept the default value.

Results

The directory `config_tool_server_install_dir/servers/cdp_ui_server` is created as the Liberty server directory, and the directory `config_tool_server_install_dir/cdpConfig` is created as the working directory, where `config_tool_server_install_dir` is the value that is specified in [“3.b”](#) on page 7.

Creating the started tasks of the Configuration Tool server and its angel server

Before you can start the IBM Z Common Data Provider Configuration Tool server, you must create the started tasks for the Configuration Tool server and its angel server by copying the sample procedures into a user procedure library, and updating the copies.

Procedure

To create the started tasks, complete the following steps:

1. Copy the procedure HBOCFGT from *hlq.SHBOSAMP* to a user procedure library.
2. Update the procedure HBOCFGT.
 - Change the value of the variable INSTDIR to the path where the WebSphere Application Server for z/OS Liberty is installed.
 - Change the value of the variable USERDIR to the path where the Configuration Tool server is installed. This value is specified in [Step 3.b](#) in Setting up a Liberty server directory and a working directory for the Configuration Tool.
3. Copy the procedure HBOCFGGA from *hlq.SHBOSAMP* to a user procedure library.
4. Change the value of the variable WLPDIR in HBOCFGGA to the path where the WebSphere Application Server for z/OS Liberty is installed.

Starting the Configuration Tool server

Start the angel server and the Configuration Tool server before you can access the Configuration Tool in a web browser.

Procedure

1. Start the angel server for the Configuration Tool server by running the following z/OS system console command:

```
START HBOCFGGA
```

2. Verify that the angel server starts successfully.
You see the following message in the job log if the server starts successfully:

```
CWWKB0069I INITIALIZATION IS COMPLETE FOR THE HBOCFGGA ANGEL PROCESS
```

3. Start the Configuration Tool server by running the following console command:

```
START HBOCFGT
```

4. Verify that the Configuration Tool server starts successfully.
You see the following messages in the job log if the server starts successfully:

```
CWWKF0011I: The server cdp_ui_server is ready to run a smarter planet.  
CWWKT0016I: Web application available (default_host): https://hostname:port/cdp/
```

5. Verify that the Configuration Tool server and the angel server are connected successfully.

In the log file `/var/local/CDPServer/servers/cdp_ui_server/logs/messages.log`, you can see the following message:

```
CWWKB0103I: Authorized service group KERNEL is available.
```

If the SAF authorized user registry services and SAF authorization services (SAFCRED) are enabled, the following message is in the log file:

```
CWWKB0103I: Authorized service group SAFCRED is available
```

Tip: The user ID that can read the file `messages.log` must belong to the **STC_GROUP** user group that is specified in the `/var/cdp-uiconfig/cdpui.properties` file that is created in [“Configuring user IDs, group IDs, and security product”](#) on page 5.

Accessing the Configuration Tool

After the Configuration Tool server and the angel server are started, you can access the Configuration Tool in a web browser.

Procedure

1. Access the following URL in a web browser:

```
https://HostName:port/cdp
```

HostName is the host name of your server that runs the Configuration Tool server. *port* is the port number that is used by the Configuration Tool server. The default port number is 17977.

2. Log in as the user ID that is connected to the user group that is defined by the **AUTHORIZED_GROUP** parameter in the `/var/cdp-uiconfig/cdpui.properties` file that is created in [“Configuring user IDs, group IDs, and security product”](#) on page 5.

Results

The **"Common Data Provider"** tab opens. Predefined policies are listed.


Creating a policy to stream SMF data

From IBM Z Common Data Provider Configuration Tool, you can create a policy to stream SMF data to various subscribers.

About this task

In the policy, select SMF_30 and SMF_110_1_KPI data streams, and add a subscriber with the **CDP on Elasticsearch via Logstash** protocol. No transforms will be applied.

Procedure


1. From the Configuration Tool, click the **Create a new policy** box.
2. In the resulting **Policy Profile Edit** window, type the required policy name and, optionally, a policy description.
3. Click the **Add Data Stream** icon  **DATA STREAM**.

The **"Select data stream"** window is shown with a list of categorized data streams. You can expand the categories to view the possible data streams that you can define for this policy.

4. Select the following data streams and click **Select**.

- Expand **Starter Sets** > **Common Data Provider for z System** > **General z/OS system monitoring** and select **SMF_030**.
- Expand **Starter Sets** > **Common Data Provider for z System** > **CICS** and select **SMF_110_1_KPI**.

Each data stream that you choose is shown as a node in the graph.

5. Click the **Subscribe**  icon on the SMF_30 data stream node, the "**Subscribe to a data stream**" window opens where you can define a new subscriber by completing the following steps.

a) Click the **Add Subscriber** icon  **ADD SUBSCRIBER**.

b) In the resulting "**Add subscriber**" window, update the associated configuration values.

You can update the following values in the "**Add subscriber**" window:

Name

The name of the subscriber.

Description

An optional description for the subscriber.

Protocol

The streaming protocol that the Data Streamer uses to send data to the subscriber. Select **CDP on Elasticsearch via Logstash**.

Host

The host name or IP address of the subscriber.

Port

The port on which the subscriber listens for data from the Data Streamer.


Auto-Qualify

A specification of whether to prepend system names and sysplex names to data source names in the data streams that are sent to the subscriber. The data source name is the value of the **dataSourceName** field in the data stream configuration. Accept the default value **None**.

c) Click **OK** to save the subscriber.

6. In the "**Subscribe to a data stream**" window, select the new subscriber, and click **Update Subscriptions**.

The subscriber that you choose is then shown on the graph.

7. Click the **Subscribe**  icon on the SMF_110_1_KPI data stream node. Select the same subscriber and click **Update Subscriptions**.

8. Click the **SYSTEM DATA ENGINE** button to set the CDP Concatenation value for your System Data Engine environment.

CDP Concatenation

This value must be the name of the SHBODEFS data set that is installed with IBM Z Common Data Provider in your environment. This data set is also referenced in the `concat.json` file, which is in the working directory for the IBM Z Common Data Provider Configuration Tool.

9. To save the policy, click **Save**.

Results

The following files are created in the Configuration Tool working directory:

- `policy_name.policy`
- `policy_name.layout`
- `policy_name.sde`
- `policy_name.zlf.conf`
- `policy_name.config.properties`

Creating the System Data Engine started task for streaming SMF data

To have the IBM Z Common Data Provider System Data Engine stream SMF data to the Data Streamer, you must create the started task for the System Data Engine by copying the sample procedure HBOSMF in the hlq.SHBOSAMP library, and updating the copy. In this scenario, all default values are accepted.

Before you begin

For the System Data Engine to gather System Management Facilities (SMF) data, the SHBOLoad library must be authorized with APF.

To authorize the SHBOLoad library, a library name and volume ID must be in the list of authorized libraries in the PROGxx member of the SYS1.PARMLIB library. To do that, add the following statement to a PROGxx member:

```
APF ADD DSNAME(hlq.SHBOLoad) VOLUME(volname)
```

Procedure

To create the started task, complete the following steps:

1. Copy the procedure HBOSMF in the hlq.SHBOSAMP library to a user procedure library.
2. Update the high-level qualifier to the one for your IBM Z Common Data Provider target libraries that were installed by using SMP/E.
3. Update the port value for IBM_UPDATE_TARGET to specify the TCP/IP port that is configured for the Data Streamer.
The default port for Data Streamer is 51401.
4. Replace the value /u/userid/cdpConfig/Sample1.sde with the policy_name.sde file path and name that is created in [“Creating a policy to stream SMF data”](#) on page 9.
5. Update the value for IBM_RESOURCE with the SMF in-memory resource name to specify the source from which SMF data is to be gathered.

Important: If the in-memory method is not allowed in your environment, consider using the SMF user exit method by specifying EXIT for this parameter. For more information about how to install the SMF user exit, see [Installing the SMF user exit](#) in the IBM Knowledge Center.

6. Verify that the user ID that is associated with the System Data Engine started task has the required authorities, as described in [Preparation checklist](#).
7. Update the SAF product that is protecting your system, such as the Resource Access Control Facility (RACF), to permit the System Data Engine started task to run in your environment.

Configuring the Data Streamer

The IBM Z Common Data Provider Data Streamer streams operational data to configured subscribers in the appropriate format. It receives the data from the data gatherers like the System Data Engine, and sends the data to the subscriber.

About this task

To configure the Data Streamer, you must create the Data Streamer started task by copying the sample procedure HBODSPRO in the hlq.SHBOSAMP library, and updating the copy. Only parameters that are related to basic functions are updated in this scenario.

The user ID that is associated with the Data Streamer started task must have the appropriate authority to access the IBM Z Common Data Provider program files, which include the installation files and the policy file. It must also have read/execute permissions to the Java libraries in the UNIX System Services file system.

Procedure

To create the started task, complete the following steps:

1. Copy the procedure HBODSPRO in the `hlq.SHBOSAMP` library to a user procedure library.
2. In your copy of the procedure HBODSPRO, customize the following parameter values for your environment:

/usr/lpp/IBM/zcdp/v2r1m0/DS/LIB

Replace this value with the directory where the Data Streamer is installed in your environment. This directory contains the `startup.sh` script for the Data Streamer.

/etc/cdpConfig/myPolicy.policy

Replace this value with the `policy_name.policy` file path and name that is created in [“Creating a policy to stream SMF data”](#) on page 9.

nnnnn

Replace this value with the port number on which the Data Streamer listens for data from the data gatherers. The default port on which the Data Streamer listens for data is 51401.

3. In your copy of the procedure HBODSPRO, set the following environment variables for your environment:

JAVA_HOME

Specify the Java installation directory.

CDP_HOME

Specify the location of the Data Streamer working directory. The Data Streamer working directory contains files that are created and used during the operation of the Data Streamer, including the Data Streamer truststore and file buffers.

Guidelines for the working directory

Use the following guidelines to help you decide which directory to use as the working directory:

- The directory must be readable and writable by the user ID that runs the Data Streamer.
- To avoid possible conflicts, do not use a directory that is defined as the Configuration Tool working directory.

Important: Do not update, delete, or move the files in the `CDP_HOME` directory.

4. Update your security software, such as the Resource Access Control Facility (RACF), to permit the Data Streamer started task to run in your environment.

Configuring and starting Logstash

To send data from IBM Z Common Data Provider to Elasticsearch, configure Logstash by using the Logstash configuration files that are provided by IBM Z Common Data Provider.

About this task

The IBM Z Common Data Provider Elasticsearch ingestion kit contains the Logstash configuration files that are provided by IBM Z Common Data Provider.

Procedure

1. From the IBM Z Common Data Provider `/usr/lpp/IBM/zcdp/v2r1m0/DS/LIB` directory, download the IBM Z Common Data Provider Elasticsearch ingestion kit, which is in the `ibm_cdpz_ELK.tar.gz` file, in binary mode.

2. Extract the Elasticsearch ingestion kit and make sure the Logstash configuration files are available.
3. Create a new directory under the Logstash installation directory and copy the Logstash configuration files that you need for your environment to the new directory.

Table 2 on page 13 indicates the prefixes that are used in the file names for the Logstash configuration files in the IBM Z Common Data Provider Elasticsearch ingestion kit. The file name prefix is an indication of the configuration file content.

<i>Table 2. Mapping of the prefix that is used in a Logstash configuration file name to the content of the file</i>	
Prefix in file name of Logstash configuration file	Content of configuration file with this prefix
B_	Input stage
E_	Preparation stage
H_	Field name annotation stage
N_	Timestamp resolution stage
Q_	Output stage

The following descriptions further explain the Logstash configuration files in the IBM Z Common Data Provider Elasticsearch ingestion kit:

B_CDPz_Input.lsh file

This file contains the input stage that specifies the TCP/IP port on which Logstash listens for data from the Data Streamer. Copy this file to your Logstash configuration directory. You might need to edit the port number after you copy the file.

E_CDPz_Index.lsh file

This file contains the preparation stage. Copy this file to your Logstash configuration directory.

Files with H_ prefix in file name

Each of these files contains a unique field name annotation stage that maps to a unique data stream that IBM Z Common Data Provider can send to Logstash. To your Logstash configuration directory, copy the H_ files for only the data streams that you want to send to Elasticsearch.

Files with N_ prefix in file name

Each of these files contains a unique timestamp resolution stage that maps to a unique data stream that IBM Z Common Data Provider can send to Logstash. To your Logstash configuration directory, copy the N_ files for only the data streams that you want to send to Elasticsearch.

Q_CDPz_Elastic.lsh file

This file contains an output stage that sends all records to a single Elasticsearch server. Copy this file to your Logstash configuration directory.

After you copy the file, edit it to add the name of the host to which the stage is sending the indexing call. The default name is localhost, which indexes the data on the server that is running the ingestion processing. Change the value of the **hosts** parameter rather than the value of the **index** parameter. The **index** value is assigned during ingestion so that the data for each source type is sent to a different index. The host determines the Elasticsearch farm in which the data is indexed. The index determines the index in which the data is held.

To split data according to sysplex, you can use the [sysplex] field in an if statement that surrounds an appropriate Elasticsearch output stage.

4. In the script for starting Logstash, specify the directory that you created in step “3” on page 13.
5. Start Logstash and Elasticsearch.

Starting components

The best practice for the order in which to start the components is from the destination end to the source end. Because you have started Logstash in the previous task, continue to start the Data Streamer, and then start the System Data Engine.

Procedure

1. Start the Data Streamer by running the following z/OS console command:

```
START HBODSPRO
```

2. After the Data Streamer is started successfully, run the following z/OS console command to start the System Data Engine:

```
START HBOSMF
```

Results

If the activation is successful, IBM Z Common Data Provider starts sending data to Elasticsearch.

Stopping components

The best practice for the order in which to stop the components is from the source end to the destination end. Stop the System Data Engine first, and then stop the Data Streamer.

Procedure

1. Run the following z/OS console command to stop the System Data Engine:

```
STOP HBOSMF
```

2. After the System Data Engine is stopped successfully, stop the Data Streamer by running the following z/OS console command:

```
STOP HBODSPRO
```