IBM Z Operations Analytics

User Guide

Version 3 Release 2

IBM Z Operations Analytics

User Guide

Version 3 Release 2

Figures

Tables

Contents

Z Operations Analytics overview

IBM® Z Operations Analytics is available on three platforms: IBM Operations Analytics - Log Analysis, Elastic Stack, and Splunk. The user interface and available functions can vary depending on the platform.

IBM Z Operations Analytics can provide IT operational insights for multiple domains of interest, including z/OS® system, databases, messaging, networks, security, transactions, or web servers. It provides the function for analyzing each unique type of z/OS operations data and producing the associated insights.

IBM Common Data Provider for z Systems® Version 1.1.0

IBM Z Operations Analytics includes IBM Common Data Provider for z Systems V1.1.0.

IBM Common Data Provider for z Systems provides the infrastructure for accessing IT operational data from z/OS systems and streaming it to the analytics platform of your choice 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.

For more information about IBM Common Data Provider for z Systems, see the IBM Common Data Provider for z Systems V1.1.0 documentation.

Planning for configuration of IBM Common Data Provider for z Systems

Verify that your environment meets the IBM Common Data Provider for z Systems system requirements. Also, determine the sources from which you want IBM Common Data Provider for z Systems to gather data so that you can correctly configure IBM Common Data Provider for z Systems to stream the operational data to IBM Z Operations Analytics.

About this task

For information about the IBM Common Data Provider for z Systems system requirements, see Planning to use IBM Common Data Provider for z Systems in the IBM Common Data Provider for z Systems V1.1.0 documentation.

For information about the correlation between the following items, see "Operational insights reference" on page 105:

- The data sources that contribute to the different types of operational insights, and the configuration that you must do in the IBM Common Data Provider for z Systems Configuration Tool to send that data to IBM Z Operations Analytics
- The dashboards that represent the operational data from the data sources
- The predefined searches for searching the operational data

Enabling support for SMF data destined for IBM Z Operations Analytics

If you plan to send System Management Facilities (SMF) data to IBM Z Operations Analytics, you must enable the IBM Common Data Provider for z Systems Configuration Tool to support the SMF record types that are destined for IBM Z Operations Analytics.

About this task

You must complete this task before you create any policies that define SMF data streams with IBM Z Operations Analytics as the target destination.

Procedure

1. Depending on your platform, copy the specified file from the IBM Z Operations Analytics installation directory to the working directory for the IBM Common Data Provider for z Systems Configuration Tool.

   If you are sending SMF data to IBM Z Operations Analytics on both the IBM Operations Analytics - Log Analysis platform AND the Elastic Stack or Splunk platform

   If you want to send SMF data from a single IBM Common Data Provider for z Systems instance to IBM Z Operations Analytics on both the IBM Operations Analytics - Log Analysis platform and the Elastic Stack or Splunk platform, copy the `glaELKSplunk.streams.json` file.

Tip: From UNIX System Services, use the following sample command, with the appropriate values for your environment, to copy the glaELKSplunk.streams.json file:

```
cp /usr/lpp/IBM/zscala/V3R2/samples/glaELKSplunk.streams.json
   config_tool_workdir
```

The following SMF record types are not supported on the IBM Operations Analytics - Log Analysis platform:

- SMF100_1
- SMF101_SUMMARY
- SMF110_1_SUMMARY

Table 1 indicates the differences between the data source types that are defined in the glasmf.streams.json file and the glaELKSplunk.streams.json file. These differences are relevant only if you were previously using the glasmf.streams.json file and must therefore migrate data to the glaELKSplunk.streams.json file. These differences do not affect your existing data, but you might need to update any custom dashboards or searches to reflect the new data source types.

Table 1. Differences between the data source types that are defined in the glasmf.streams.json file and the glaELKSplunk.streams.json file

Data source name	Data source type in **glasmf.streams.json**	Data source type in **glaELKSplunk.streams.json**
SMF80_COMMAND	zOS-SMF80	zOS-SMF80_COMMAND
SMF80_LOGON	zOS-SMF80	zOS-SMF80_LOGON
SMF80_OPERATION	zOS-SMF80	zOS-SMF80_OPERATION
SMF80_OMVS_RES_1	zOS-SMF80	zOS-SMF80_OMVS_RES_1
SMF80_OMVS_RES_2	zOS-SMF80	zOS-SMF80_OMVS_RES_2
SMF80_OMVS_SEC_1	zOS-SMF80	zOS-SMF80_OMVS_SEC_1
SMF80_OMVS_SEC_2	zOS-SMF80	zOS-SMF80_OMVS_SEC_2
SMF80_RESOURCE	zOS-SMF80	zOS-SMF80_RESOURCE
SMF120_REQAPPL	zOS-SMF120	zOS-SMF120_REQAPPL
SMF120_REQCONT	zOS-SMF120	zOS-SMF120_REQCONT

If you are sending SMF data to IBM Z Operations Analytics on only the IBM Operations Analytics - Log Analysis platform

Copy the glasmf.streams.json file.

Tip: From UNIX System Services, use the following sample command, with the appropriate values for your environment, to copy the glasmf.streams.json file:

```
cp /usr/lpp/IBM/zscala/V3R2/samples/glasmf.streams.json
   config_tool_workdir
```

If you are sending SMF data to IBM Z Operations Analytics on either the Elastic Stack or Splunk platform

Copy the glaELKSplunk.streams.json file.

Tip: From UNIX System Services, use the following sample command, with the appropriate values for your environment, to copy the glaELKSplunk.streams.json file:

```
cp /usr/lpp/IBM/zscala/V3R2/samples/glaELKSplunk.streams.json
   config_tool_workdir
```

2. Verify that the `concats.json` file is in the working directory for the IBM Common Data Provider for z Systems Configuration Tool. Also, verify that the file contains the appropriate values for your installation, as described in the following example:

Line in `concats.json` file	Explanation
`"CDP" : "CDP.SHBODEFS"`	A reference to the `SHBODEFS` data set that is installed with IBM Common Data Provider for z Systems. The value in quotation marks (in this example, `CDP.SHBODEFS`) must be the data set name for your installation.
`"IZOA" : "ZSCALA.V3R2M0.SGLADEFS"`	A reference to the `SGLASAMP` data set that is installed with IBM Z Operations Analytics. The value in quotation marks (in this example, `ZSCALA.V3R2M0.SGLADEFS`) must be the data set name for your installation.

Results

In the IBM Common Data Provider for z Systems Configuration Tool, you can now create policies that define SMF data streams with IBM Z Operations Analytics as the target destination. In the Configuration Tool, different SMF data streams are listed under the following two categories:

- **Common Data Provider for z Systems**
- **IBM Z Operations Analytics**

The SMF data streams in each category are similar, but not the same. For example, the **Common Data Provider for z Systems** category includes the **SMF_030** data stream, and the **IBM Z Operations Analytics** category includes the **SMF30** data stream.

The difference between the categories is that the **IBM Z Operations Analytics** category includes the custom SMF data streams that are provided by IBM Z Operations Analytics. These data streams condense the data flows to only the data that is required by the IBM Z Operations Analytics. IBM Z Operations Analytics does not use data from the SMF data streams in the **Common Data Provider for z Systems** category.

To use the dashboards and predefined searches that are provided by IBM Z Operations Analytics, you must configure the SMF data streams in the **IBM Z Operations Analytics** category.

Requirements for gathering WebSphere Application Server for z/OS log data

If you plan to gather log data for WebSphere® Application Server for z/OS, you must determine the application servers from which to gather log data.

For each of the application servers, you must then determine where to retrieve the log data.

On the IBM Operations Analytics - Log Analysis platform, if the application server is configured to use High Performance Extensible Logging (HPEL) mode, a best practice is to retrieve the log data by using the HPEL API.

If the application server is configured to use basic logging, the log data is retrieved from JES job logs, z/OS UNIX log files, or both, depending on how the server is configured.

On the IBM Operations Analytics - Log Analysis platform, if the application server is configured to use HPEL mode

On the IBM Operations Analytics - Log Analysis platform, for each application server that is configured to use HPEL mode, complete the following steps:

1. Use the WebSphere Integrated Solutions Console to determine the HPEL logging and trace directories. Logging and trace information is typically in the same directory, but it can be configured to be in different directories.
2. Determine whether logging data only, trace data only, or both, is gathered.

   Logging data includes data from the `java.util.logging` package (the level DETAIL and higher), the `System.out` stream, and the `System.err` stream.

   Trace data includes data from the `java.util.logging` package (the level DETAIL and lower).
3. Ensure that the user ID that is associated with the IBM Common Data Provider for z Systems procedure is authorized to read the HPEL logging and trace directories and files.

If the application server is configured to log to JES job logs

For each application server that is configured to log to JES job logs, complete the following steps:

1. Determine which regions of the application server to gather log data from.

Region	Focus area
Controller region	Inbound and outbound communication, security, and transaction control
Servant region	Most of the application server components
Adjunct region	Internal messaging

2. For each application server region, determine the job name.

Region	Typical job name
Controller region	Server short name
Servant region	Job name for the controller region with an “S” appended
Adjunct region	Job name for the controller region with an “A” appended

3. For each application server region, determine whether to gather SYSOUT data, SYSPRINT data, or both types of data.

   SYSOUT data includes Java™ logs (non-trace levels) and native message logs.

   SYSPRINT data includes Java logs (with trace levels) and native trace.

If the application server is configured to log to z/OS UNIX log files

For each application server that is configured to log to z/OS UNIX log files, complete the following steps:

1. Determine which regions of the application server to gather log data from.

Region	Focus area
Controller region	Inbound and outbound communication, security, and transaction control
Servant region	Most of the application server components
Adjunct region	Internal messaging

2. For each application server region, determine whether to gather SYSOUT data, SYSPRINT data, or both types of data.

   SYSOUT data includes Java logs (non-trace levels) and native message logs.

   SYSPRINT data includes Java logs (with trace levels) and native trace.
3. Determine the file path of each z/OS UNIX log file to be gathered.
4. For each z/OS UNIX file to be gathered, determine whether the path name is constant or varies. The path name varies in the following situations:
   - When date and time substitution is used in the file name that is specified in the data definition (DD) statement. The use of date and time substitution causes a new log file to be created for each server instance.
   - When the WebSphere environment variable *redirect_server_output_dir* is used to redirect output to files. The use of this variable causes a new log file to be created for each server instance. It also gives you the capability to use the **ROLL_LOGS** parameter of the modify command to create a new set of log files.

   For more information about file paths for rolling logs, see the IBM Common Data Provider for z Systems V1.1.0 documentation.
5. Ensure that the user ID that is associated with the IBM Common Data Provider for z Systems Log Forwarder procedure is authorized to read the z/OS UNIX log files.

Determination of time zone information for z/OS log records

IBM Z Operations Analytics determines time zone information for z/OS log records. The time zone information varies depending on the source of the log data.

If IBM Z Operations Analytics cannot determine the time zone information for each log record, it might not identify the correct relative placement in time for log records from different sources.

The following information describes how time zone is determined for z/OS log records, depending on the source of the log data:

z/OS SYSLOG data
The time stamps for z/OS SYSLOG messages include time zone information.

CICS® Transaction Server for z/OS log data
The time stamps for CICS Transaction Server for z/OS EYULOG and MSGUSR log messages do not include time zone information. These time stamps are based on the local z/OS system time zone.

NetView® for z/OS messages

The time stamps for the NetView for z/OS messages that are provided by the NetView message provider do not include time zone information. These time stamps are based on Coordinated Universal Time (UTC).

SMF data

The time stamps for SMF messages do not include time zone information. These time stamps are based on the local z/OS system time zone.

UNIX System Services system log (`syslogd`) messages

The time stamps for `syslogd` messages do not include time zone information. These time stamps are based on the local z/OS system time zone.

WebSphere Application Server for z/OS log data

The time stamps for the WebSphere Application Server for z/OS log messages do not include time zone information, with the following exceptions:

- WebSphere Application Server for z/OS log messages data that is produced in distributed format contains time stamps with time zone information.
- **On the IBM Operations Analytics - Log Analysis platform only**, WebSphere Application Server for z/OS log messages data that is retrieved from High Performance Extensible Logging (HPEL) contains time stamps with time zone information.

By default, time stamps in the WebSphere Application Server for z/OS logs are based on UTC. However, if the WebSphere Application Server for z/OS variable *ras_time_local* is set to 1, time stamps are based on the local z/OS system time zone. WebSphere Application Server for z/OS variables can be set at the cell, cluster, node, or server scope level.

For each WebSphere Application Server for z/OS data set that is written to a JES job log or z/OS UNIX log file, determine whether the time stamps in the log data are based on the local z/OS system time zone or on UTC.

Z Operations Analytics on the Log Analysis platform

IBM Z Operations Analytics extends the capabilities of IBM Operations Analytics - Log Analysis to help you more quickly identify, isolate, and resolve problems in a z/OS-based IT operations environment.

IBM Z Operations Analytics provides the following capabilities:

- Capability to gather z/OS logs across the System z® enterprise and to forward them to the IBM Operations Analytics - Log Analysis server for analysis
- Capability to index, search, and analyze application, middleware, and infrastructure log data across the System z enterprise
- Capability to quickly search and visualize errors across thousands of log records
- Expert advice that is based on linking search results to available troubleshooting information, such as best practices and previously documented solutions
- Continuous streaming of z/OS logs

Component overview

In addition to IBM Common Data Provider for z Systems, IBM Z Operations Analytics includes IBM Operations Analytics - Log Analysis Version 1.3.5 and a IBM z Advanced Workload Analysis Reporter (IBM zAware) Version 3.1 feature.

IBM Operations Analytics - Log Analysis

IBM Operations Analytics - Log Analysis is a cross-platform solution for identifying and resolving problems in an IT operations environment.

In IBM Operations Analytics - Log Analysis, an *Insight Pack* is software that extends the capabilities of IBM Operations Analytics - Log Analysis to provide support for loading and analyzing data from sources that share common characteristics. Examples include log sources for a specific operating system or for a specific application, such as IBM WebSphere Application Server.

IBM Z Operations Analytics includes z/OS Insight Packs that extend IBM Operations Analytics - Log Analysis function to analyze the various types of z/OS log data.

IBM Z Operations Analytics also includes a version of IBM Operations Analytics - Log Analysis.

IBM z Advanced Workload Analysis Reporter (IBM zAware)

IBM zAware monitors software and models normal system behavior to detect and diagnose anomalies in z/OS and Linux on z Systems environments. It creates a model of normal system behavior that is based on prior system data, and it uses pattern recognition techniques to identify unexpected messages from the systems that it is monitoring. This analysis of events provides near real-time detection of anomalies that you can easily view through a GUI, which you can also use to determine the cause of past or current anomalies. This early detection can help IT personnel correct problems before the problems affect system processing.

IBM Z Operations Analytics includes the separately installable IBM zAware feature.

Flow of source data

Figure 1 illustrates the flow of data among the following primary components of IBM Z Operations Analytics on the IBM Operations Analytics - Log Analysis platform:

- IBM Common Data Provider for z Systems V1.1.0
- "z/OS Insight Packs" on page 11
- "IBM Operations Analytics - Log Analysis V1.3.5" on page 12

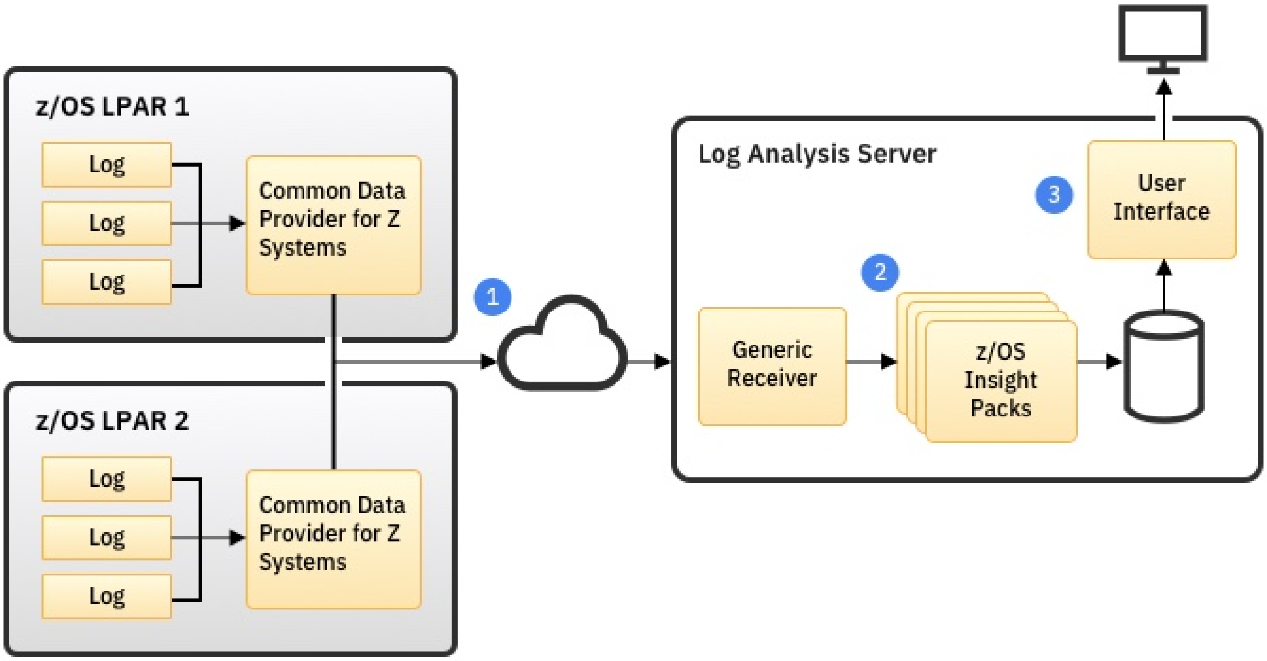

Figure 1. Flow of source data among IBM Z Operations Analytics components on the Log Analysis platform

The following steps describe the data flow among components, which is indicated by arrows in the illustration:

1. In each z/OS logical partition (LPAR), the IBM Common Data Provider for z Systems retrieves the data from the respective source and sends it to the IBM Operations Analytics - Log Analysis server.
2. The source data is processed by the respective z/OS Insight Pack. Insights are provided for data from the following source types:
   - z/OS system log (SYSLOG)
   - CICS Transaction Server for z/OS EYULOG or MSGUSR log data
   - Network data, such as data from UNIX System Services system log (`syslogd`) or z/OS Communications Server
   - NetView for z/OS message data
   - SMF data
   - WebSphere Application Server for z/OS logs that include SYSOUT, SYSPRINT, or HPEL log data
3. Users can see visualizations of the analyzed data in the Log Analysis user interface.

Insight is also provided for IBM zAware interval anomaly data. The flow of interval anomaly data includes the following steps:

1. The IBM zAware data gatherer of IBM Z Operations Analytics resides on the same distributed system as the Log Analysis server. It retrieves interval anomaly data from the IBM zAware Representational State Transfer (REST) API and provides it to the Log Analysis server.
2. The source data is processed by the z/OS SYSLOG Insight Pack.
3. Users can see visualizations of the analyzed data in the Log Analysis user interface.

z/OS Insight Packs

Each z/OS Insight Pack in IBM Z Operations Analytics extends IBM Operations Analytics - Log Analysis function to analyze a unique type of z/OS log data.

The following z/OS Insight Packs are available:

z/OS SYSLOG Insight Pack

This Insight Pack enables Log Analysis to ingest, and perform searches against, data that is retrieved from the following sources:

- z/OS SYSLOG console, which includes data from the following software:
  - IBM CICS Transaction Server for z/OS
  - IBM Db2® for z/OS
  - IBM IMS for z/OS
  - IBM MQ for z/OS
  - IBM Resource Access Control Facility (RACF®)
  - z/OS Communications Server
- CICS Transaction Server for z/OS EYULOG and MSGUSR log
- UNIX System Services system log (`syslogd`)
- IBM z Advanced Workload Analysis Reporter (IBM zAware)

Summary: Install this Insight Pack if you want to analyze any of the following data:

- z/OS SYSLOG console data
- CICS Transaction Server for z/OS EYULOG or MSGUSR log data
- UNIX System Services system log (`syslogd`)
- IBM zAware interval anomaly data
- SMF data

z/OS Network Insight Pack

This Insight Pack enables Log Analysis to perform searches against z/OS network data. It also enables Log Analysis to ingest, and perform searches against, IBM Tivoli® NetView for z/OS message data.

To analyze z/OS network data, you must use both the z/OS Network Insight Pack and the z/OS SYSLOG Insight Pack. The z/OS SYSLOG Insight Pack enables Log Analysis to ingest z/OS network data.

Summary: Install this Insight Pack if you want to analyze any of the following data:

- Network data, such as data from UNIX System Services system log (`syslogd`) or z/OS Communications Server
- NetView for z/OS message data

z/OS SMF Insight Pack

This Insight Pack enables Log Analysis to ingest, and perform searches against, data that is retrieved from IBM z/OS System Management Facilities (SMF).

Summary: Install this Insight Pack if you want to analyze SMF data.

WebSphere Application Server for z/OS Insight Pack

This Insight Pack enables Log Analysis to ingest, and perform searches against, logging data that is retrieved from IBM WebSphere Application Server for z/OS.

Summary: Install this Insight Pack if you want to analyze WebSphere Application Server for z/OS log data.

Table 2 lists each z/OS Insight Pack with its associated installation package file.

Table 2. Installation package files for z/OS Insight Packs

Insight Pack	**Installation package**
WebSphere Application Server for z/OS Insight Pack	`WASforzOSInsightPack_v3.2.0.0.zip`
z/OS Network Insight Pack **Important:** To analyze network data, install both the z/OS SYSLOG Insight Pack and this insight pack.	`zOSNetworkInsightPack_v3.2.0.0.zip`
z/OS SMF Insight Pack **Important:** To analyze SMF data, install both the z/OS SYSLOG Insight Pack and this insight pack.	`SMFforzOSInsightPack_v3.2.0.0.zip`
z/OS SYSLOG Insight Pack	`SYSLOGforzOSInsightPack_v3.2.0.0.zip`

IBM Operations Analytics - Log Analysis V1.3.5

IBM Operations Analytics - Log Analysis Version 1.3.5 is used with the z/OS Insight Packs to provide support for analyzing z/OS log data.

IBM Operations Analytics - Log Analysis includes the following functions that you can configure in the Log Analysis UI:

Role-based access control

You can create and modify users and roles to assign role-based access control to individual users.

For more information, see Users and roles in the Log Analysis documentation.

Alerting

You can create alerts that are based on events, and you can define the actions for the system to take when an alert is triggered. For example, you can define an action to send an email notification to one or more people when an alert is triggered.

For more information, see Managing alerts in the Log Analysis documentation.

Some components that are preinstalled with Log Analysis

Although many components are preinstalled with IBM Operations Analytics - Log Analysis, the following components might be especially useful with your z/OS Insight Packs:

Expert advice custom search dashboard

The Expert Advice provides links to contextually relevant information to help you resolve problems quickly. Using this application, you can select any column or cells in the Grid view and can launch a search of the IBM Support Portal. The application searches for matches to unique terms that are contained in the column that you select.

You can start the Expert Advice application by clicking **Search Dashboards** in the navigation pane of the Search workspace.

For more information about this component, see Custom Search Dashboards in the Log Analysis documentation.

Tip: To use this Expert Advice, the Log Analysis server must have access to the Internet. With the client-side Expert Advice extension that is provided by IBM Z Operations Analytics, you can access expert advice even if the Log Analysis server does not have access to the Internet. For more information about the client-side Expert Advice, see "Log Analysis extensions for z/OS Problem Insights and client-side Expert Advice."

WebSphere Application Server Insight Pack

This Insight Pack is different from the WebSphere Application Server for z/OS Insight Pack. It is intended for use with WebSphere Application Server on distributed platforms. It does not support native logging formats for WebSphere Application Server for z/OS. However, if you configure your WebSphere Application Server for z/OS environment to use a distributed logging format, this Insight Pack provides you with the required annotation and indexing capabilities.

For more information about this component, see WebSphere Application Server Insight® Pack in the Log Analysis documentation.

Log Analysis extensions for z/OS Problem Insights and client-side Expert Advice

IBM Z Operations Analytics provides extensions to IBM Operations Analytics - Log Analysis for z/OS Problem Insights and client-side Expert Advice.

Problem Insights

The Problem Insights extension provides near real-time insight into problems in your IT environment, with suggested actions to help resolve the problems.

If the Problem Insights component of IBM Z Operations Analytics is installed, the Log Analysis UI includes a new tab that is titled **Problem Insights**. For each sysplex from which data is being forwarded to the Log Analysis server, the Problem Insights page includes insight about certain problems that are identified in the ingested data.

You can select the time range for which you want to see insights. For example, if you want to know about certain problems that were identified in the last hour, you can select the last hour as the time range. The default time range is the last 15 minutes.

You can also reload the Problem Insights page and associated data by clicking the **Refresh** button that is located in the area of the UI where you select the time range.

Each sysplex in the monitored environment is represented by a button that indicates the number of problems that are found in that sysplex. You can toggle the showing or hiding of data for a sysplex by clicking the button for the respective sysplex.

The Problem Insights and Suggested Actions table includes the following information about each problem that is discovered:

Severity
The indication of the severity of the problem.

Sysplex
The sysplex where the problem occurred.

System
The system where the problem occurred.

Interval score
Indicates unusual patterns of message IDs within an analysis interval, in comparison to the model of normal system behavior for the respective system. This score is extracted from the IBM zAware *interval anomaly score*. A score of 99.5 or higher indicates an anomaly.

You can click a score to navigate in context to IBM zAware for more information about the interval.

Subsystem
The subsystem or system resources manager where the problem occurred.

Time The last time that the problem occurred in the selected time range. For example, if you select the last 15 minutes as the time range, this column shows the last time that the problem occurred in the last 15 minutes.

Problem Summary
A summary of the problem that provides insight into the cause.

Count The total number of occurrences of the problem in the selected time range.

If you hover over this number, a pop-up window shows the following information about each occurrence:

- Time
- Severity
- Sysplex
- System
- Interval score. You can click a score to navigate in context to IBM zAware for more information about that interval.

Suggested Actions
Click a question mark (?) to go to more insight about the problem, including actions that you can take to resolve the problem. The Suggested Actions information also includes links to other sources of information that might help you resolve the problem, such as relevant topics in the IBM Knowledge Center.

Evidence
The message number of the message that identifies the problem, which you can click to view more information.

After the Problem Insights and Suggested Actions table, a bar chart for each system within the selected sysplex shows the interval scores for the last 24 hours. If you hover over an individual bar in a bar chart, a pop-up window shows the following information:

- The start and stop time for the interval
- The total number of unique message IDs for the interval
- Interval score

You can click an individual bar to navigate in context to IBM zAware for more information about the interval.

Client-side Expert Advice

If the IBM Operations Analytics - Log Analysis server is behind a firewall, and therefore does not have access to the Internet, you can use the Client-side Expert Advice (**IBMSupportPortal-ExpertAdvice on Client**) rather than the default Expert Advice (**IBMSupportPortal-ExpertAdvice**) in the Log Analysis UI.

After you install the IBM Z Operations Analytics extensions, the following applications are shown under **Expert Advice** in the Custom Search Dashboards panel of the left navigation pane of the Search workspace:

- **IBMSupportPortal-ExpertAdvice**, which is the default Expert Advice in Log Analysis. When this Expert Advice is launched, the Log Analysis server sends search requests to the IBM Support Portal and displays the query search results in the Log Analysis UI.
- **IBMSupportPortal-ExpertAdvice on Client**, which is the extension for client-side Expert Advice that is provided by IBM Z Operations Analytics. When this Expert Advice is launched, the client browser sends search requests directly to the IBM Support Portal and opens a new browser tab to display the query search results.

IBM zAware V3.1 feature

IBM Z Operations Analytics includes the separately installable IBM z Advanced Workload Analysis Reporter (IBM zAware) Version 3.1 feature.

For more information about this feature, see the IBM zAware Guide.

This feature must be installed before you can use the IBM zAware data gatherer.

IBM zAware data gatherer

IBM Z Operations Analytics includes the IBM zAware data gatherer, which is a client that gathers interval anomaly data from an IBM zAware server.

The data is ingested into IBM Operations Analytics - Log Analysis and is shown in the Log Analysis user interface under the data source that is named `zOS Anomaly Interval`. Visualizations of this interval anomaly data are also shown on the new **Problem Insights** page, if the Problem Insights extension is installed.

Before you can use the IBM zAware data gatherer, both the IBM zAware V3.1 feature and the data gatherer must be installed, and the data gatherer must be configured with an IBM zAware server definition. Python 2.7.9 or later must be installed so that the data gatherer can connect to an IBM zAware server.

The data gatherer is installed as part of the installation of the z/OS Insight Packs and extensions. You can also configure the data gatherer during this installation by responding to prompts for the values of the configuration parameters.

Logstash and the `ioaz` Logstash output plugin

Logstash is an open source data collection engine that in near real-time, can dynamically unify data from disparate sources and normalize the data into the destinations of your choice for analysis and visualization. IBM Z Operations Analytics provides a copy of Logstash 2.3.4 and an `ioaz` Logstash output plugin that forwards data to the IBM Operations Analytics - Log Analysis server for processing by the z/OS Insight Packs.

Because IBM Common Data Provider for z Systems cannot forward data directly to the Log Analysis server, the z/OS log data and SMF data that is gathered by IBM Common Data Provider for z Systems must be forwarded to a Logstash instance. The Logstash instance then forwards the data to the Log Analysis server. Through Logstash, you can also forward data to other destinations, or route data through a message broker such as Apache Kafka.

For more information about Logstash, see the Logstash documentation.

Installation and configuration checklists

These checklists summarize the important steps for installing and configuring the Z Operations Analytics components, including IBM Operations Analytics - Log Analysis, the z/OS Insight Packs, the Log Analysis extensions, the IBM zAware data gatherer, and Logstash and the `ioaz` Logstash output plugin.

The following steps outline the installation and configuration process:

__ 1. Install IBM Operations Analytics - Log Analysis V1.3.5. See “Checklist for IBM Operations Analytics - Log Analysis” on page 17.

__ 2. On the Log Analysis server, install the z/OS Insight Packs, the Log Analysis extensions, and the IBM zAware data gatherer. See Checklist for z/OS Insight Packs, Log Analysis extensions, and the IBM zAware data gatherer.

__ 3. Install Logstash 2.3.4 and the `ioaz` Logstash output plugin. See “Checklist for Logstash and the ioaz Logstash output plugin” on page 17.

__ 4. Determine the sources from which you want to gather log data. For more information, see “Planning for configuration of IBM Common Data Provider for z Systems” on page 3.

__ 5. Install the IBM Common Data Provider for z Systems in each z/OS logical partition (LPAR) for which you want to process z/OS log messages. Then, in each LPAR where the IBM Common Data Provider for z Systems is installed, configure it to gather and forward z/OS log data to Logstash. For more information, see the IBM Common Data Provider for z Systems V1.1.0 documentation.

If you plan to gather System Management Facilities (SMF) data, also complete the configuration steps that are described in “Enabling support for SMF data destined for IBM Z Operations Analytics” on page 3.

__ 6. Secure communication between IBM Common Data Provider for z Systems and Logstash. For more information, see “Securing communication between IBM Common Data Provider for z Systems and Logstash” on page 46.

__ 7. Ensure that z/OS data sources are configured and grouped appropriately in Log Analysis. For more information, see “Grouping data sources to optimize troubleshooting in your IT environment” on page 48.

Checklist for IBM Operations Analytics - Log Analysis

__ 1. Verify that the system requirements for Log Analysis, including the installation of the prerequisite software, are met. For more information, see Hardware and software requirements.

__ 2. Using the **`ulimit`** command, tune the Linux operating system so that the number of concurrent files is `4096`, and the virtual memory is `unlimited`. For more information, see Hardware and software requirements.

__ 3. To install and run Log Analysis, you must be logged in to the Linux computer system with a non-root user ID. Either create this non-root user ID, or use an existing non-root user ID. For more information, see "Planning for installation of Log Analysis" on page 19.

__ 4. Review the network port assignments for Log Analysis. If the default ports are not available, determine alternative port assignments. For more information, see Default ports in the Log Analysis documentation.

__ 5. The IBM Tivoli Monitoring Log File Agent is an optional component that is provided with Log Analysis and other IBM products. If you plan to gather logs from Linux for z Systems, install this agent as part of the Log Analysis installation. For more information, see Installing and configuring the IBM Tivoli Monitoring Log File Agent in the Log Analysis documentation.

__ 6. Verify that you are logged in to the Linux computer system with the non-root user ID, and complete the steps in "Installing Log Analysis" on page 23.

Checklist for z/OS Insight Packs, Log Analysis extensions, and the IBM zAware data gatherer

__ 1. Verify that Log Analysis is running.

__ 2. Verify that you are logged in to the Linux computer system with the non-root user ID that was used to install Log Analysis.

__ 3. Install the z/OS Insight Packs with sample searches, the extensions for Problem Insights and client-side Expert Advice, and the IBM zAware data gatherer. For more information, see "Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer" on page 25.

Tip: If the IBM zAware data gatherer is not configured, you cannot see visualizations of the interval anomaly data in the Log Analysis user interface.

For more information about configuring and getting started with the data gatherer, see the following topics:

- "Configuring the IBM zAware data gatherer" on page 37
- "Getting started with the IBM zAware data gatherer" on page 55

Checklist for Logstash and the `ioaz` Logstash output plugin

__ 1. Verify that the Logstash system requirements are met. For more information, see "Logstash requirements" on page 19.

__ 2. Install Logstash and the `ioaz` Logstash output plugin. For more information, see "Installing Logstash and the `ioaz` Logstash output plugin" on page 30.

__ 3. Configure Logstash and the `ioaz` Logstash output plugin. For more information, see "Configuring Logstash" on page 41.

Planning to use Z Operations Analytics

In planning for the installation of IBM Z Operations Analytics, you must also plan for the installation of IBM Operations Analytics - Log Analysis and the installation and configuration of both Logstash and the IBM Common Data Provider for z Systems.

System requirements

Ensure that your environment meets the system requirements for IBM Z Operations Analytics.

IBM Common Data Provider for z Systems requirements

For information about the IBM Common Data Provider for z Systems system requirements, see Planning to use IBM Common Data Provider for z Systems in the IBM Common Data Provider for z Systems V1.1.0 documentation.

Log Analysis requirements

For information about the system requirements for IBM Operations Analytics - Log Analysis, see Hardware and software requirements.

Restriction: If you deploy Log Analysis on a Linux on System z system, you can use the IBM InfoSphere® BigInsights® Hadoop Version 3.0 service, but at this time, the Hadoop Distributed File System (HDFS) must be installed on an x86 Linux system. The support for integrating Log Analysis on a Linux on System z system with other Hadoop distributions is dependent on the platform support that is provided by the respective Hadoop distribution.

z/OS Insight Pack requirements

You must install the z/OS Insight Packs on a Log Analysis server with Log Analysis Version 1.3.5. The IBM Z Operations Analytics Version 3.2.0 package includes IBM Operations Analytics - Log Analysis Version 1.3.5 Standard Edition.

Table 3 indicates the software versions that are supported by each z/OS Insight Pack.

Table 3. Software versions that are supported by each z/OS Insight Pack

Insight Pack	Supports data ingestion from this software
z/OS Network Insight Pack	• IBM z/OS 2.2 and 2.3 • IBM Tivoli NetView for z/OS 6.2 and 6.2.1
z/OS SMF Insight Pack	• IBM z/OS 2.2 and 2.3
z/OS SYSLOG Insight Pack	• IBM z/OS 2.2 and 2.3 • IBM CICS Transaction Server for z/OS 5.1.1, 5.2, 5.3, and 5.4 • IBM Db2 for z/OS 11.1 and 12.1 • IBM IMS for z/OS 13.1, 14.1, and 15.1 • IBM MQ for z/OS 8.0 and 9.0 • Access Monitor component of IBM Security zSecure™ Admin 2.2.1 with APAR OA52273, 2.3.0, or later

Table 3. Software versions that are supported by each z/OS Insight Pack (continued)

Insight Pack	Supports data ingestion from this software
WebSphere Application Server for z/OS Insight Pack	• IBM WebSphere Application Server for z/OS 8.5.5 and 9.0

Logstash requirements

You must install Logstash 2.3.4 on a Linux on System p, System x, or System z system with the following software:

- Red Hat Enterprise Linux 6 or 7 or SUSE Linux Enterprise Server 11 or 12
- Java Runtime Environment (JRE) 8

  Restriction: For this version of Logstash, only an IBM JRE is supported. The following file on the Log Analysis server includes a supported JRE:

  LA_INSTALL_DIR/unity_images/ibm-java-sdk-*.tgz

Logstash requires 460 MB of disk space, plus more disk space for logs and the cache. The cache is used for storing data that cannot be sent if the IBM Operations Analytics - Log Analysis server is down.

Planning for installation of Log Analysis

Before you install IBM Operations Analytics - Log Analysis, review the requirements for the installation user ID, domain name resolution, and network connectivity. Also, decide whether you need to use the optional IBM Tivoli Monitoring Log File Agent component.

Before you begin

Verify that your environment meets the system requirements that are described in "System requirements" on page 18.

About this task

Installation user ID

To install and run Log Analysis, you must be logged in to the Linux computer system with a non-root user ID.

Domain name resolution

On the system where you plan to install Log Analysis, verify that the details for the Log Analysis server are maintained correctly in the /etc/hosts file.

Network connectivity between the Logstash server and the IBM Common Data Provider for z Systems LPARs

The LPARs with IBM Common Data Provider for z Systems must communicate with the Logstash server on the port that is defined during Logstash pipeline configuration.

The default value for this port is 8080. Ensure that communication on this port is not blocked by a firewall or by other aspects of your network configuration.

Network connectivity between the Logstash server and the Log Analysis server

The Logstash server must communicate with the Log Analysis server on the port that is defined during Log Analysis installation.

The default value for this port is 9987. Ensure that communication on this port is not blocked by a firewall or by other aspects of your network configuration.

Optional Log File Agent component

The IBM Tivoli Monitoring Log File Agent is an optional component that is provided with Log Analysis and other IBM products. You can use this agent to collect logs from Linux for z Systems and from other Linux and UNIX operating systems. For more information, see Installing and configuring the IBM Tivoli Monitoring Log File Agent in the Log Analysis documentation.

Planning for Log Analysis to support LDAP interaction with RACF for sign-on authentication

You can configure IBM Operations Analytics - Log Analysis to support LDAP interaction with RACF for sign-on authentication.

Procedure

You must plan for the following configuration steps:

1. In Log Analysis, verify that no data sources are created and that no permissions are assigned to data sources.
2. In the Log Analysis UI, create at least one user ID for use in doing administrative tasks in Log Analysis.
3. Define other user IDs for use in doing non-administrative tasks in Log Analysis.
4. Because LDAP and RACF authenticate these user IDs during sign-on, ensure that these user IDs are also defined to LDAP and RACF.

   Tip: In this configuration with LDAP and RACF, Log Analysis cannot use the default user IDs `unityadmin` and `unityuser` due to RACF restrictions on the length of a user ID (it must be 7 characters or less).
5. Configure Log Analysis to use LDAP for authentication.

   For instructions, see LDAP configuration in the Log Analysis documentation.

   For z Systems LDAP support, choose the Tivoli Directory Server as the type of LDAP server.
6. Save a backup copy of the following Log Analysis configuration files so that you can update these files.
   - `LA_INSTALL_DIR/utilities/datacollector-client/javaDatacollector.properties`
   - `LA_INSTALL_DIR/remote_install_tool/config/rest-api.properties`
   - `LA_INSTALL_DIR/UnityEIFReceiver/config/unity.conf`
   - `LA_INSTALL_DIR/solr_install_tool/scripts/register_solr_instance.sh`
7. Update the passwords in the four previously listed Log Analysis configuration files.

   For more information, see Changing a user password in the Log Analysis documentation.
8. Map your LDAP groups to the Log Analysis security role so that the users can access Log Analysis.

   Individual users can also be granted user or administrative privileges.

For more information, see LDAP configuration in the Log Analysis documentation.

9. If you use the delete utility to prune data sources, provide a user ID and password that are known to RACF to this utility.

   You can encrypt the password if you choose to store it in the Log Analysis server.

Planning for configuration of Logstash

You must install and configure at least one Logstash instance to receive z/OS data from one or more IBM Common Data Provider for z Systems instances and to forward this data to the IBM Operations Analytics - Log Analysis server.

About this task

You can install Logstash on the same system as the Log Analysis server or on a separate server.

A self-extracting installer installs Logstash and the `ioaz` Logstash output plugin and updates the Logstash configuration file with the values that are provided during the installation. You can use that Logstash configuration, or you can set up a more advanced Logstash configuration.

Procedure

Plan to use either the basic Logstash configuration that results from running the self-extracting installer, or set up your configuration as described in one of the following two examples:

Configuration description	More information
Basic Logstash configuration	Use the single Logstash instance that is configured by the self-extracting installer. This configuration is a good way to start.
Example 1 of an advanced Logstash configuration	Use the following two different Logstash instances: **Receiver instance** A Logstash instance that receives data from IBM Common Data Provider for z Systems. **Sender instance** A Logstash instance that forwards data to the Log Analysis server. With this configuration, you can route the data through a message broker such as Apache Kafka.
Example 2 of an advanced Logstash configuration	For each z/OS sysplex or other grouping of LPARs, use a different Logstash instance. With this configuration, each Logstash instance processes a smaller volume of data.

Important: Regardless of the Logstash configuration that you use, the data that is read by IBM Common Data Provider for z Systems for each data stream must arrive at the Log Analysis server in the order in which it was read. Otherwise, partial records might be incorrectly combined in the Log Analysis server.

Data sources

In the IBM Common Data Provider for z Systems configuration for each z/OS logical partition (LPAR), you define a data stream for each separate source of z/OS log data or SMF data. For each data stream that you define, a corresponding data source is created in IBM Operations Analytics - Log Analysis.

A *data source* is metadata about log data that enables the log data to be ingested for analysis. The data source includes, for example, the data source name, the log type, the origin of the log, and an annotation function that enables the content to be more easily searched, filtered, and used to create visualizations in dashboards.

How the length of data source names is affected by the Auto-Qualify value in the subscriber configuration

Subscriber configuration in the IBM Common Data Provider for z Systems documentation describes the subscriber configuration values that must be defined in the policy, including the **Auto-Qualify** value. The value that you select in the **Auto-Qualify** field can affect the length of the data source name. Depending on the value that you select, IBM Common Data Provider for z Systems adds more characters to the data source name.

`Sysplex`
If you select `Sysplex`, IBM Common Data Provider for z Systems can add a maximum of 18 characters to the data source name that is derived from the data stream name. For example, it adds the sysplex name (which can be a maximum of 8 characters), the system name (which can be a maximum of 8 characters), and 2 hyphens.

`System`
If you select `System`, IBM Common Data Provider for z Systems can add a maximum of 9 characters to the data source name that is derived from the data stream name. For example, it adds the system name (which can be a maximum of 8 characters) and 1 hyphen.

Why you might need to edit a data source name: IBM Operations Analytics - Log Analysis does not support data source names that have more than 30 characters. Therefore, in the IBM Common Data Provider for z Systems data stream definition, you might need to edit the value of the data source name to shorten the default name that is given by IBM Common Data Provider for z Systems. For example, in the following scenario, you might need to shorten the data source name to `SMF_30_ACCT`:

- If the data source name that is given by IBM Common Data Provider for z Systems is `SMF_30_ACCOUNTING`
- If your **Auto-Qualify** value is `Sysplex`

Data source types

The configuration artifacts that are provided with each z/OS Insight Pack include data source types. A log file splitter and a log record annotator are provided for each type of data source.

Log file splitters
The log file splitters split data into records.

Log record annotators
The annotators annotate fields in the log records so that data in those fields can be more easily searched, filtered, and used to create visualizations in dashboards. Each field corresponds to part of a log record.

The fields are defined in the index configuration file, and each field is assigned index configuration attributes.

The annotated fields are displayed in the IBM Operations Analytics - Log Analysis Search workspace and can be used to filter or search the log records.

Installing Z Operations Analytics

As part of installing Z Operations Analytics, you install IBM Operations Analytics - Log Analysis, the z/OS Insight Packs, Logstash and the `ioaz` Logstash output plugin, and the IBM Common Data Provider for z Systems.

About this task

You must install the following software:

1. Log Analysis
2. z/OS Insight Packs

   The following software is installed as part of installing the z/OS Insight Packs:

   - Log Analysis extensions for Problem Insights and client-side Expert Advice
   - IBM zAware data gatherer
3. Logstash, including the `ioaz` Logstash output plugin
4. IBM Common Data Provider for z Systems

   For information about installing IBM Common Data Provider for z Systems, see the IBM Common Data Provider for z Systems V1.1.0 documentation.

Tip: To verify that your version of the Z Operations Analytics is the latest available version, check for updates on the IBM Software Support site.

Installing Log Analysis

IBM Z Operations Analytics can be used with an existing instance of IBM Operations Analytics - Log Analysis Version 1.3.5. The IBM Z Operations Analytics Version 3.2.0 package includes IBM Operations Analytics - Log Analysis Version 1.3.5 Standard Edition. You can also install Log Analysis Version 1.3.5 Fix Pack 1 or Fix Pack 2.

Before you begin

Tip: To get Log Analysis Version 1.3.5 Fix Pack 1 (1.3.5.1) or Fix Pack 2 (1.3.5.2) , complete the following steps:

1. Go to IBM Fix Central.
2. In the **Product selector** field, start typing `IBM Operations Analytics - Log Analysis`, and when the correct product name is shown in the resulting list, select it. More fields are then shown.
3. In the **Installed Version** field, select `1.3.5`.
4. In the **Platform** field, select `All`.
5. Click **Continue**.
6. In the resulting “Identify fixes” window, select **Browse for fixes**, and click **Continue**.
7. In the “Select fixes” window, you should see the Fix Packs `1.3.5-TIV-IOALA-FP001` and `1.3.5-TIV-IOALA-FP002`, which you can select and download. For installation instructions, see the readme file.

Complete the planning task that is described in “Planning for installation of Log Analysis” on page 19.

The Log Analysis installation media is included in the Z Operations Analytics offering on the DVDs with the following labels:

IBM Operations Analytics - Log Analysis 1.3.5 Linux 64 bit (LCD8-2724)

IBM Operations Analytics - Log Analysis 1.3.5 Linux on System z 64 bit (LCD7-6059)

IBM Operations Analytics - Log Analysis 1.3.5 Linux on Power® 8 64 bit (LCD8-2737)

For more information about installing Log Analysis, see Installing in the Log Analysis documentation.

Procedure

1. Insert the DVD media into the DVD drive of (or mount the corresponding ISO image on) the Linux system that you plan to use as the Log Analysis server. If the DVD or image is not mounted automatically, mount it by using one of the utilities that are provided with the Linux operating system.
2. In a terminal window, change to the DVD or image directory by issuing the following command, where *media_mountpoint* is the directory in which the media is mounted:

   ```
   cd media_mountpoint
   ```
3. Run the installation process in either graphical or console mode, and provide values as necessary.

Option	Description
Graphical mode	Run the `install.sh` script.
Console mode	Run the `install.sh` script with the `-c` option, as shown in the following example: `./install.sh –c`

4. Wait for installation to complete.

   When the installation is complete, the Log Analysis server starts automatically.
5. To verify the status of the Log Analysis server, issue the following command:

   ```
   LA_INSTALL_DIR/utilities/unity.sh -status
   ```

   The following example shows sample output of this command:

   ```
   Fri April 27 09:43:13 EDT 2018
   IBM Operations Analytics - Log Analysis v1.3.5 STANDARD EDITION Application Services Status:
   ---------------------------------------------------------
   No.  Service                        Status     Process ID
   ---------------------------------------------------------
   1    Derby Network Server           UP         26279
   2    ZooKeeper                      UP         26317
   3    Websphere Liberty Profile      UP         26440
   4    EIF Receiver                   UP         26579
   ---------------------------------------------------------
   Getting status of Solr on myhost.mydomain.com
   Status of Solr Nodes:
   -----------------------------------------------------------------
   No. Instance Name            Host                 Status  State
   -----------------------------------------------------------------
   1   SOLR_NODE_LOCAL          myhost.mydomain.com  UP      ACTIVE
   -----------------------------------------------------------------
   All Application Services are in Running State
   Checking server initialization status: Server has initialized!
   ```

Depending on the components that are installed, different services are displayed in the output for this command. Verify that all services display an `UP` status. Also, verify that the last message in the output indicates that the server has initialized.

Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer

For IBM Operations Analytics - Log Analysis to provide insight about z/OS operational data, you must install the IBM Z Operations Analytics components for Log Analysis, including the z/OS Insight Packs, the extensions for Problem Insights and client-side Expert Advice, and the IBM zAware data gatherer.

Before you begin

IBM Operations Analytics - Log Analysis Version 1.3.5 must be installed. For more information, see "Installing Log Analysis" on page 23.

For more information about the z/OS Insight Packs, the extensions, and the IBM zAware data gatherer, see the following topics:

- "z/OS Insight Packs" on page 11
- "Log Analysis extensions for z/OS Problem Insights and client-side Expert Advice" on page 13
- "IBM zAware data gatherer" on page 15

Each Insight Pack includes optional sample searches (sometimes called *Quick Search samples*) to help you find common errors in the software products that the Insight Pack supports.

The installation packages are on the product DVD with the following label:

IBM Z Operations Analytics (LCD7-6544)

When you install the z/OS Insight Packs and extensions, Log Analysis must be running, and you must be logged in to the Linux computer system with the non-root user ID that was used to install Log Analysis.

About this task

The self-extracting installer file `izoa_install.run` simplifies the installation of the Insight Packs with sample searches and the extensions for Problem Insights and client-side Expert Advice. All four z/OS Insight Packs and their sample searches are installed. The IBM zAware data gatherer is also installed.

To install the Insight Packs, the installer uses the **`pkg_mgmt.sh`** command with the `-upgrade` option. If previous versions of the z/OS Insight Packs are installed on the system, the installer upgrades those Insight Packs as part of the installation process. For information about the **`pkg_mgmt.sh`** command, see pkg_mgmt.sh command in the Log Analysis documentation.

To install Insight Packs, extensions, sample searches, and the IBM zAware data gatherer, the installer must have the directory where Log Analysis is installed (*LA_INSTALL_DIR*) and the user name and password for logging in to Log Analysis.

The installer prompts you to respond whether you want to configure the IBM zAware data gatherer. If you choose to configure this data gatherer at installation time, you are then prompted for other parameters that are needed to complete the configuration. If you choose not to configure at installation time, you can configure this data gatherer later from the command line by using the `zAwareDataGathererConfig.py` script. For more information about this configuration, see "Configuring the IBM zAware data gatherer" on page 37.

Important: During the installation of the Problem Insights and client-side Expert Advice extensions with the installer file, the Log Analysis server is stopped and restarted.

Procedure

1. Insert the Insight Pack DVD media into the DVD drive of the Log Analysis server. If the DVD is not mounted automatically, mount it by using one of the utilities that are provided with the Linux operating system.

   Restriction: If you obtained the z/OS Insight Packs and extensions as a `.tar` file from IBM Shopz or IBM Fix Central, unpack the `.tar` file into a temporary directory on the target computer, and complete the remaining steps in this procedure.
2. Run the command `sh izoa_install.run`, and specify the requested information.

   If you try to install with an incorrect user ID: Remember that to install the Insight Packs and extensions, you must be logged in to the Linux computer system with the non-root user ID that was used to install Log Analysis.

   If you try to install these with some other user ID, and then rerun the installer with the correct user ID, the installer might stop with the following message:

   ```
   Creating directory /tmp/install_izoa
   Verifying archive integrity... All good.
   Uncompressing Install Extensions and Insight Packs  100%  Extraction failed.
   Terminated
   ```

   To resolve this error, remove the directory and log file by running the following commands:

   ```
   rm -rf /tmp/install_izoa
   rm /tmp/izoa_install.log
   ```

What to do next

If you installed any of the optional predefined searches (sometimes also called *sample searches*, *saved searches*, or *Quick Search samples*) in IBM Operations Analytics - Log Analysis, they are available in the `zos` folder in the **Saved Searches** navigator of the Log Analysis user interface.

Some sample searches are organized into subfolders under the `zos` folder.

To use a sample search to search logs, double-click the title of the search in the **Saved Searches** navigator.

Installing Insight Packs for SMF data source types that are provided by IBM Common Data Provider for z Systems

If you want to process operational data from System Management Facilities (SMF) data source types that are provided by IBM Common Data Provider for z Systems, you must install the separate Insight Packs that are available for these data source types.

Before you begin

The z/OS Insight Packs that are provided with IBM Z Operations Analytics support the data source types that are described in "Operational insights reference" on page 105. For information about installing those Insight Packs, see "Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer" on page 25.

About this task

To install the separate Insight Packs for the SMF data source types that are provided by IBM Common Data Provider for z Systems, you use the `installCDPSourceTypes.sh` script, which is described in "`installCDPSourceTypes.sh` script for installing Insight Packs" on page 28.

Procedure

To install the Insight Packs, complete the following steps:

1. Determine the additional SMF data source types from which you want to process operational data.

   Restriction: Because each unique data source type uses system resources for Logstash and IBM Operations Analytics - Log Analysis, limit the number of installed Insight Packs to the minimum that you need. For example, if you have more than 30 Insight Packs installed, system resource problems might occur.
2. Install Logstash, as described in "Installing Logstash and the `ioaz` Logstash output plugin" on page 30.
3. Copy the `izoaCDPSourceTypes.zip` file from the IBM Z Operations Analytics installation package to the system where Logstash and Log Analysis are installed.

   Important: After you extract the `izoaCDPSourceTypes.zip` file, do not move the `installCDPSourceTypes.sh` file separately to another location because the `installCDPSourceTypes.sh` script runs only from the `bin` directory within the directory structure of the `.zip` file.
4. Optional: Set the following environment variables:

   LA_INSTALL_DIR
   Directory where Log Analysis is installed

   LOGSTASH_HOME
   Directory where Logstash is installed
5. To run the `installCDPSourceTypes.sh` script, issue one of the following commands, depending on whether you also want to collect log output for the script:
   - To run the script without collecting log output:
     ```
     ./installCDPSourceTypes.sh [DATA TYPES] [SYSTEM TYPE]
     ```
   - To run the script, and collect log output:

```
./installCDPSourceTypes.sh [DATA TYPES] [SYSTEM TYPE] 2>&1 | tee -a ipinstall.log
```

6. In the IBM Common Data Provider for z Systems Configuration Tool policy, add the appropriate SMF data streams so that you can receive the data.

   Subscriber configuration in the IBM Common Data Provider for z Systems documentation describes the subscriber configuration values that must be defined in the policy. Use the following values for the following fields:

   Protocol
   `CDP Logstash`

   Host The host name or IP address of the system where Logstash is installed.

   Port The port that is indicated in the Logstash configuration file `config/B_logstash-ioaz-input.conf`.

   Auto-Qualify
   If data is being sent from multiple systems, use the value `System`.

   "Data sources" on page 22 includes information about how the value that you select in the **Auto-Qualify** field can affect the length of the data source name, and why you might need to edit the data source name due to restrictions on the length of the name.

   Send As
   `Unsplit`

`installCDPSourceTypes.sh` script for installing Insight Packs

The `installCDPSourceTypes.sh` script, which is in the IBM Z Operations Analytics installation package in the `izoaCDPSourceTypes.zip` file, simplifies the installation of Insight Packs for System Management Facilities (SMF) data source types that are provided by IBM Common Data Provider for z Systems.

To install the Insight Packs, the `installCDPSourceTypes.sh` script uses the IBM Operations Analytics - Log Analysis DSV toolkit. The script automatically overwrites any existing Insight Pack with the same name, if it detects that the existing Insight Pack was not previously run for the data source type. However, if the script finds evidence in the *`LA_INSTALL_DIR`*`/izoa_sourcetypes` directory (which is described in "Predefined values that can be changed in the script" on page 29) that the respective data source type was previously processed, it does not re-create or reinstall the Insight Pack for that data source type.

Also, for each data source type, the `installCDPSourceTypes.sh` script moves a configuration file to the *`LOGSTASH_INSTALL_DIR`*`/config` directory.

When you run the `installCDPSourceTypes.sh` script, you must provide input for some parameters, as described in "Required script parameters" and "Directories that must be provided to the script" on page 29. If you provide input that is blank, or if you provide an invalid data source type, the script ends, and you must rerun it.

The script also contains some predefined values that can be changed, as described in "Predefined values that can be changed in the script" on page 29.

Required script parameters

Data source types for which you want insights
Enter the value for each data source type for which you want insights. For

example, if you want insights for all System Management Facilities (SMF) record type 110 data (CICS Transaction Server for z/OS), the data source type is `SMF_110*`.

The value for each data source type must be the same as the corresponding file name in the `properties` directory in the extracted `izoaCDPSourceTypes.zip` file.

You can use the asterisk (*) as a wildcard character.

Restriction: Because each unique data source type uses system resources for Logstash and IBM Operations Analytics - Log Analysis, limit the number of installed Insight Packs to the minimum that you need. For example, if you have more than 30 Insight Packs installed, system resource problems might occur.

System type
The following values are valid:

A If both Logstash and Log Analysis are installed on the system where you are running the script, use this value.

B If only Logstash is installed on the system where you are running the script, use this value.

C If only Log Analysis is installed on the system where you are running the script, use this value.

Important: If you do not use value `A`, the script must be run twice, once with value `B`, and once with value `C`, in that order. Also, before the script is run, the `izoaCDPSourceTypes.zip` file must be extracted on both the Logstash system and the Log Analysis system.

Directories that must be provided to the script

If you do not set environment variables for the following directories, as described in step 4 on page 27 of “Installing Insight Packs for SMF data source types that are provided by IBM Common Data Provider for z Systems” on page 27, the script prompts you for these values:

- Directory where Logstash is installed
- Directory where Log Analysis is installed

Predefined values that can be changed in the script

The script contains the following predefined values that can be changed:

`IZOA_PROPERTIES_DIR="izoa_sourcetypes"`
This value is the name of the Log Analysis directory where the script copies each `properties` file when it creates or installs the Insight Packs. If a `properties` file was previously moved into this directory, then, to prevent the overwriting of any updates that were made to the file, the script does not re-create or reinstall the Insight Pack for the corresponding data source type.

`AUTOMATICALLYDEPLOY="y"`
Automatic deployment occurs only if this value is `y`.

Otherwise, you must deploy the Insight Packs by using the `properties` files that are in the directory that is specified by `IZOA_PROPERTIES_DIR`. For information about how to deploy your own Insight Packs, see Generate an Insight Pack in the Log Analysis documentation.

`MAXIMUMTODEPLOY=5`
: This value is the maximum number of Insight Packs to create or deploy without prompting the user to verify the number of data source types that are indicated.

  Restriction: Because each unique data source type uses system resources for Logstash and IBM Operations Analytics - Log Analysis, limit the number of installed Insight Packs to the minimum that you need. For example, if you have more than 30 Insight Packs installed, system resource problems might occur.

`DSVTOOLKIT="DSVToolkit_v1.1.0.4"`
: This value is the name of the DSV toolkit directory, which is in the *`LOGSTASH_INSTALL_DIR`*`/unity_content` directory.

Installing Logstash and the `ioaz` Logstash output plugin

To send data from the IBM Common Data Provider for z Systems to the IBM Operations Analytics - Log Analysis server, you must install the `ioaz` Logstash output plugin with Logstash 2.3.4. The Logstash 2.3.4 version that is provided with IBM Z Operations Analytics is optimized for use with Linux on z Systems.

Before you begin

Review the following information, and verify that the system where you plan to install Logstash meets the Logstash system requirements:

- "Logstash and the `ioaz` Logstash output plugin" on page 16
- Logstash system requirements
- "Planning for configuration of Logstash" on page 21

About this task

To install Logstash and the `ioaz` Logstash output plugin, run the self-extracting installer file `logstash_install.run` on the target system. For the installation of Logstash, you do not need administrator privileges, and you can use a non-root user ID.

During the installation, you are prompted for the following information:

Logstash installation directory
: The directory where you want to install Logstash. You must have system permissions to create this directory, and adequate free space must be available on the file system for the installation.

Log Analysis server name
: The fully qualified domain name of the IBM Operations Analytics - Log Analysis server that the `ioaz` Logstash output plugin connects to.

  If you use the default value of `localhost`, the Log Analysis server must be running on the same system where you install Logstash.

Log Analysis server port number
: The port number that is used by the `ioaz` Logstash output plugin to communicate with the Log Analysis server. The default value is `9987`.

Log Analysis user name

The user name that the ioaz Logstash output plugin uses to log in to the Log Analysis server. The default value is unityadmin.

Password for Log Analysis user name

The password for the Log Analysis user name that the ioaz Logstash output plugin uses to log in to the Log Analysis server. The default value is unityadmin.

Plugin cache directory

The directory that the ioaz Logstash output plugin uses for caching events when it cannot communicate with the Log Analysis server. The default value is *LOGSTASH_INSTALL_DIR*/cache_dir. Verify that adequate space for the cache is available in the specified cache directory.

Plugin log directory

The directory to which the ioaz Logstash output plugin writes logs. The default value is *LOGSTASH_INSTALL_DIR*/logs.

Port on which Logstash listens for z/OS data

The port on which the ioaz Logstash output plugin listens for data from the IBM Common Data Provider for z Systems. The default value is 8080.

Trust all certificates

A true or false indication of whether the ioaz Logstash output plugin should trust all security certificates when it sends data to the Log Analysis server. The default value is true.

A value of false instructs the ioaz Logstash output plugin to compare the security certificate from the Log Analysis server with the certificates that are stored in the cacerts keystore file. Therefore, if you specify a value of false, you must manually import a security certificate for the Log Analysis server into the cacerts keystore file for the Java Runtime Environment (JRE) that Logstash uses. For more information, see “Verifying the identity of the Log Analysis server” on page 41.

Procedure

To install Logstash and the ioaz Logstash output plugin, complete the following steps:

1. For specification during the installation, gather the information that is described in About this task.
2. Transfer the Logstash logstash_install.run file to the system where you plan to install Logstash and the ioaz Logstash output plugin.
3. Complete one of the following actions to give the self-extracting installer file logstash_install.run access to the Java Runtime Environment (JRE) that is used for installing and running Logstash:
   - Set the *JAVA_HOME* environment variable to the JRE location.
   - Add the fully qualified path of the JRE executable to the *PATH* environment variable.

   For example, if Log Analysis is installed in the /opt/IBM/LogAnalysis directory on the system where you plan to install Logstash and the plugin, and you want to use the JRE 8 that is packaged with Log Analysis, issue one of the following commands:

   To set *JAVA_HOME*

   ```
   export JAVA_HOME=/opt/IBM/LogAnalysis/ibm-java
   ```

To add the JRE path to *PATH*

```
export PATH=/opt/IBM/LogAnalysis/ibm-java/bin:$PATH
```

4. Issue the following command to run the installer:

   ```
   sh logstash_install.run
   ```

   After Logstash and the `ioaz` Logstash output plugin are installed, the Logstash configuration files are updated with the values that were provided during the installation, and Logstash is started.

   The Logstash configuration files are in the *LOGSTASH_INSTALL_DIR*/`config` directory. Table 4 indicates the prefixes that are used in the file names for the Logstash configuration files. The file name prefix is an indication of the configuration file content.

Table 4. Mapping of the prefix that is used in a Logstash configuration file name to the content of the file

Prefix in file name of Logstash configuration file	Content of configuration file with this prefix
B_	Input stage
H_	Field name annotation stage
N_	Timestamp resolution stage
Q_	Output stage

The following descriptions further explain these Logstash configuration files:

`B_logstash-ioaz-input.conf` file
This file contains the input stage that specifies the TCP/IP port on which Logstash listens for data from the IBM Common Data Provider for z Systems Data Streamer. This file is always present and is required for all data source types. The file is configured for you when you install Logstash, but you might need to update the port number in the file at a later time.

`H_logstash_ioaz.conf` file
This file contains rules that create a unique index name for each SMF data stream that is provided by IBM Common Data Provider for z Systems, based on the data source type and the name of the system from which the operational data is collected. This file is present in the *LOGSTASH_INSTALL_DIR*/`config` directory only if you install support for the associated SMF data source types, as described in “Installing Insight Packs for SMF data source types that are provided by IBM Common Data Provider for z Systems” on page 27.

Files with `N_` prefix in file name
Each of these files contains a unique timestamp resolution stage that maps to a unique SMF data stream that is provided by IBM Common Data Provider for z Systems. These files are present in the *LOGSTASH_INSTALL_DIR*/`config` directory only if you install support for the associated SMF data source types, as described in “Installing Insight Packs for SMF data source types that are provided by IBM Common Data Provider for z Systems” on page 27.

`Q_logstash-ioaz-output.conf` file
This file contains an output stage that sends all records to a single Log Analysis server. This file is always present and is required for all data source types. The file is configured for you when you install Logstash, but you might need to update the Log Analysis server information in the file at a later time.

What to do next

You can update the Logstash configuration files in the *LOGSTASH_INSTALL_DIR*/config directory by using a text editor. If you update these files, you must restart Logstash.

You can start and stop the ioaz Logstash output plugin as needed by using the *LOGSTASH_INSTALL_DIR*/bin/logstash_util.sh script.

Related reference:

"Configuration options for the ioaz Logstash output plugin" on page 42
The ioaz Logstash output plugin, which is provided by IBM Z Operations Analytics, forwards event data from the IBM Common Data Provider for z Systems to the IBM Operations Analytics - Log Analysis server. This reference lists the configuration options that can be set for the plugin in the Logstash configuration file, which is *LOGSTASH_INSTALL_DIR*/config/logstash-ioaz.conf.

Uninstalling the Insight Packs

When you uninstall an Insight Pack, all data sources that are associated with the Insight Pack are deleted. If ingested data exists for any data source types that are defined by the Insight Pack, you must remove that data before you can uninstall the Insight Pack.

Before you begin

If you are uninstalling the WebSphere Application Server for z/OS Insight Pack: The WASSystemOut data source type is defined by the WebSphere Application Server Insight Pack that is provided with Log Analysis. Data sources that are associated with this data source type are not deleted when you uninstall the WebSphere Application Server for z/OS Insight Pack. You do not have to remove data for the WASSystemOut data source type.

Procedure

To uninstall an Insight Pack, complete the following steps:

1. If ingested data exists for any data source types that are defined by the Insight Pack, remove that data from IBM Operations Analytics - Log Analysis by following the instructions in "Removing data from Log Analysis" on page 34.
2. Uninstall the Insight Pack by using the **pkg_mgmt.sh** command that is provided with Log Analysis. For example, issue one of the following commands, depending on the Insight Pack that you are uninstalling:

```
LA_INSTALL_DIR/utilities/pkg_mgmt.sh -uninstall
    LA_INSTALL_DIR/unity_content/SMFforzOS/
    SMFforzOSInsightPack_v3.2.0.0.zip
```

```
LA_INSTALL_DIR/utilities/pkg_mgmt.sh -uninstall
    LA_INSTALL_DIR/unity_content/SYSLOGforzOS/
    SYSLOGforzOSInsightPack_v3.2.0.0.zip
```

```
LA_INSTALL_DIR/utilities/pkg_mgmt.sh -uninstall
    LA_INSTALL_DIR/unity_content/WASforzOS/
    WASforzOSInsightPack_v3.2.0.0.zip
```

```
LA_INSTALL_DIR/utilities/pkg_mgmt.sh -uninstall
    LA_INSTALL_DIR/unity_content/zOSNetwork/
    zOSNetworkInsightPack_v3.2.0.0.zip
```

Removing data from Log Analysis

To remove data from IBM Operations Analytics - Log Analysis, use the deletion tool.

Before you begin

Before you remove the data, complete the following steps:

1. Stop each instance of the IBM Common Data Provider for z Systems.
2. Ensure that Log Analysis is running.

About this task

The deletion tool provides the following options (or use cases) for deleting data:

1. Delete all data from a single data source.
2. Delete all data from a single collection.
3. For a specified time period, delete data from all data sources.
4. At regular intervals, delete data that is older than the specified retention period.

For more information about removing data from Log Analysis, see delete.properties.

Procedure

To remove data from Log Analysis by using the deletion tool, complete the following steps:

1. In the `LA_INSTALL_DIR/utilities/deleteUtility` directory, open the `delete.properties` file.
2. For the use case that you want to run, specify the use case number and the variables that are associated with that use case.

   The use cases are summarized in About this task.
3. Save the `delete.properties` file.
4. Run the following command, where *Python_path* represents the location where Python is installed, and *password* represents the password that is defined in the `delete.properties` file for the associated user name:

   ```
   Python_path deleteUtility.py password
   ```

Uninstalling Log Analysis

You can uninstall IBM Operations Analytics - Log Analysis by using the IBM Installation Manager.

Before you begin

Before you uninstall Log Analysis, uninstall any remote installations of Apache Solr.

For more information about uninstalling Log Analysis, see Removing Log Analysis.

Procedure

Run the uninstallation process in either graphical or console mode.

Option	Description
Graphical mode	1. Go to the following directory, where *im_install_dir* represents the directory where the IBM Installation Manager is installed: *im_install_dir*/IBM/InstallationManager/eclipse. 2. Run the following command: `./launcher`
Console mode	1. Go to the following directory, where *im_install_dir* represents the directory where the IBM Installation Manager is installed: *im_install_dir*/IBM/InstallationManager/eclipse/tools. 2. Run the following command: `./imcl -c`

Upgrading Z Operations Analytics

You can upgrade from IBM Z Operations Analytics Version 3.1.0 to Version 3.2.0. IBM Z Operations Analytics Version 3.2.0 includes IBM Operations Analytics - Log Analysis Version 1.3.5 Standard Edition and IBM Common Data Provider for z Systems Version 1.1.0.

Before you begin

If your version of IBM Z Operations Analytics is earlier than V3.1.0, you must first upgrade to IBM Z Operations Analytics V3.1.0. For instructions, see Upgrading Operations Analytics for z Systems in the IBM Z Operations Analytics 3.1.0 documentation.

Upgrading from Log Analysis V1.3.3.1 to V1.3.5

Log data, System Management Facilities (SMF) data, and data sources that are associated with the z/OS Insight Packs are retained during the upgrade from IBM Operations Analytics - Log Analysis Version 1.3.3.1 to Version 1.3.5.

Procedure

To upgrade from Log Analysis V1.3.3.1 to V1.3.5, complete the following steps:

1. In each logical partition (LPAR) from which you are gathering data, stop the IBM Common Data Provider for z Systems components.
2. Stop the Log Analysis server.
3. Back up the data in Log Analysis.

   For information about backing up the data, see Upgrading, backing up, and migrating data in the V1.3.5 Log Analysis documentation.

4. Uninstall Log Analysis V1.3.3.

   See “Uninstalling Log Analysis” on page 34.
5. Install Log Analysis V1.3.5.

   See “Installing Log Analysis” on page 23.
6. Restore the backed-up data to Log Analysis.

   For information about restoring the data, see Restoring data in the V1.3.5 Log Analysis documentation.
7. Optional: Install any Log Analysis V1.3.5 fix packs.
8. In each LPAR from which you are gathering operational data, start the IBM Common Data Provider for z Systems components.
9. Verify that data is being properly sent to, and ingested by, the Log Analysis server.

Upgrading the z/OS Insight Packs and extensions

To get the latest insights about z/OS operational data, you must upgrade several IBM Z Operations Analytics components, including the z/OS Insight Packs and the extensions for Problem Insights and client-side Expert Advice.

Before you begin

Your IBM Operations Analytics - Log Analysis version must be Version 1.3.5. If this is not your version, upgrade your version according to the instructions in “Upgrading from Log Analysis V1.3.3.1 to V1.3.5” on page 35.

About this task

The self-extracting installer file `izoa_install.run` upgrades the four Insight Packs with sample searches and the extensions for Problem Insights and client-side Expert Advice. If the extensions were previously installed, they are upgraded. If they were not previously installed, they are installed as part of the upgrade process.

If the IBM zAware data gatherer was previously installed, it is also upgraded. If it was not previously installed, you can optionally install it as part of the upgrade process.

Procedure

To upgrade the z/OS Insight Packs and the extensions for Problem insights and client-side Expert Advice, complete the following steps:

1. In each logical partition (LPAR) from which you are gathering data, stop the IBM Common Data Provider for z Systems components.
2. Because Custom Search Dashboard applications are reinstalled as part of this upgrade, back up any IBM-provided Custom Search Dashboard applications that you customized.

   These Custom Search Dashboard applications are in the following directories:

   WebSphere Application Server for z/OS Insight Pack directory for dashboard applications

   LA_INSTALL_DIR/AppFramework/Apps/WASforzOSInsightPack_v3.1.0.0

   z/OS Network Insight Pack directory for dashboard applications

   LA_INSTALL_DIR/AppFramework/Apps/zOSNetworkInsightPack_v3.1.0.0

z/OS SMF Insight Pack directory for dashboard applications
: *LA_INSTALL_DIR*/AppFramework/Apps/SMFforzOSInsightPack_v3.1.0.0

z/OS SYSLOG Insight Pack directory for dashboard applications
: *LA_INSTALL_DIR*/AppFramework/Apps/SYSLOGforzOSInsightPack_v3.1.0.0

3. To upgrade the z/OS Insight Packs and extensions, complete the steps for installing the z/OS Insight Packs and extensions, which are described in "Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer" on page 25.
4. Update the IBM-provided Custom Search Dashboard applications in the following directories to include any custom changes that you made to the previous version of these Apps.

   WebSphere Application Server for z/OS Insight Pack directory for dashboard applications
   LA_INSTALL_DIR/AppFramework/Apps/WASforzOSInsightPack_v3.2.0.0

   z/OS Network Insight Pack directory for dashboard applications
   LA_INSTALL_DIR/AppFramework/Apps/zOSNetworkInsightPack_v3.2.0.0

   z/OS SMF Insight Pack directory for dashboard applications
   LA_INSTALL_DIR/AppFramework/Apps/SMFforzOSInsightPack_v3.2.0.0

   z/OS SYSLOG Insight Pack directory for dashboard applications
   LA_INSTALL_DIR/AppFramework/Apps/SYSLOGforzOSInsightPack_v3.2.0.0
5. In each LPAR from which you are gathering operational data, start the IBM Common Data Provider for z Systems components.
6. Verify that data is being properly sent to, and ingested by, the Log Analysis server.

What to do next

- If the Problem Insights page does not get updated, or does not render correctly, clear the browser cache.
- To get the latest Logstash updates, reinstall Logstash and the `ioaz` Logstash output plugin according to the instructions in "Installing Logstash and the `ioaz` Logstash output plugin" on page 30.

Configuring Z Operations Analytics

The primary configuration tasks for Z Operations Analytics are to configure the IBM zAware data gatherer and to secure communication between IBM Common Data Provider for z Systems and Logstash, and between Logstash and the IBM Operations Analytics - Log Analysis server. Optionally, you can configure some aspects of the Problem Insights extension for your environment.

Configuring the IBM zAware data gatherer

During the installation of the z/OS Insight Packs, extensions, and IBM zAware data gatherer, you are prompted about whether you want to configure the IBM zAware data gatherer. If you choose not to configure this data gatherer at installation time, you can configure it later from the command line by using the `zAwareDataGathererConfig.py` script.

Before you begin

For more information about the installation process, see "Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer" on page 25.

About this task

For more information about the configuration process, see "Configuration reference for the IBM zAware data gatherer." The following sections list and describe the information that you must provide during the configuration, regardless of whether you configure at installation time or later from the command line:

- "IBM zAware server information in the zAwareConfig.json configuration file" on page 39
- "Log Analysis server information in the unityClientConfig.json configuration file" on page 40

Procedure

To configure the data gatherer from the command line, complete the following steps:

1. Verify that you are logged in to the Linux computer system with the non-root user ID that was used to install Log Analysis.
2. Run the `zAwareDataGathererConfig.py` script.

   When you run the `zAwareDataGathererConfig.py` script, it first prompts you for whether you are adding or deleting an IBM zAware server definition.

   Tips:

   - The data gatherer gathers interval anomaly data from only one IBM zAware server. In this configuration, you must specify the host name or IP address of the IBM zAware server from which you want to gather interval anomaly data. That server must be actively collecting z/OS log data and deriving interval anomaly scores.
   - If you want to change the server definition after it is configured, you must first delete the existing definition, and then, add a new definition.
   - You can configure this data gatherer from the command line without using the options. Then, you are prompted for the parameters that are needed to complete the configuration. Otherwise, the following examples illustrate the use of the options:
     - The following example illustrates how to add a new IBM zAware server definition to the configuration:

       ```
       ./zAwareDataGathererConfig.py add -z server01.anycompany.com
           -u admin -w password -l unityadmin -s unityadmin -t 9987
       ```
     - The following example illustrates how to delete an existing IBM zAware server definition from the configuration:

       ```
       ./zAwareDataGathererConfig.py delete -z server01.mycompany.com
       ```
     - To access the Help for adding or deleting a server definition, use the `-h` option, as shown in the following example for getting help for the add operation:

       ```
       ./zAwareDataGathererConfig.py add -h
       ```

What to do next

See "Getting started with the IBM zAware data gatherer" on page 55.

Configuration reference for the IBM zAware data gatherer

The configuration of the IBM zAware data gatherer results in the creation of two JavaScript Object Notation (JSON) files, `zAwareConfig.json` and unityClientConfig.json, on the Log Analysis server in the *LA_INSTALL_DIR*/zAwareDataGatherer/config directory. These files contain the configuration information that was provided for the data gatherer either at installation time or later by using the zAwareDataGathererConfig.py script.

This reference includes the following information:
- "Verification of the configuration"
- "Situations in which you must reconfigure the data gatherer"
- "IBM zAware server information in the zAwareConfig.json configuration file"
- "Log Analysis server information in the unityClientConfig.json configuration file" on page 40

Verification of the configuration

If the two configuration files are present in the following directory, the IBM zAware data gatherer is configured, and you can open the files to view or verify their content:

```
LA_INSTALL_DIR/zAwareDataGatherer/config
```

Situations in which you must reconfigure the data gatherer

If an external change occurs to the values for any of the entries in the configuration files, you must reconfigure the IBM zAware data gatherer. For example, assume that a system administrator changes the password for the administrator user ID for the IBM zAware server server01.anycompany.com, which is defined in the data gatherer configuration. For the IBM zAware data gatherer to have continued access to server01.anycompany.com, the system administrator must reconfigure the data gatherer according to the following steps:

1. Stop the running process for the data gatherer.
2. Use the zAwareDataGathererConfig.py script to delete the existing definition for the server.
3. Use the zAwareDataGathererConfig.py script to add a new definition for the server, and specify the administrator user ID with the new password.
4. Restart the data gatherer.

IBM zAware server information in the zAwareConfig.json configuration file

The zAwareConfig.json file contains the information that is required by the IBM zAware data gatherer to access the IBM zAware server that provides interval anomaly data.

The following information about the IBM zAware server is defined in this file:

Host name or IP address

The host name or IP address of the IBM zAware server from which you want to gather interval anomaly data. The server must be actively collecting z/OS log data and deriving interval anomaly scores.

If the host name or IP address that you provide is not valid, the IBM zAware data gatherer does not start.

Required or optional?

Required for adding or deleting a IBM zAware server definition

Default value
None

Command line option for specifying this information
`-z`

User ID
The user ID that is defined to the IBM zAware server in the Administrator role and is authorized to retrieve interval anomaly data.

Required or optional?
Required for adding a new IBM zAware server definition

Default value
`admin`

Command line option for specifying this information
`-u`

Password
The password for the user ID. This password is encrypted in the configuration file.

Required or optional?
Required for adding a new IBM zAware server definition

Default value
None

Command line option for specifying this information
`-w`

Log Analysis server information in the `unityClientConfig.json` configuration file

The `unityClientConfig.json` file contains the information that is required by the IBM zAware data gatherer to access the Log Analysis server to provide the interval anomaly data for display in the Log Analysis UI.

The following information about the Log Analysis server is defined in this file:

User ID
The user ID for the Log Analysis server that is authorized to receive interval anomaly data.

Required or optional?
Required for adding a new IBM zAware server definition

Default value
`unityadmin`

Command line option for specifying this information
`-l`

Password
The password for the user ID. This password is encrypted in the configuration file.

Required or optional?
Required for adding a new IBM zAware server definition

Default value
None

Command line option for specifying this information

-s

Port number

The number of the Log Analysis server port that is listening for inbound communication.

Required or optional?

Optional

Default value

9987

Command line option for specifying this information

-t

Configuring Logstash

After Logstash and the `ioaz` Logstash output plugin are installed, the Logstash configuration file is updated with the values that were provided during the installation, and Logstash is started. The remaining basic configuration task is to secure communication between Logstash and the IBM Operations Analytics - Log Analysis server.

Before you begin

Review "Planning for configuration of Logstash" on page 21.

For more information about the configuration options for the `ioaz` Logstash output plugin, see "Configuration options for the `ioaz` Logstash output plugin" on page 42.

Verifying the identity of the Log Analysis server

For communication between Logstash and the IBM Operations Analytics - Log Analysis server to be secure, the identity of the Log Analysis server must be verified by the `ioaz` Logstash output plugin. Communication between Logstash and the Log Analysis server occurs over the Hypertext Transfer Protocol Secure (HTTPS). By default, the identity of the Log Analysis server is trusted without validation.

About this task

In the Logstash configuration file, specifying a value of `false` for the `trust_all_certificates` option directs the `ioaz` Logstash output plugin to compare the security certificate from the Log Analysis server with the certificates that are stored in the `cacerts` keystore file. Therefore, you must manually import a security certificate for the Log Analysis server into the `cacerts` keystore file for the Java Runtime Environment (JRE) that Logstash uses.

Location of `cacerts` keystore file

The `cacerts` keystore file is in the following path, where *`java.home`* represents the directory for the JRE that Logstash uses:

```
java.home/lib/security
```

You can configure and manage this file by using the keytool utility. The initial password of the `cacerts` file is `changeit`.

Default location of Log Analysis server certificate file

The default location of the Log Analysis server certificate file is the following path:

```
LA_INSTALL_DIR/wlp/usr/servers/Unity/resources/security/client.crt
```

Procedure

To secure communication between Logstash and the Log Analysis server, complete the following steps:

1. Copy the Log Analysis server certificate file from the Log Analysis server to the Logstash server.

   Attention: To prevent file corruption, the file transfer must occur in binary mode.
2. To import the Log Analysis server certificate, issue the following command (all on one line):

   ```
   JRE_path/bin/keytool -import -alias IOAz -file
       certificate_path/client.crt -keystore
       java.home/lib/security/cacerts -storepass changeit
   ```
3. In the Logstash configuration file, change the value of the `trust_all_certificates` option from `true` to `false`.
4. Restart Logstash by using the *LOGSTASH_INSTALL_DIR*/bin/logstash_util.sh script.

Updating the Logstash configuration with the user password for the Log Analysis server

For communication between Logstash and the IBM Operations Analytics - Log Analysis server to be secure, you must change the user password for the Log Analysis server on a regular schedule, encrypt the password, and update the Logstash configuration with the encrypted password.

Procedure

To change the user password in the Logstash configuration, complete one of the following steps:

- In the directory where Logstash is installed, run the `update_pass.sh` script.
- Insert the password into the script either by using the environment variable *LA_PASSWORD* or by running the following command with the `-p` option:

  ```
  update_pass.sh -p:password
  ```

The *LOGSTASH_INSTALL_DIR* and *JAVA_HOME* variables are set in the script when Logstash is installed. Depending on your installation, you might need to update these values.

Configuration options for the `ioaz` Logstash output plugin

The `ioaz` Logstash output plugin, which is provided by IBM Z Operations Analytics, forwards event data from the IBM Common Data Provider for z Systems to the IBM Operations Analytics - Log Analysis server. This reference lists the configuration options that can be set for the plugin in the Logstash configuration file, which is *LOGSTASH_INSTALL_DIR*/config/logstash-ioaz.conf.

Update the *LOGSTASH_INSTALL_DIR*/config/logstash-ioaz.conf file by using a text editor.

After you update the file, to have your configuration changes take effect, run the following command to restart Logstash:

```
LOGSTASH_INSTALL_DIR/bin/logstash_util.sh restart
```

Configuration options

disk_cache_path

The file system path to the directory where data is temporarily held. The ioaz Logstash output plugin writes data to this directory before transmission. The available disk space under the path must be large enough to store bursts of data that are not immediately handled by the Log Analysis server.

Input type
String

Required value?
Yes

Default value
None

idle_flush_time

The maximum time (in seconds) between successive transmissions for a data source.

Input type
Number

Required value?
No

Default value
5

keystore_path

The path to the Log Analysis keystore that contains the key for decrypting the Log Analysis user password.

Input type
String

Required value?
This value is required only if the value of the password option is encrypted.

Default value
None

log_file_count

The maximum number of log files that are stored (in the directory that is specified by the log_path configuration option) before older log files are overwritten.

Tip: The log_file_count value, multiplied by the log_file_limit value, equals the maximum amount of storage space that is used for the log files in the log_path directory.

Input type
Integer

Required value?
No

Default value
20

log_file_limit

The maximum size for each log file that is stored in the directory that is specified by the log_path configuration option.

Tip: The log_file_count value, multiplied by the log_file_limit value, equals the maximum amount of storage space that is used for the log files in the log_path directory.

Input type

Integer

Required value?

No

Default value

10485760 bytes

log_level

The level of logging information. The following values are valid:

- severe
- warning
- info
- event
- debug
- trace

The value of the log_level option is applied each time that Logstash is started.

Input type

String

Required value?

No

Default value

info

log_path

The path to the directory that is used for logging information from the ioaz Logstash output plugin.

Tip: The log_file_count value, multiplied by the log_file_limit value, equals the maximum amount of storage space that is used for the log files in the log_path directory.

Input type

String

Required value?

No

Default value

/tmp/ioaz

password

The Log Analysis user password. The password can be cleartext or encrypted. If the password is encrypted, the value of the keystore_path option must also be specified.

Use the unity_securityUtility.sh command on the Log Analysis server to encrypt the password.

Input type
String

Required value?
No

Default value
`unityadmin`

port
The port on which the `ioaz` Logstash output plugin listens for data from the IBM Common Data Provider for z Systems.

Input type
String

Required value?
No

Default value
`8080`

`trust_all_certificates`
A `true` or `false` indication of whether the `ioaz` Logstash output plugin should trust all security certificates when it sends data to the IBM Operations Analytics - Log Analysis server.

A value of `true` directs the plugin to trust all security certificates that are provided by the Log Analysis server. A value of `false` directs the plugin to use only a specific security certificate.

The default value is `true` because the default behavior is to trust all security certificates.

Input type
Boolean

Required value?
No

Default value
`true`

`url`
The Representational State Transfer (REST) endpoint for the Log Analysis ingestion REST API.

Input type
String

Required value?
Yes

Default value
None

`user`
The Log Analysis user name.

Input type
String

Required value?
No

Default value
`unityadmin`

Securing communication between IBM Common Data Provider for z Systems and Logstash

You must configure a secure data connection for streaming operational data from IBM Common Data Provider for z Systems to Logstash.

Procedure

Depending on your environment and Logstash setup, use one of the following options to secure the data connection:

Option	More information
You want to use the Logstash that is provided by IBM Z Operations Analytics, and you have OpenSSL installed.	To enable SSL, uncomment the SSL statements in the Logstash configuration file (`*LOGSTASH_INSTALL_DIR*/config/logstash-ioaz.conf`). As part of the Logstash installation, the Logstash installer then secures the Logstash end of the data connection. To secure the IBM Common Data Provider for z Systems end of the data connection, complete the steps that are outlined in the IBM Common Data Provider for z Systems V1.1.0 documentation.
You want to use the Logstash that is provided by IBM Z Operations Analytics, but you do *not* have OpenSSL installed.	To secure the data connection, complete the steps that are outlined in the IBM Common Data Provider for z Systems V1.1.0 documentation.
You want to use an existing installation of Logstash rather than the Logstash that is provided by IBM Z Operations Analytics.	To secure the data connection, complete the steps that are outlined in the IBM Common Data Provider for z Systems V1.1.0 documentation.

Configuring the Problem Insights extension

You can update the configuration for the Problem Insights extension to change the interval for refreshing the Problem Insights data cache, or to improve the search capabilities for the Problem Insights.

Before you begin

Install the Problem Insights extension, as described in "Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer" on page 25.

Procedure

1. To update your configuration, change the values of the following parameters in the file `*LA_INSTALL_DIR*/wlp/usr/servers/Unity/apps/Unity.war/WEB-INF/unitysetup.properties`.

Option	Parameter value to change
To change the interval for refreshing the Problem Insights data cache	`INSIGHTS_DATA_CACHE_REFRESH_INTERVAL` The default (and minimum) value for this parameter is 5, which means 5 minutes. If you specify a value that is less than 5, the default value of 5 is used instead.
To improve the search capabilities for the Problem Insights	• `MIN_SOLR_HEAP_SIZE` Set this value to a minimum of 2048. • `MAX_SOLR_HEAP_SIZE` Set this value to a minimum of 4096.

2. To have updates take effect, restart the Log Analysis server and all indexing engines.

Preparing to analyze z/OS log data

To begin analyzing z/OS log data, you must log in to IBM Operations Analytics - Log Analysis. To optimize troubleshooting in your IT operations environment, assign data sources to groups within the Log Analysis user interface, and customize the Custom Search Dashboards that are provided in the z/OS Insight Packs.

About this task

You might also want to use the Problem Insights page and the client-side Expert Advice in the Log Analysis UI.

Logging in to Log Analysis

To create or update data sources (including to group data sources), you must log in to IBM Operations Analytics - Log Analysis with a user name that has administrator authority. If you are only analyzing z/OS log data, the user name that you use does not require administrator authority.

Before you begin

For information about administering, using, and troubleshooting IBM Operations Analytics - Log Analysis, see the Log Analysis documentation.

Procedure

To log in to IBM Operations Analytics - Log Analysis for the first time, use the following URL:

```
https://fully_qualified_ioala_hostname:secure_port/Unity
```

where *`fully_qualified_ioala_hostname`* is the fully qualified domain name of the IBM Operations Analytics - Log Analysis server, and *`secure_port`* is the Web Console secure port that is defined during the installation of IBM Operations Analytics - Log Analysis. The default value for this port is `9987`.

Use of cookies in the Log Analysis UI

In the IBM Operations Analytics - Log Analysis UI, you can enable the search history. Log Analysis then uses a cookie that is saved in the temporary directory of the browser to remember the last 10 terms that are entered in the search field.

By default, the cookie that saves the search terms expires every 30 days.

For information about how to enable or clear this search history, see the information about enabling the GUI search history in the Enabling the GUI search history.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases, no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled “Cookies, Web Beacons and Other Technologies,” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

Grouping data sources to optimize troubleshooting in your IT environment

By defining logical groups of data sources within IBM Operations Analytics - Log Analysis, you can more easily apply searches to related sets of data sources to optimize troubleshooting in your IT environment.

About this task

Determine whether any data sources can be organized into meaningful logical groups.

For example, assume that your IT environment contains an online banking application with a WebSphere Application Server for z/OS front end and a Db2 for z/OS back end. Organizing the associated data sources into a group can help you to quickly focus your troubleshooting efforts on the online banking application.

Although you can configure the `ioaz` Logstash output plugin to create the necessary data sources if they do not exist, the plugin does not group the data sources. To group the data sources that are created by the `ioaz` Logstash output plugin, you must manually assign the data sources to groups by using Log Analysis user interface.

For more information about grouping data sources, see Editing service topology information in Group JSON in the Log Analysis documentation.

Extending troubleshooting capability with Custom Search Dashboard applications

The Custom Search Dashboard applications that are provided in the z/OS Insight Packs are samples of custom logic that you can use to extend the troubleshooting capability of IBM Operations Analytics - Log Analysis. You can use these samples to create dashboards for presenting the generated data in a useful visual format, such as a chart or graph, and for presenting HTML content.

Before you begin

For information about the content of each dashboard, see "Dashboards that represent the operational data" on page 141.

For information about the data source types that are used in populating the dashboards, see "Data sources that contribute to the operational insights" on page 109.

For more information about Custom Search Dashboards, see Custom Search Dashboards in the Log Analysis documentation.

About this task

Before you run these applications, you must customize them for your environment.

Each dashboard application is defined in a JavaScript Object Notation (JSON) file that specifies the following information:

- The script that runs the custom logic
- The parameters to pass to the script
- The output charts to display in the dashboard

To update the parameters for a dashboard application, edit the appropriate `*.app` file in the following directories:

WebSphere Application Server for z/OS Custom Search Dashboard applications
`LA_INSTALL_DIR/AppFramework/Apps/WASforzOSInsightPack_3.2.0.0`

z/OS Network Custom Search Dashboard applications
`LA_INSTALL_DIR/AppFramework/Apps/zOSNetworkInsightPack_v3.2.0.0`

z/OS SMF Custom Search Dashboard applications
`LA_INSTALL_DIR/AppFramework/Apps/SMFforzOSInsightPack_v3.2.0.0`

z/OS SYSLOG Custom Search Dashboard applications
`LA_INSTALL_DIR/AppFramework/Apps/SYSLOGforzOSInsightPack_3.2.0.0`

Customizing the dashboard applications

Before you run the Custom Search Dashboard applications, customize them for your environment.

Before you begin

In this topic, the Db2 for z/OS Troubleshooting dashboard is used as an example to show you how to customize a dashboard application.

If you are customizing the SYSLOG for z/OS Time Comparison dashboard, also see "Customizing the SYSLOG for z/OS Time Comparison dashboard" on page 53.

About this task

To customize a Troubleshooting application, you must update the parameters for the following items:

data source
: The data source, or group of data sources, to use as input for the Troubleshooting dashboard

time interval
: The time interval for which to extract data to present in the Troubleshooting dashboard

host name (optional)
: The host name for which to group log data in the Troubleshooting dashboard. Specifying this parameter is optional, but grouping log data by host name can help you identify where problems are occurring.

Procedure

1. To update the parameters, edit the following file:

   LA_INSTALL_DIR/AppFramework/Apps/SYSLOGforzOSInsightPack_3.2.0.0/DB2_for_zOS_Troubleshooting.app

   In the **search** parameter, you specify the data source, or group of data sources, to be used as input for the Db2 for z/OS Troubleshooting dashboard.

   Tip: The name of the data source or group of data sources that you specify must match the name that was defined in the **Data Sources** tab of the Administrative Settings workspace in IBM Operations Analytics - Log Analysis.
2. To specify the data sources to be used for the dashboard data, use one of the following formats, depending on whether you want to specify data sources individually or as a group:
   - To specify individual data sources, use the following format:

     ```
     "parameters": [
         {
         "name": "search",
         "type": "SearchQuery",
         "value": {
             "logsources": [
                 {
                     "type": "logSource",
                     "name": "SYSLOG1"
                 },
                 {
                     "type": "logSource",
                     "name": "SYSLOG2"
                 }
             ]
         }
     },
     ```

     Specify the following values for the **search** parameter:

     type logSource

     name The name of the data source that was defined in the **Data Sources** tab of the Administrative Settings workspace in IBM Operations Analytics - Log Analysis. An example value is SYSLOG.
   - To specify a group of data sources, use the following format:

     ```
     "parameters": [
         {
         "name": "search",
         "type": "SearchQuery",
     ```

```
    "value": {
        "logsources": [
            {
                "type": "tag",
                "name": "/day trader/trading application1/"
            },
            {
                "type": "tag",
                "name": "/day trader/trading application2/"
            }
        ]
    }
},
```

Specify the following values for the **search** parameter:

type tag

name The tag path that was defined in the **Data Sources** tab of the Administrative Settings workspace in IBM Operations Analytics - Log Analysis. An example value is /day trader/trading application/.

3. To specify the time interval for which to extract data to present in the dashboard, define either a relative time interval or a custom time interval.
   - To specify a relative time interval, use the **relativeTimeInterval** parameter, as shown in the following format:

```
"parameters": [
    {
    "name": "search",
    "type": "SearchQuery",
    "value": {
        "logsources": [
            {
                "type": "logSource",
                "name": "SYSLOG"
            }
        ]
    }
},
{
    "name": "relativeTimeInterval",
    "type": "string",
    "value": "LastDay"
},
{
    "name": "timeFormat",
    "type": "data",
    "value": {
        "timeUnit": "hour",
        "timeUnitFormat": "MM-dd HH:mm"
    }
},
{
   "name": "hostnameField",
    "type": "string",
    "value": "SystemName"
}
],
```

The following values are valid for the **relativeTimeInterval** parameter:

- LastQuarterHour
- LastHour
- LastDay
- LastWeek

    - `LastMonth`
    - `LastYear`
  - To specify a custom time interval, use the *timestamp* value in the **`search`** parameter, as shown in the following format:

```
"parameters": [
    {
    "name": "search",
    "type": "SearchQuery",
    "value": {
        "filter":{
            "range":{
                "timestamp":{
                    "from":"01/01/2018 00:00:00.000 -0400",
                    "to":"04/25/2018 00:00:00.000 -0400",
                    "dateFormat":"MM/dd/yyyy HH:mm:ss.SSS Z"
                }
            }
        },
        "logsources": [
            {
                "type": "logSource",
                "name": "SYSLOG"
            }
        ]
    }
},
{
    "name": "timeFormat",
    "type": "data",
    "value": {
        "timeUnit": "hour",
        "timeUnitFormat": "yyyy-MM-dd HH:mm:ss"
    }
},
{
    "name": "hostnameField",
    "type": "string",
    "value": "SystemName"
}
],
```

    Specify the following values:

    from The start of the time interval for which data is extracted to present in the dashboard

    to The end of the time interval for which data is extracted to present in the dashboard

    dateFormat
    The date format string that is used in the *from* and *to* fields

4. The time format defines how the time interval (either relative or custom) is displayed on the x-axis of a chart in the dashboard. To specify the time format, use the **`timeFormat`** parameter, as shown in the following format:

```
{
    "name": "timeFormat",
    "type": "data",
    "value": {
        "timeUnit": "hour",
        "timeUnitFormat": "MM-dd HH:mm"
    }
},
```

   Specify the following values:

timeUnit

The discrete time unit that is displayed on the x-axis of a chart in the dashboard. The following values are valid:

- `minute`
- `hour`
- `day`
- `week`
- `month`
- `year`

The time unit should be a smaller unit than the time interval. For example, if the chart displays a time interval of `LastWeek`, a reasonable time unit is `day`.

timeUnitFormat

The date format of the time unit that is displayed on the x-axis of a chart in the dashboard, as specified by the Java `SimpleDateFormat` class. For example, a date format might be `yyyy-MM-dd HH:mm:ssZ`.

5. Optional: To specify the host name, use the **`hostnameField`** parameter, as shown in the following format:

```
{
   "name": "hostnameField",
   "type": "string",
   "value": "SystemName"
 }
```

Specify one of the following values for the host name, or accept the default value, which is *`SystemName`*:

`hostname`

The host name that is specified in the Service Topology that is associated with the data source.

`SystemName`

The system name.

6. Ensure that the port on which IBM Operations Analytics - Log Analysis listens is set to the default value of 9987 in the following file:

`LA_INSTALL_DIR``/AppFramework/Apps/SYSLOGforzOSInsightPack_3.2.0.0/CommonAppMod.py`

For example, ensure that the following line of the file contains the correct port number:

```
baseurl = 'https://localhost:9987/Unity'
```

Customizing the SYSLOG for z/OS Time Comparison dashboard:

To customize the SYSLOG for z/OS Time Comparison dashboard, you must first complete the basic customization for the SYSLOG for z/OS dashboard. Then, update the parameter for the time filter in the `*.app` file.

Before you begin

For more information about the basic customization, see “Customizing the dashboard applications” on page 49.

About this task

To specify a time filter, use the **timefilters** parameter, as shown in the following format. The *lastnum* value specifies the number of days between the two time periods for the comparison.

```
"filter": {
      "timefilters":
            {
               "lastnum": 1,
               "granularity": "days",
               "type": "relative"
             }
           }
```

You can use the SYSLOG for z/OS Time Comparison dashboard in either of the following ways:

- If you have an active set of search results in the Log Analysis user interface, this application uses the time filter, the data set filter, and the query from that search result as the base for time period 1. Time period 2 is then based on what you specify for the time filter in the *.app file. In the preceding example, time period 2 is one day.
- If you have no active search results in the Log Analysis user interface, time period 1 is based on the values for the **logsources** and the **relativeTimeInterval** parameters.

Procedure

To update the parameters, edit the appropriate *.app file in the following directory:

```
LA_INSTALL_DIR/AppFramework/Apps/SYSLOGforzOSInsightPack_3.2.0.0
```

Running the dashboard applications

You run the Custom Search Dashboard applications from the IBM Operations Analytics - Log Analysis user interface.

Procedure

To run the applications, complete the following steps:

1. Log in to the IBM Operations Analytics - Log Analysis Search workspace.
2. In the **Search Dashboards** section, expand one of the following folders, depending on which applications you want to run:
   - **SMFforzOSInsightPack_v3.2.0.0**
   - **SYSLOGforzOSInsightPack_v3.2.0.0**
   - **WASforzOSInsightPack_v3.2.0.0**
   - **zOSNetworkInsightPack_v3.2.0.0**

   Tip: If you do not see the folder, refresh the contents of this section.
3. Double-click the name of the application.

   Tip: If you hover over the application tab in the workspace, the tooltip displays the time that the application was run, which is an indication of the time that the data for the charts was generated.

Getting started with the IBM zAware data gatherer

Before you can use the IBM zAware data gatherer and see visualizations of the resulting interval anomaly data in the IBM Operations Analytics - Log Analysis user interface, you must verify that all prerequisites are met, and start the data gatherer.

Before you begin

Verify that the following prerequisites for using the IBM zAware data gatherer are met:

__ 1. The IBM zAware V3.1 feature is installed. See "IBM zAware V3.1 feature" on page 15.

__ 2. IBM Operations Analytics - Log Analysis Version 1.3.5 Fix Pack 1 (V1.3.5.1) is installed. See "Installing Log Analysis" on page 23.

__ 3. The z/OS Insight Packs, extensions, and IBM zAware data gatherer are installed. See "Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer" on page 25.

__ 4. The IBM zAware data gatherer is correctly configured with the definition for the IBM zAware server from which you plan to gather interval anomaly data. See "Configuring the IBM zAware data gatherer" on page 37.

Tip: The data gatherer requires a version of Python that is supported by Log Analysis. For more information about the versions, see Verifying the Python version in the Log Analysis documentation.

About this task

The IBM zAware data gatherer files are in the following directory on the Log Analysis server:

```
LA_INSTALL_DIR/zAwareDataGatherer
```

Logs are in the following directory:

```
LA_INSTALL_DIR/zAwareDataGatherer/logs
```

Usage notes for the data gatherer:

- The data gatherer creates a data source with the name "`zOS Anomaly Interval`" and the type `zOS-Anomaly-Interval`. All interval anomaly data that is retrieved by the data gatherer is ingested into this data source. If another data source with the same name but with a different type was previously defined to the Log Analysis server, the IBM zAware data gatherer cannot gather data.
- The data gatherer gathers data only if the following conditions are met:
  - The data gatherer is running.
  - The IBM zAware server from which data is being gathered is accessible.
  - The Log Analysis server to which data is being sent for analysis and visualization is accessible.
- The data gatherer does not support the retrieval of historical interval anomaly data.
- The data gatherer obtains interval anomaly data only for z/OS systems that are monitored by IBM zAware. It does not obtain Linux system data.

Procedure

To start the IBM zAware data gatherer, complete the following steps:

1. Verify that you are logged in to the Linux computer system with the non-root user ID that was used to install Log Analysis.
2. Run the following command from the *LA_INSTALL_DIR*/zAwareDataGatherer/bin directory, where *zaware_server* represents either the host name or the IP address of the IBM zAware server from which you want to gather interval anomaly data:

   ```
   python zAwareDataGatherer.py zaware_server
   ```

   Tip: The data gatherer gathers interval anomaly data from only one IBM zAware server. Therefore, the data gatherer script requires only one argument, which is for specifying the host name or IP address of the IBM zAware server from which you want to gather interval anomaly data. This server must have a configuration entry in the following file:

   ```
   LA_INSTALL_DIR/zAwareDataGatherer/config/zAwareConfig.json
   ```

Getting started with Problem Insights for z/OS

If the Problem Insights component of IBM Z Operations Analytics is installed, the IBM Operations Analytics - Log Analysis UI includes a new tab that is titled **Problem Insights**. For each sysplex from which data is being forwarded to the Log Analysis server, the Problem Insights page includes insight about certain problems that are identified in the ingested data.

Before you begin

The Problem Insights extension must be installed.

For more information about the Problem Insights component and the system requirements and installation information for this component, see the following topics:

- "Log Analysis extensions for z/OS Problem Insights and client-side Expert Advice" on page 13
- "System requirements" on page 18
- "Installing Log Analysis" on page 23
- "Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer" on page 25

Procedure

To use the Problem Insights page in the Log Analysis UI, click the tab that is titled **Problem Insights**.

Getting started with client-side Expert Advice

If the client-side Expert Advice component of IBM Z Operations Analytics is installed, **IBMSupportPortal-ExpertAdvice on Client** is a choice under **Expert Advice** in the Search workspace of the IBM Operations Analytics - Log Analysis UI. With this extension, you can access Expert Advice even if the Log Analysis server does not have access to the Internet.

Before you begin

Ensure that the client-side Expert Advice component is installed. Also, ensure that you are using Log Analysis Version 1.3.5. For more information about the client-side Expert Advice component and the system requirements and installation information for this component, see the following topics:

- "Log Analysis extensions for z/OS Problem Insights and client-side Expert Advice" on page 13
- "System requirements" on page 18
- "Installing Log Analysis" on page 23
- "Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer" on page 25

About this task

The only field that you can update in the `IBMSupportPortal-ExpertAdvice on Client.app` file is `_TERM_LIMIT`, which defines the maximum number of keywords that a user can search for at one time.

The following data is removed from the search keywords by the client-side Expert Advice (`IBMSupportPortal-ExpertAdvice on Client.app`) to increase the chance of a successful search on the client:

- URLs
- File names
- File paths
- IP addresses
- Numbers
- Double spaces

Tips for using the client-side Expert Advice:

1. Verify that the client browser settings allow pop-up windows from the Log Analysis server.
2. If the Log Analysis server cannot be reached, clear cookies for that domain, and again, try to connect.

Procedure

To use the client-side Expert Advice, complete the following steps:

1. In the `_TERM_LIMIT` field of the `IBMSupportPortal-ExpertAdvice on Client.app` file, define the maximum number of keywords that a user can search for at one time.
2. To launch the client-side Expert Advice, click **IBMSupportPortal-ExpertAdvice on Client** under **Expert Advice** in the Custom Search Dashboards panel of the left navigation pane of the Search workspace.

   The client browser sends search requests directly to the IBM Support Portal and opens a new browser tab to display the query search results.

Z Operations Analytics on the Elastic Stack and Splunk platforms

On both the Elastic Stack and Splunk platforms, IBM Z Operations Analytics provides dashboards and searches for Z operational insights.

- "Component overview"
- "Flow of source data among Elastic Stack components"
- "Flow of source data among Splunk components" on page 60
- "Summary of system requirements" on page 61

Component overview

For each platform, IBM Z Operations Analytics includes the following components, which must be deployed in this order:

1. IBM Common Data Provider for z Systems

   See "Planning for configuration of IBM Common Data Provider for z Systems" on page 3.
2. IBM Z Operations Analytics application

   See "Deploying the Z Operations Analytics application on the Elastic Stack platform" on page 62 or "Deploying the Z Operations Analytics application on the Splunk platform" on page 66.
3. Problem Insights Framework

   See "Problem Insights Framework" on page 80.

Flow of source data among Elastic Stack components

Figure 2 on page 60 illustrates the flow of data among the primary components of IBM Z Operations Analytics on the Elastic Stack platform.

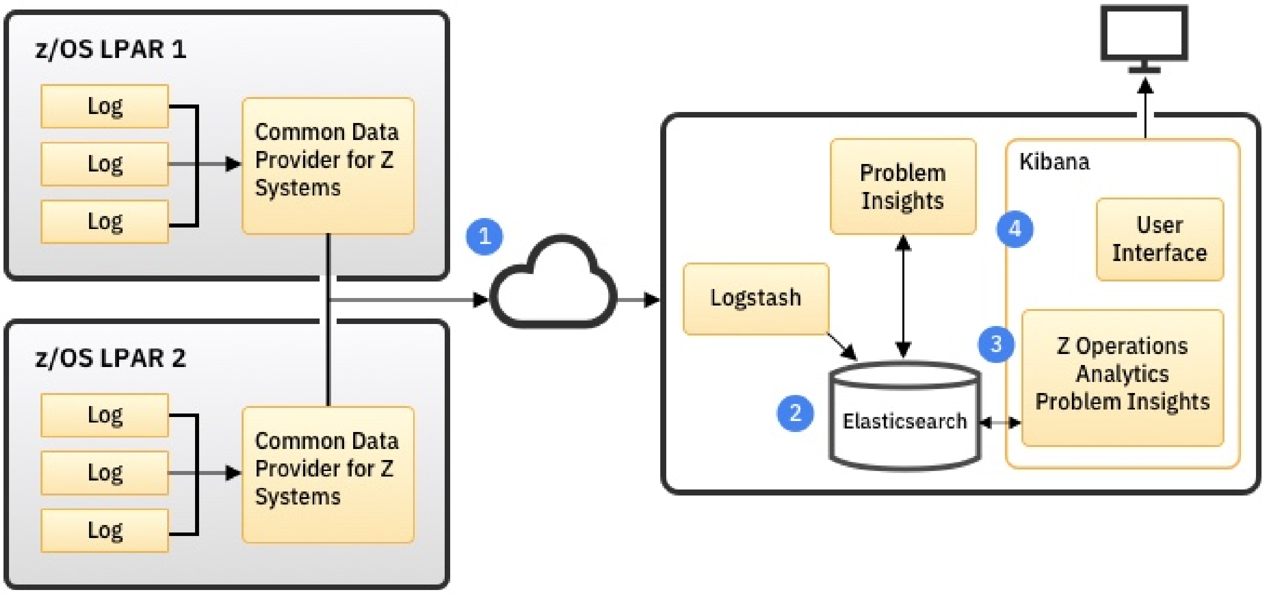
z/OS LPAR 1
Log
Log
Log
Common Data Provider for Z Systems
z/OS LPAR 2
Log
Log
Log
Common Data Provider for Z Systems
1
Logstash
Problem Insights
Kibana
4
User Interface
3
Z Operations Analytics Problem Insights
2
Elasticsearch

Figure 2. Flow of source data among IBM Z Operations Analytics components on the Elastic Stack platform

The following steps describe the data flow among components, which is indicated by arrows in the illustration:

1. In each z/OS logical partition (LPAR), the IBM Common Data Provider for z Systems retrieves the data from the respective source and sends it to the Elastic Stack server.
2. The source data is processed by Logstash according to IBM Z Operations Analytics definitions and is forwarded to Elasticsearch.
3. The system searches the Problem Insights for known problems and presents them in the IBM Z Operations Analytics Problem Insights dashboard.
4. Users can see predefined searches and visualizations of the data in Kibana. Insights are provided for data from the following source types:
   - z/OS system log (SYSLOG)
   - CICS Transaction Server for z/OS EYULOG or MSGUSR log data
   - Network data, such as data from UNIX System Services system log (`syslogd`) or z/OS Communications Server
   - NetView for z/OS message data
   - SMF data
   - WebSphere Application Server for z/OS logs that include SYSOUT or SYSPRINT log data

Flow of source data among Splunk components

Figure 3 on page 61 illustrates the flow of data among the primary components of IBM Z Operations Analytics on the Splunk platform.

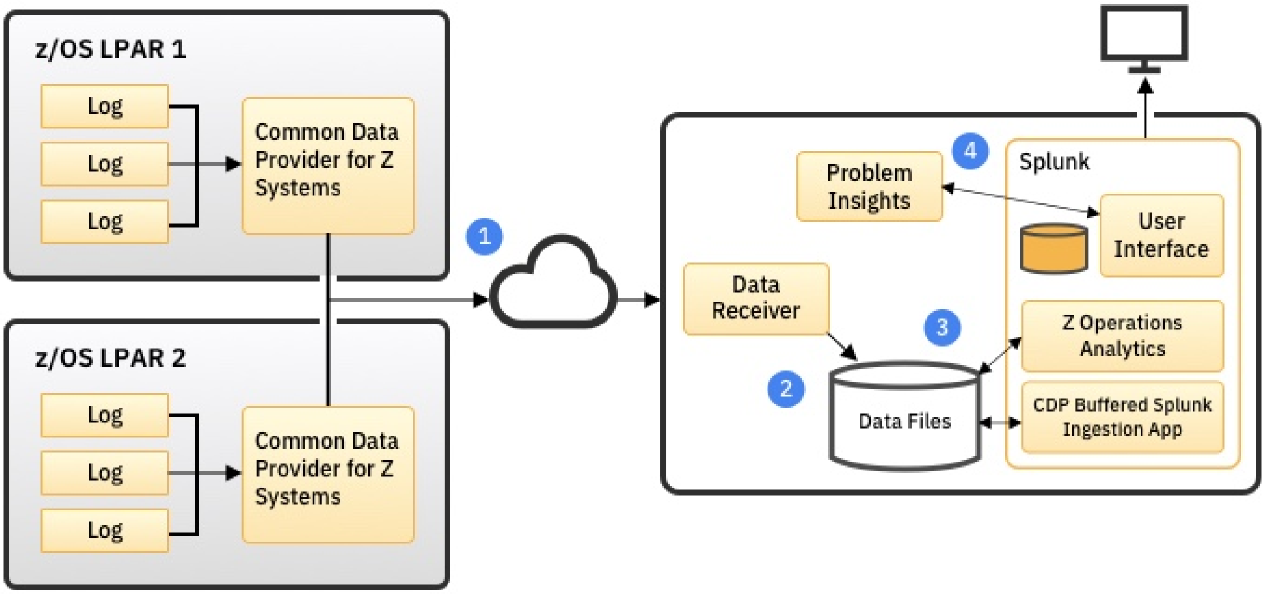

Figure 3. Flow of source data among IBM Z Operations Analytics components on the Splunk platform

The following steps describe the data flow among components, which is indicated by arrows in the illustration:

1. In each z/OS logical partition (LPAR), the IBM Common Data Provider for z Systems retrieves the data from the respective source and sends it to the Splunk Enterprise server.
2. The source data is received by the IBM Common Data Provider for z Systems Data Receiver and is written to local data files.
3. Splunk reads and processes the local data files based on rules that are provided by IBM Z Operations Analytics and the IBM Common Data Provider for z Systems Buffered Splunk Ingestion App.
4. Users can see predefined searches and visualizations of the data in the Splunk user interface. Insights are provided for data from the following source types:
   - z/OS system log (SYSLOG)
   - CICS Transaction Server for z/OS EYULOG or MSGUSR log data
   - Network data, such as data from UNIX System Services system log (`syslogd`) or z/OS Communications Server
   - NetView for z/OS message data
   - SMF data
   - WebSphere Application Server for z/OS logs that include SYSOUT or SYSPRINT log data
5. The system searches the Problem Insights for known problems and presents them in the IBM Z Operations Analytics Problem Insights dashboard.

Summary of system requirements

IBM Common Data Provider for z Systems

See Planning to use IBM Common Data Provider for z Systems in the IBM Common Data Provider for z Systems V1.1.0 documentation.

z/OS systems and subsystems: version requirements

See "Software version requirements for the z/OS systems and subsystems from which operational data is gathered" on page 62.

IBM Z Operations Analytics application on the Elastic Stack platform
See "System requirements for deploying the IBM Z Operations Analytics application on the Elastic Stack platform" on page 66.

IBM Z Operations Analytics application on the Splunk platform
See "System requirements for deploying the IBM Z Operations Analytics application on the Splunk platform" on page 66.

IBM Z Operations Analytics Problem Insights Framework
See "System requirements for the Problem Insights Framework" on page 81.

Software version requirements for the z/OS systems and subsystems from which operational data is gathered

Your environment must meet the software version requirements for the z/OS systems and subsystems from which you want to gather operational data.

The following software versions are required:

- IBM z/OS 2.2 and 2.3
- IBM CICS Transaction Server for z/OS 5.1.1, 5.2, 5.3, and 5.4
- IBM Db2 for z/OS 11.1 and 12.1
- IBM IMS for z/OS 13.1, 14.1, and 15.1
- IBM MQ for z/OS 8.0 and 9.0
- IBM Tivoli NetView for z/OS 6.2 and 6.2.1
- IBM WebSphere Application Server for z/OS 8.5.5 and 9.0
- Access Monitor component of IBM Security zSecure Admin 2.2.1 with APAR OA52273, 2.3.0, or later

Deploying the Z Operations Analytics application on the Elastic Stack platform

To deploy the IBM Z Operations Analytics application, you must extract the IBM Z Operations Analytics package `ZOA-Elastic-V3.2.0.zip` on the Kibana server.

Before you begin

Verify that the system requirements are met, as described in "System requirements for deploying the IBM Z Operations Analytics application on the Elastic Stack platform" on page 66, and that all prerequisite software is configured and is running.

Tips:

- Elasticsearch can be installed by a root user, but it must be run by a non-root user.
- Python must be in the system path.

About this task

You must use the Logstash configuration files from the IBM Z Operations Analytics Elastic Stack ingestion kit. The IBM Z Operations Analytics Elastic Stack ingestion kit is a simpler version of the IBM Common Data Provider for z Systems

Elasticsearch ingestion kit. It contains only the Logstash configuration files that support the IBM Z Operations Analytics dashboards and visualization.

The IBM Z Operations Analytics Elastic Stack ingestion kit contains a Logstash configuration file `Q_elasticsearch.conf` that has the following parameters:

```
index => "zoa-%{sourceType}-%{host}-%{+YYYYMMdd}"
```

These default values can be used when you start Logstash.

If you are using the Elasticsearch ingestion kit for IBM Common Data Provider for z Systems: You must replace the Logstash configuration files in the IBM Common Data Provider for z Systems Elasticsearch ingestion kit with the Logstash configuration files in the IBM Z Operations Analytics Elastic Stack ingestion kit. The IBM Common Data Provider for z Systems Elasticsearch ingestion kit contains a Logstash configuration file `Q_CDPz_Elastic.lsh` that has the following parameters:

```
index => "cdp-%([@metadata][indexname]}-%"{yyyyMMdd}"
```

In the Logstash configuration file `Q_elasticsearch.conf` from the IBM Z Operations Analytics Elastic Stack ingestion kit, the value `cdp` in the parameter is changed to `zoa`.

If you have existing host data that was previously sent to Elasticsearch by using the IBM Common Data Provider for z Systems Elasticsearch ingestion kit, that data is not affected by this installation. You can continue to view your existing data. You can start Logstash to ingest new operational data for either IBM Common Data Provider for z Systems or IBM Z Operations Analytics, but this new data uses the following index value:

```
index => "zoa-%{sourceType}-%{host}-%{+YYYYMMdd}"
```

Procedure

To deploy the IBM Z Operations Analytics application, complete the following steps. These steps are based on deployment on a Linux system. Use comparable steps if you are deploying on a Windows system.

1. Log in to the Logstash server, and extract the IBM Z Operations Analytics Elastic Stack ingestion kit, which is in the file `ZOA-IngestionKit-3.2.0.zip`.
2. Configure Logstash according to the following instructions, depending on the type of your Logstash image.

Type of Logstash image	Instructions

.deb	Run the following commands, where *config_file_dir* is the directory where you extracted the IBM Z Operations Analytics Elastic Stack ingestion kit:
	`# /etc/init.d/logstash stop`
	`# mv config_file_dir /etc/logstash/conf.d`
	`# /etc/init.d/logstash start`

.rpm	Run the following commands, where *config_file_dir* is the directory where you extracted the IBM Z Operations Analytics Elastic Stack ingestion kit:
	`# /etc/init.d/logstash stop`
	`# mv config_file_dir /etc/logstash/conf.d`
	`# /etc/init.d/logstash start`

Type of Logstash image	Instructions
.tar.gz	From the installation directory, run the following command, where *config_file_dir* is the directory where you extracted the IBM Z Operations Analytics Elastic Stack ingestion kit: `# bin/logstash -f config_file_dir`
.zip	From the installation directory, run the following command, where *config_file_dir* is the directory where you extracted the IBM Z Operations Analytics Elastic Stack ingestion kit: `# bin/logstash -f config_file_dir`

3. Verify all index names.

   The Elasticsearch index name is defined by the following parameter in the Logstash configuration file `Q_elasticsearch.conf`:

   ```
   index => "zoa-%{sourceType}-%{host}-%{+YYYYMMdd}"
   ```

   IBM Z Operations Analytics supports this default index name and all index names that match the pattern `zoa-*`. If you use customized index names, they must also use the pattern `zoa-*`.
4. Log in to the Kibana server.
5. Verify that Elasticsearch, Logstash, and Kibana are running. To verify the status of Elasticsearch, Logstash, and Kibana, run the following commands, which list each process, if that process is running:

   Elasticsearch
   ```
   # ps -ef | grep elasticsearch
   ```

   Logstash
   ```
   # ps -ef | grep logstash
   ```

   Kibana
   ```
   # ps -ef | grep node
   ```
6. Extract the IBM Z Operations Analytics package `ZOA-Elastic-V3.2.0.zip` on the Kibana server.
7. Verify that the following files are in the directory where they were extracted on the Kibana server:
   - License
   - `setup.conf`
   - `setup.py`
   - `resource`
8. In the directory where the files are extracted, complete the following steps to import the IBM Z Operations Analytics dashboards:
   a. In the `setup.conf` file, provide the following parameter values, which are the configuration values for the `setup.py` file:

      `host_kibana`
      The IP address where Kibana is bound.

      `host_es`
      The IP address where Elasticsearch is bound.

      `port_kibana`
      The port number that is used by Kibana.

      `port_es`
      The port number that is used by Elasticsearch.

dir_kibana
The absolute path for the Kibana home directory.

login
The default value of **login** is an empty string. If authentication, such as X-Pack, is enabled, provide credentials by adding *user*:*password* for the value of **login**.

pi_ip
The IP address where the Problem Insights Framework is bound.

b. Import the IBM Z Operations Analytics dashboards by running the following command:

```
# python setup.py import
```

9. Open the Kibana URL in a browser to verify that the index pattern, dashboards, and visualizations were created. The index pattern must be `zoa-*`.

   Verify that the following dashboards are included:

   - CICS Transaction Server for z/OS Enterprise Dashboard by Region
   - CICS Transaction Server for z/OS Enterprise Dashboard by System
   - CICS Transaction Server for z/OS System Dashboard
   - CICS Transaction Server for z/OS Region Dashboard
   - CICS Transaction Server for z/OS Transaction Dashboard
   - CICS Transaction Server for z/OS Job Dashboard
   - Db2 for z/OS Enterprise Dashboard by Subsystem
   - Db2 for z/OS Enterprise Dashboard by System
   - Db2 for z/OS System Dashboard
   - Db2 for z/OS Subsystem Dashboard
   - Db2 for z/OS Job Dashboard
   - IMS for z/OS Job Dashboard
   - MQ for z/OS Job Dashboard
   - Saved Searches Dashboard
   - Systems Dashboard
   - Welcome Dashboard
   - z/OS Job Dashboard
   - z/OS Security Server RACF Dashboard
   - zSecure Access Monitor Dashboard

10. Optional: Install the IBM Z Operations Analytics Problem Insights Framework. For more information, see "Problem Insights Framework" on page 80.

Results

When the IBM Z Operations Analytics application is successfully deployed, operational insights are available in the Kibana dashboards.

Tip: If you ever need to delete the IBM Z Operations Analytics application, run the following command:

```
# python setup.py delete
```

System requirements for deploying the IBM Z Operations Analytics application on the Elastic Stack platform

Your environment must meet the system requirements for deploying the IBM Z Operations Analytics application on the Elastic Stack platform.

The IBM Z Operations Analytics application can be run on a Linux or Windows system and must be run with the following software:

- Elastic Stack 6.1, 6.2, or 6.3
- Java Runtime Environment (JRE) 8 or later
- Python 2.6 or later

Deploying the Z Operations Analytics application on the Splunk platform

You can configure the Splunk environment in different ways depending on volume of data, number of users and searches, system availability, and disaster recovery. Two options for deploying the Z Operations Analytics application are highlighted.

Before you begin

Verify that the system requirements are met, as described in "System requirements for deploying the IBM Z Operations Analytics application on the Splunk platform," and that all prerequisite software is configured and is running.

About this task

The following deployment options are highlighted. You can also use this information to configure your Splunk environment by using other options, as described in Types of distributed deployments in the Splunk documentation.

Single Splunk Enterprise system
See "Deploying Z Operations Analytics on a single Splunk Enterprise system" on page 67.

Clustered Splunk environment
See "Deploying Z Operations Analytics in a clustered Splunk environment" on page 68.

As part of your deployment, you can also install the following items:

Problem Insights Framework
For information about the IBM Z Operations Analytics Problem Insights Framework, see "Problem Insights Framework" on page 80.

Splunk IT Service Intelligence (ITSI) module
For information about the Splunk ITSI module for IBM Z Operations Analytics, see "Splunk ITSI module for IBM Z Operations Analytics" on page 70.

System requirements for deploying the IBM Z Operations Analytics application on the Splunk platform

Your environment must meet the system requirements for deploying the IBM Z Operations Analytics application on the Splunk platform.

The IBM Z Operations Analytics application can be run on a Linux or Windows system and must be run with the following software:

- Splunk Enterprise 6.6, 7.0, or 7.1
- For Splunk IT Service Intelligence (ITSI) module for IBM Z Operations Analytics, Splunk ITSI 3.01 or later is required.
- For the IBM Z Operations Analytics Problem Insights Framework, Java Runtime Environment (JRE) 8 or later is required.

Tip: If you use Splunk Enterprise 7.1, you must use Splunk ITSI 3.1.

Index schema in IBM Z Operations Analytics application on the Splunk platform

To help improve search performance in the IBM Z Operations Analytics application on the Splunk platform, a unique index is defined for each data source type. Each file monitor input is configured so that the same value is used for both the source type and the index name. A macro is also provided for each index, and the name of the macro is the same as the source type and index name values.

The data source type for z/OS SYSLOG Console data is an exception to this naming convention because this source type is included with IBM Common Data Provider for z Systems as part of the Buffered Splunk Ingestion App. The index name for z/OS SYSLOG Console data is `zosdex`.

In addition to the individual source type macros, the macro `cdp_index` is provided and defined as `index=zos*`. This macro can be used to search for data across all z/OS source types.

All of the IBM-provided dashboards and predefined searches use macros to query indexed data. If you change the name of an IBM-provided index, you must also change the corresponding macro definition to point to your new index so that the dashboards and predefined searches continue to show the expected results.

Deploying Z Operations Analytics on a single Splunk Enterprise system

The advantage of deploying IBM Z Operations Analytics on a single Splunk Enterprise system is that the deployment is simple and quick.

About this task

The steps in this procedure must be done on the system where the web browser is running rather than on the Splunk Enterprise server.

Procedure

To deploy the IBM Z Operations Analytics application, complete the following steps:

1. Install and configure the IBM Common Data Provider for z Systems Data Receiver and Buffered Splunk Ingestion App. For more information, see Preparing to send data to Splunk in the IBM Common Data Provider for z Systems V1.1.0 documentation.
2. Mount the IBM Z Operations Analytics ISO installation image, or extract the IBM Z Operations Analytics `.tar` file.
3. Log in to Splunk.

4. From the Splunk Web Home page, click the gear icon that is next to the word "Apps."
5. Select **Install app from file**.
6. Navigate to the ISO image, select the `ibm_zoa_insights.spl` file, and click **Upload**.
7. If you are prompted to restart Splunk Enterprise server, restart it.
8. Verify that the application is shown in the list of apps and add-ons. The application is also in the following directory on the Splunk Enterprise server:

   `$SPLUNK_HOME/etc/apps/ibm_zoa_insights`
9. Optional: Install the Splunk IT Service Intelligence (ITSI) module for IBM Z Operations Analytics. For more information, see "Installing the Splunk ITSI module for IBM Z Operations Analytics" on page 78.
10. Optional: Install the IBM Z Operations Analytics Problem Insights Framework. For more information, see "Problem Insights Framework" on page 80.

Deploying Z Operations Analytics in a clustered Splunk environment

The advantage of deploying IBM Z Operations Analytics in a clustered Splunk environment is that you have increased capacity for analyzing operational data. You also have capability for disaster recovery and environmental redundancy. For example, when you cluster the indexer, some indexers can go offline without having any impact on the capability to search for data.

Procedure

To deploy the IBM Z Operations Analytics application in an indexer cluster environment, complete the following steps:

1. **On the master indexer**, complete the following steps.
   a. Install the IBM Common Data Provider for z Systems Buffered Splunk Ingestion App, as indicated in the following instructions:
      1) From the IBM Common Data Provider for z Systems `/usr/lpp/IBM/cdpz/v1r1m0/DS/LIB` directory, download the IBM Common Data Provider for z Systems Buffered Splunk Ingestion App (which is a part of your SMP/E installation package) in binary mode. The following files contain the App:
         - UNIX system: `ibm_cdpz_buffer_nix.spl`
         - Windows system: `ibm_cdpz_buffer_win.spl`
      2) To install the Buffered Splunk Ingestion App in Splunk, complete the following steps:
         a) Log in to Splunk.
         b) Click the gear icon that is next to the word "Apps."
         c) Select **Install app from file**.
         d) Browse for the file that you downloaded in step 1a1, select that file, and click **Upload**.
         e) When you are prompted, select **Enable now**.
   b. Copy the app files from `$SPLUNK_HOME/etc/apps/ibm_cdpz_buffer` to `$SPLUNK_HOME/etc/master-apps/ibm_cdpz_buffer`.
   c. In the `$SPLUNK_HOME/etc/master-apps/ibm_cdpz_buffer/default` directory, edit the `indexes.conf` file, and add the line `repFactor=auto`.

   d. Use Splunk Web or command line interface (CLI) to distribute the configuration bundle to the peer nodes.

   e. If you are prompted to restart the Splunk indexers, restart them.

   For more information about this process, see Update common peer configurations and apps in the Splunk documentation.

2. **On the search head**, complete the following steps.

   a. Install the IBM Common Data Provider for z Systems Buffered Splunk Ingestion App, as indicated in the following instructions:

      1) From the IBM Common Data Provider for z Systems `/usr/lpp/IBM/cdpz/v1r1m0/DS/LIB` directory, download the IBM Common Data Provider for z Systems Buffered Splunk Ingestion App (which is a part of your SMP/E installation package) in binary mode. The following files contain the App:

         - UNIX system: `ibm_cdpz_buffer_nix.spl`
         - Windows system: `ibm_cdpz_buffer_win.spl`

      2) To install the Buffered Splunk Ingestion App in Splunk, complete the following steps:

         a) Log in to Splunk.

         b) Click the gear icon that is next to the word “Apps.”

         c) Select **Install app from file**.

         d) Browse for the file that you downloaded in step 2a1, select that file, and click **Upload**.

         e) When you are prompted, select **Enable now**.

   b. Install the IBM Z Operations Analytics application, as indicated in the following instructions:

      1) Mount the IBM Z Operations Analytics ISO installation image, or extract the IBM Z Operations Analytics `.tar` file.

      2) Log in to Splunk.

      3) From the Splunk Web Home page, click the gear icon that is next to the word “Apps.”

      4) Select **Install app from file**.

      5) Navigate to the ISO image, select the `ibm_zoa_insights.spl` file, and click **Upload**.

      6) If you are prompted to restart Splunk Enterprise, restart it.

      7) Verify that the add-on is shown in the list of apps and add-ons. The add-on is also in the following directory on the Splunk Enterprise server:

         ```
         $SPLUNK_HOME/etc/apps/ibm_zoa_insights
         ```

3. **On the heavy forwarder**, complete the following steps.

   a. Install and configure the IBM Common Data Provider for z Systems Data Receiver and Buffered Splunk Ingestion App. For more information, see Preparing to send data to Splunk in the IBM Common Data Provider for z Systems V1.1.0 documentation.

   b. Install the IBM Z Operations Analytics application, as indicated in the following instructions:

      1) Mount the IBM Z Operations Analytics ISO installation image, or extract the IBM Z Operations Analytics `.tar` file.

      2) Log in to Splunk.

3) From the Splunk Web Home page, click the gear icon that is next to the word "Apps."
4) Select **Install app from file**.
5) Navigate to the ISO image, select the `ibm_zoa_insights.spl` file, and click **Upload**.
6) If you are prompted to restart Splunk Enterprise, restart it.
7) Verify that the add-on is shown in the list of apps and add-ons. The add-on is also in the following directory on the Splunk Enterprise server:

   ```
   $SPLUNK_HOME/etc/apps/ibm_zoa_insights
   ```

4. Optional: Install the Splunk IT Service Intelligence (ITSI) module for IBM Z Operations Analytics. For more information, see "Installing the Splunk ITSI module for IBM Z Operations Analytics" on page 78.
5. Optional: Install the IBM Z Operations Analytics Problem Insights Framework. For more information, see "Problem Insights Framework" on page 80.

Splunk ITSI module for IBM Z Operations Analytics

The Splunk IT Service Intelligence (ITSI) module for IBM Z Operations Analytics provides key performance indicators (KPIs) for monitoring IBM Z® systems. After you install this module, you can create your own Splunk ITSI service, and add the prebuilt KPIs from this module to your service.

Overview of Splunk ITSI modules and services

Splunk ITSI modules are service templates that provide prebuilt KPIs, entity definitions, and dashboard visualizations. They help ITSI users understand and act on the data that is generated from monitoring services within ITSI.

Splunk ITSI services contain the KPIs that provide the capability to monitor IT service health, perform root cause analysis, receive alerts, and ensure that IT operations are in compliance with service-level agreements (SLAs) for the business. A KPI is a recurring saved search that returns the value of an IT performance metric, such as CPU utilization, response time, or paging rate.

For more information about Splunk ITSI concepts and features, see ITSI concepts and features in the Splunk documentation.

The Splunk ITSI module for IBM Z Operations Analytics includes the service "IBM Z Operations Analytics," which contains the KPIs for monitoring IBM Z systems.

View of Splunk ITSI modules and services in the Splunk ITSI application

In the Splunk ITSI application, you can view the ITSI modules and services in your deployment.

Viewer for all ITSI Modules

In the Splunk ITSI application, click **Configure** > **Modules** to view configuration information about the ITSI modules that are installed in your deployment.

If you install the Splunk ITSI Module for IBM Z Operations Analytics, it is shown in the viewer for all ITSI modules.

Viewer for all Services

In the Splunk ITSI application, click **Configure** > **Services** to view the services in your deployment.

If you install the Splunk ITSI Module for IBM Z Operations Analytics, the "IBM Z Operations Analytics" service is shown in the viewer for all services.

KPIs in the IBM Z Operations Analytics service

The IBM Z Operations Analytics service in the Splunk IT Service Intelligence (ITSI) module for IBM Z Operations Analytics contains 33 key performance indicators (KPIs) that apply to four IBM Z subsystems.

The following list indicates the four IBM Z subsystems with the applicable KPIs:

CICS Transaction Server for z/OS

- 4 KPIs search for data with a source type of `zOS-SMF30`.
- 1 KPI searches for data with a source type of `zOS-SMF110_S_10`.
- 8 KPIs search for data with a source type of `zOS-SMF110_1_SUMMARY`.

Db2 for z/OS

- 4 KPIs search for data with a source type of `zOS-SMF30`.
- 2 KPIs search for data with a source type of `zOS-SMF100_1`.
- 6 KPIs search for data with a source type of `zOS-SMF101_SUMMARY`.

IMS for z/OS

- 4 KPIs search for data with a source type of `zOS-SMF30`.

MQ for z/OS

- 4 KPIs search for data with a source type of `zOS-SMF30`.

These IBM Z KPI searches are defined at the enterprise level. You might want to create entities for your IBM Z system to filter each IBM Z KPI search for your environment. For information about defining entities, see Define entities in ITSI in the Splunk documentation.

Table 5 on page 72 describes the KPIs in the IBM Z Operations Analytics service. It includes each KPI title with a description of the KPI, the name of the associated base search for the KPI (if applicable), and the properties that are defined for the Search and Calculate section. Default values are used for the Threshold and Anomaly Detection sections.

KPI base searches let you share a search definition across multiple KPIs. Where possible, they are used to consolidate similar IBM Z KPIs, reduce search load, and improve search performance. Table 6 on page 77 describes the KPI base searches that are listed in column 3 of Table 5 on page 72, and includes the associated metrics that are defined for each search.

Table 5. KPIs in the Splunk ITSI module for IBM Z Operations Analytics

KPI	Description	KPI base search	Search and calculate
CICS Abend Count	Sum current abends for your CICS regions	Not applicable	• Source type: `zOS-SMF110_1_SUMMARY` • Threshold field: `RECORD_COUNT` • Entity split by field: `sysplex + "." + system + "." + CICS_GEN_APPLID` • Calculation: Calculating sum per entity, sum of aggregate over the last 15 minutes every 15 minutes
CICS CPU Time	Average CPU time for your CICS regions	`IBMZ.CICS.Transaction_Summary`	• Source type: `zOS-SMF110_1_SUMMARY` • Threshold field: `CPU_TIME` • Entity split by field: `sysplex + "." + system + "." + CICS_GEN_APPLID` • Calculating: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
CICS CPU Utilization	Average CPU utilization for your CICS jobs	`IBMZ.CICS.Job`	• Source type: `zOS-SMF30` • Threshold field: `CPU_UTLIZATION` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
CICS Dispatch Time	Average dispatch time for your CICS regions	`IBMZ.CICS.Transaction_Summary`	• Source type: `zOS-SMF110_1_SUMMARY` • Threshold field: `DISPATCH_TIME` • Entity split by field: `sysplex + "." + system + "." + CICS_GEN_APPLID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
CICS Elapsed Time	Average elapsed time for your CICS regions	`IBMZ.CICS.Transaction_Summary`	• Source type: `zOS-SMF110_1_SUMMARY` • Threshold field: `ELAPSED_TIME` • Entity split by field: `sysplex + "." + system + "." + CICS_GEN_APPLID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
CICS I/O Rate	Average I/O rate for your CICS jobs	`IBMZ.CICS.Job`	• Source type: `zOS-SMF30` • Threshold field: `IO_RATE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes

Table 5. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	Description	KPI base search	Search and calculate
CICS Paging Rate	Average paging rate for your CICS jobs	`IBMZ.CICS.Job`	• Source type: `zOS-SMF30` • Threshold field: `PAGING_RATE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
CICS QR TCB CPU to Dispatch Time Ratio	Average QR TCB CPU to dispatch time ratio for your CICS regions	`IBMZ.CICS.Transaction_Summary`	• Source type: `zOS-SMF110_1_SUMMARY` • Threshold field: `QR_CPU_TIME/QR_DISP_TIME` • Entity split by field: `sysplex + "." + system + "." + CICS_GEN_APPLID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
CICS QR TCB Dispatch Time	Average QR TCB dispatch time for your CICS regions	`IBMZ.CICS.Transaction_Summary`	• Source type: `zOS-SMF110_1_SUMMARY` • Threshold field: `QR_DISP_TIME` • Entity split by field: `sysplex + "." + system + "." + CICS_GEN_APPLID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
CICS Response Time	Average response time for your CICS regions	`IBMZ.CICS.Transaction_Summary`	• Source type: `zOS-SMF110_1_SUMMARY` • Threshold field: `DISPATCH_TIME + SUSP_TIME` • Entity split by field: `sysplex + "." + system + "." + CICS_GEN_APPLID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
CICS Transaction Count	Sum transaction count for your CICS systems	Not applicable	• Source type: `zOS-SMF110_S_10` • Threshold field: `TRAN_COUNT` • Entity split by field: `sysplex + "." + system` • Calculation: Calculating sum per entity, sum of aggregate over the last 15 minutes every 15 minutes
CICS Wait Time	Average wait time for your CICS regions	`IBMZ.CICS.Transaction_Summary`	• Source type: `zOS-SMF110_1_SUMMARY` • Threshold field: `SUSP_TIME` • Entity split by field: `sysplex + "." + system + "." + CICS_GEN_APPLID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes

Table 5. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	Description	KPI base search	Search and calculate
CICS Working Set Size	Average working set size for your CICS jobs	`IBMZ.CICS.Job`	• Source type: `zOS-SMF30` • Threshold field: `WORKING_SET_SIZE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 CPU Time	Average CPU time for your Db2 subsystems	`IBMZ.Db2.Accounting_Summary`	• Source type: `zOS-SMF101_SUMMARY` • Threshold field: `CP_CPU_SEC_CL1 + SQLCALL_SEC_INSP + UDF_REQUESTS_SEC` • Entity split by field: `sysplex + "." + system + "." + SSID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 CPU Utilization	Average CPU utilization for your Db2 jobs	`IBMZ.Db2.Job`	• Source type: `zOS-SMF30` • Threshold field: `CPU_UTLIZATION` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 Elapsed Time	Average elapsed time for your Db2 subsystems	`IBMZ.Db2.Accounting_Summary`	• Source type: `zOS-SMF101_SUMMARY` • Threshold field: `ELAPSED_SEC_CL1` • Entity split by field: `sysplex + "." + system + "." + SSID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 GETPAGE Requests	Average GETPAGE requests for your Db2 subsystems	`IBMZ.Db2.Accounting_Summary`	• Source type: `zOS-SMF101_SUMMARY` • Threshold field: `BP4K_GETPAGE + BP32_GETPAGE + BP8K_GETPAGE + BP16_GETPAGE` • Entity split by field: `sysplex + "." + system + "." + SSID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 I/O Elapsed Wait Time	Average I/O elapsed wait time for your Db2 subsystems	`IBMZ.Db2.Accounting_Summary`	• Source type: `zOS-SMF101_SUMMARY` • Threshold field: `IO_WAIT_SEC` • Entity split by field: `sysplex + "." + system + "." + SSID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes

Table 5. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	Description	KPI base search	Search and calculate
Db2 I/O Rate	Average I/O rate for your Db2 jobs	`IBMZ.Db2.Job`	• Source type: `zOS-SMF30` • Threshold field: `IO_RATE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 Lock/Latch Wait Time	Average lock/latch wait time for your Db2 subsystems	`IBMZ.Db2.Accounting_Summary`	• Source type: `zOS-SMF101_SUMMARY` • Threshold field: `LOCK_LATCH_SEC` • Entity split by field: `sysplex + "." + system + "." + SSID` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 Lock Suspends	Average lock suspends per minute for your Db2 systems	Not applicable	• Source type: `zOS-SMF100_1` • Threshold field: Calculated in the search query • Entity split by field: `sysplex + "." + system` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 Lock Timeouts	Average lock timeouts per minute for your Db2 systems	Not applicable	• Source type: `zOS-SMF100_1` • Threshold field: Calculated in the search query • Entity split by field: `sysplex + "." + system` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 Paging Rate	Average paging rate for your Db2 jobs	`IBMZ.Db2.Job`	• Source type: `zOS-SMF30` • Threshold field: `PAGING_RATE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
Db2 Transaction Count	Sum transaction count for your Db2 subsystems	`IBMZ.Db2.Accounting_Summary`	• Source type: `zOS-SMF101_SUMMARY` • Threshold field: `COMMIT_COUNT` • Entity split by field: `sysplex + "." + system + "." + SSID` • Calculation: Calculating sum per entity, sum of aggregate over the last 15 minutes every 15 minutes

Table 5. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	Description	KPI base search	Search and calculate
Db2 Working Set Size	Average working set size for your Db2 jobs	`IBMZ.Db2.Job`	• Source type: `zOS-SMF30` • Threshold field: `WORKING_SET_SIZE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
IMS CPU Utilization	Average CPU utilization for your IMS jobs	`IBMZ.IMS.Job`	• Source type: `zOS-SMF30` • Threshold field: `CPU_UTLIZATION` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
IMS I/O Rate	Average I/O rate for your IMS jobs	`IBMZ.IMS.Job`	• Source type: `zOS-SMF30` • Threshold field: `IO_RATE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
IMS Paging Rate	Average paging rate for your IMS jobs	`IBMZ.IMS.Job`	• Source type: `zOS-SMF30` • Threshold field: `PAGING_RATE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
IMS Working Set Size	Average working set size for your IMS jobs	`IBMZ.IMS.Job`	• Source type: `zOS-SMF30` • Threshold field: `WORKING_SET_SIZE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
MQ CPU Utilization	Average CPU utilization for your MQ jobs	`IBMZ.MQ.Job`	• Source type: `zOS-SMF30` • Threshold field: `CPU_UTLIZATION` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes

Table 5. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	Description	KPI base search	Search and calculate
MQ I/O Rate	Average I/O rate for your MQ jobs	`IBMZ.MQ.Job`	• Source type: `zOS-SMF30` • Threshold field: `IO_RATE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
MQ Paging Rate	Average paging rate for your MQ jobs	`IBMZ.MQ.Job`	• Source type: `zOS-SMF30` • Threshold field: `PAGING_RATE` • Entity split by field: `sysplex + "." + system + "."+ JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes
MQ Working Set Size	Average working set size for your MQ jobs	`IBMZ.MQ.Job`	• Source type: `zOS-SMF30` • Threshold field: `WORKING_SET_SIZE` • Entity split by field: `sysplex + "." + system + "." + JOB_NAME` • Calculation: Calculating average per entity, average of aggregate over the last 15 minutes every 15 minutes

Table 6. KPI base searches in the Splunk ITSI module for IBM Z Operations Analytics

KPI base search	Description	Defined metrics
`IBMZ.CICS.Job`	Search that is used by KPIs that track accounting interval data for CICS jobs	• CPU Utilization • I/O Rate • Paging Rate • Working Set Size
`IBMZ.CICS.Transaction_Summary`	Search that is used by KPIs that track transaction summary data for CICS regions	• CPU Time • Dispatch Time • Elapsed Time • QR TCB CPU to Dispatch Time Ratio • QR TCB Dispatch Time • Response Time • Wait Time
`IBMZ.Db2.Accounting_Summary`	Search that is used by KPIs that track accounting summary data for Db2 subsystems	• CPU Time • Elapsed Time • GETPAGE Requests • I/O Elapsed Wait Time • Lock/Latch Wait Time • Transaction Count

Table 6. KPI base searches in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI base search	Description	Defined metrics
IBMZ.Db2.Job	Search that is used by KPIs that track accounting interval data for Db2 jobs	• CPU Utilization • I/O Rate • Paging Rate • Working Set Size
IBMZ.IMS.Job	Search that is used by KPIs that track accounting interval data for IMS jobs	• CPU Utilization • I/O Rate • Paging Rate • Working Set Size
IBMZ.MQ.Job	Search that is used by KPIs that track accounting interval data for MQ jobs	• CPU Utilization • I/O Rate • Paging Rate • Working Set Size

Other features in the Splunk ITSI module for IBM Z Operations Analytics

The Splunk IT Service Intelligence (ITSI) module for IBM Z Operations Analytics includes a custom Glass Table and Deep Dive for IBM Z systems, which can be integrated into your existing Z operations.

In the Splunk ITSI application, the Service Analyzer provides an overview of ITSI service health scores and KPI search results that are currently trending at the highest severity levels. You can view the IBM Z Operations Analytics service and KPIs in the default Service Analyzer and can quickly view the status of Z operations and identify services and KPIs running outside expected norms. By default, you can click on any tile in the Service Analyzer to drill down to the deep dives dashboard for further analysis and comparison of search results over time.

Glass Table for IBM Z systems
The IBM Z Operations Analytics custom glass table helps you visualize KPIs for monitoring your Z systems. You can click on any KPI in the glass table to navigate to the IBM Z Operations Analytics deep dive.

Deep Dive for IBM Z systems
The IBM Z Operations Analytics custom deep dive enables the display of KPI search results in a swim lane graphic, and lets you see the variations in your Z system metrics over time. You can click on any KPI metric in the deep dive to display a list of IBM Z Operations Analytics dashboards. Select a dashboard in the list to navigate to the IBM Z Operations Analytics Splunk application in a new browser window.

Installing the Splunk ITSI module for IBM Z Operations Analytics

To install the Splunk IT Service Intelligence (ITSI) module for IBM Z Operations Analytics, you must install the Splunk ITSI module `DA-ITSI-IBMZOA_3.2.01.spl` as a Splunk application.

Before you begin

Before you install the module, the following system requirements must be met:

- IBM Common Data Provider for z Systems V1.1.0 must be running on your z/OS system.

- On the Splunk Enterprise server or Splunk heavy forwarder, the following IBM Common Data Provider for z Systems components must be installed, configured, and running:
  - Data Receiver
  - Buffered Splunk Ingestion App
- The IBM Z Operations Analytics application must be installed.
- Splunk ITSI must be installed.

Procedure

1. Install the Splunk ITSI module `DA-ITSI-IBMZOA_3.2.01.spl` as a Splunk application. You can install from Splunk Web, from the Splunk command line, or by extracting the `DA-ITSI-IBMZOA_3.2.01.spl` file to the `$SPLUNK_HOME/etc/apps/` directory. In Splunk Web, messages are shown in the Message menu to indicate that `DA-ITSI-IBMZOA_3.2.01.spl` was created and that you must restart Splunk.
2. Restart Splunk.
3. Use one of the following methods to restore the IBM-provided backup JavaScript Object Notation (JSON) data for the Splunk ITSI module for IBM Z Operations Analytics:
   - From the Splunk ITSI UI, use the backup and restore function to create a restore job for reading the `backup_json_data_for_DA-ITSI-IBMZOA_3.2.01.zip` file. This `.zip` file must be on the system where the web browser is running rather than on the Splunk Enterprise server.
   - From the Splunk command line, run the `kvstore_to_json.py` script. If you use this method, you must unzip the `backup_json_data_for_DA-ITSI-IBMZOA_3.2.01.zip` file, and specify the file path to the JSON files.

   For more information, see Backup and restore ITSI data in the Splunk documentation.

What to do next

Troubleshooting tips:

Symptom	Problem
In Splunk ITSI, the following situations occur: • The IBM Z Operations Analytics glass table is missing. • The IBM Z Operations Analytics deep dive is missing. • The IBM Z Operations Analytics service is missing in **Configure** > **Services**. • The IBM Z Operations Analytics service is not displayed in the default Service Analyzer.	Unsuccessful restoration of the IBM-provided backup JSON data for the Splunk ITSI module for IBM Z Operations Analytics
In Splunk ITSI, the following situations occur: • The ITSI Module for IBM Z Operations Analytics is missing in **Configure** > **Modules**. • The custom drilldowns from the IBM Z Operations Analytics deep dive are missing.	Unsuccessful installation of the Splunk ITSI module for IBM Z Operations Analytics (`DA-ITSI-IBMZOA_3.2.01.spl`)

Symptom	Problem
In Splunk Web, the following message occurs when you navigate from the IBM Z Operations Analytics deep dive to a dashboard in IBM Z Operations Analytics: `The app "ibm_zoa_insights" is not available`	Unsuccessful installation of the IBM Z Operations Analytics application (`ibm_zoa_insights.spl`)

Problem Insights Framework

IBM Z Operations Analytics includes an extensible Problem Insights Framework that provides insights for a defined set of potential problems in your IT environment. In the Elastic Stack or Splunk UI, these problem insights are organized by sysplex so that you can easily view the insights for a specific sysplex within a specified time period. The UI also provides multiple ways of sorting and filtering problem insights.

Problem insights

Each problem insight includes the following information:

Severity
The severity of the problem.

Sysplex
The sysplex where the problem occurred.

System
The system where the problem occurred.

Subsystem
The subsystem or system resources manager where the problem occurred.

Problem Summary
A summary of the problem.

Count The number of occurrences of the problem in the specified time period.

Time The time when the problem last occurred in the specified time period.

Suggested Actions
The suggested actions that you can take to resolve the problem.

Evidence
Information that further identifies the problem.

Message Libraries

The problem insights content in the Elastic Stack or Splunk UI is provided by Message Libraries in the Problem Insights Framework. A Message Library is an XML file that applies to a specific domain of interest. A domain of interest might be, for example, the network, the z/OS system, or a z/OS subsystem such as CICS Transaction Server for z/OS or MQ for z/OS. The XML file contains messages about a defined set of potential problems for the applicable domain of interest. The Problem Insights Framework uses the Message Libraries to find information about potential problems and provide information about how to fix those problems.

You can customize the messages in the Message Libraries to better represent problems that can occur in your environment.

For more information about the Message Libraries, see "Message Library reference for the Problem Insights Framework" on page 102.

System requirements for the Problem Insights Framework

Your Elastic Stack or Splunk environment must meet the system requirements for deploying the IBM Z Operations Analytics Problem Insights Framework.

Operating system requirements

Table 7. Operating system requirements for the Problem Insights Framework

Operating system	Version
Windows Server	2008, 2012, or 2016
Red Hat Enterprise Linux: for IBM POWER8®	6 or 7
Red Hat Enterprise Linux: for IBM Z Systems	6 or 7
Red Hat Enterprise Linux: for x86	6 or 7
SUSE Linux Enterprise Server: for IBM POWER8	11 or 12
SUSE Linux Enterprise Server: for IBM Z Systems	11 or 12
SUSE Linux Enterprise Server: for x86	11 or 12

Runtime requirements

- Java™ Runtime Environment (JRE) 8 or later
- For a Linux or UNIX system, the Bash shell

Port requirement

HTTPS port 9446

Data storage requirements on your platform

The Problem Insights Framework pulls data from either the Elastic Stack or Splunk platform. The results of data matches (between incoming operational data and problem insights content) are stored in different ways, depending on your platform.

Elastic Stack platform: data storage

On the Elastic Stack platform, the data match results are stored in the Elasticsearch database.

Splunk platform: data storage

On the Splunk platform, the data match results are stored locally. Therefore, the system where the Problem Insights Framework is installed must have enough storage space to hold data for an extended period of time.

Deployment planning for the Problem Insights Framework

Guidelines are provided for planning your deployment of the IBM Z Operations Analytics Problem Insights Framework, based on your platform (Elastic Stack or Splunk) and your environment.

Use the following guidelines for deployment, depending on your platform:

Elastic Stack platform

Deploy the Problem Insights Framework either on the same server as the

Elastic Stack Kibana component, or on a server that is located near that Kibana server. The Problem Insights Framework stores the results of data matches (between incoming operational data and problem insights content) from Elastic Stack in the Elasticsearch database. The data is later retrieved from the Elasticsearch database and presented in the Problem Insights dashboard of the IBM Z Operations Analytics application, based on the specified time period. The request goes directly to the Elasticsearch database.

Figure 4 on page 83 illustrates this data flow from each Message Library XML file, through the Problem Insights Framework, to the presentation of the problem insights content in the dashboard.

The default time period is 15 minutes. The Elasticsearch administrator must prune the data as needed.

Splunk platform

Deploy the Problem Insights Framework either on the Splunk search head component, or as close to the search head component as possible. The Problem Insights Framework stores the results of data matches (between incoming operational data and problem insights content) from Splunk locally. The data is later retrieved from the Problem Insights Framework and presented in the Problem Insights dashboard of the IBM Z Operations Analytics application, based on the specified time period. The request goes through the Splunk custom endpoint to the Problem Insights Framework by using basic authentication over HTTPS.

Figure 4 on page 83 illustrates this data flow from each Message Library XML file, through the Problem Insights Framework, to the presentation of the problem insights content in the dashboard.

The default time period is 15 minutes. A best practice is to not request data for a time period greater than 30 days. The Command reference for the Problem Insights Framework includes a **`purgeResults`** command that you can use to delete data match results that are over 30 days old.

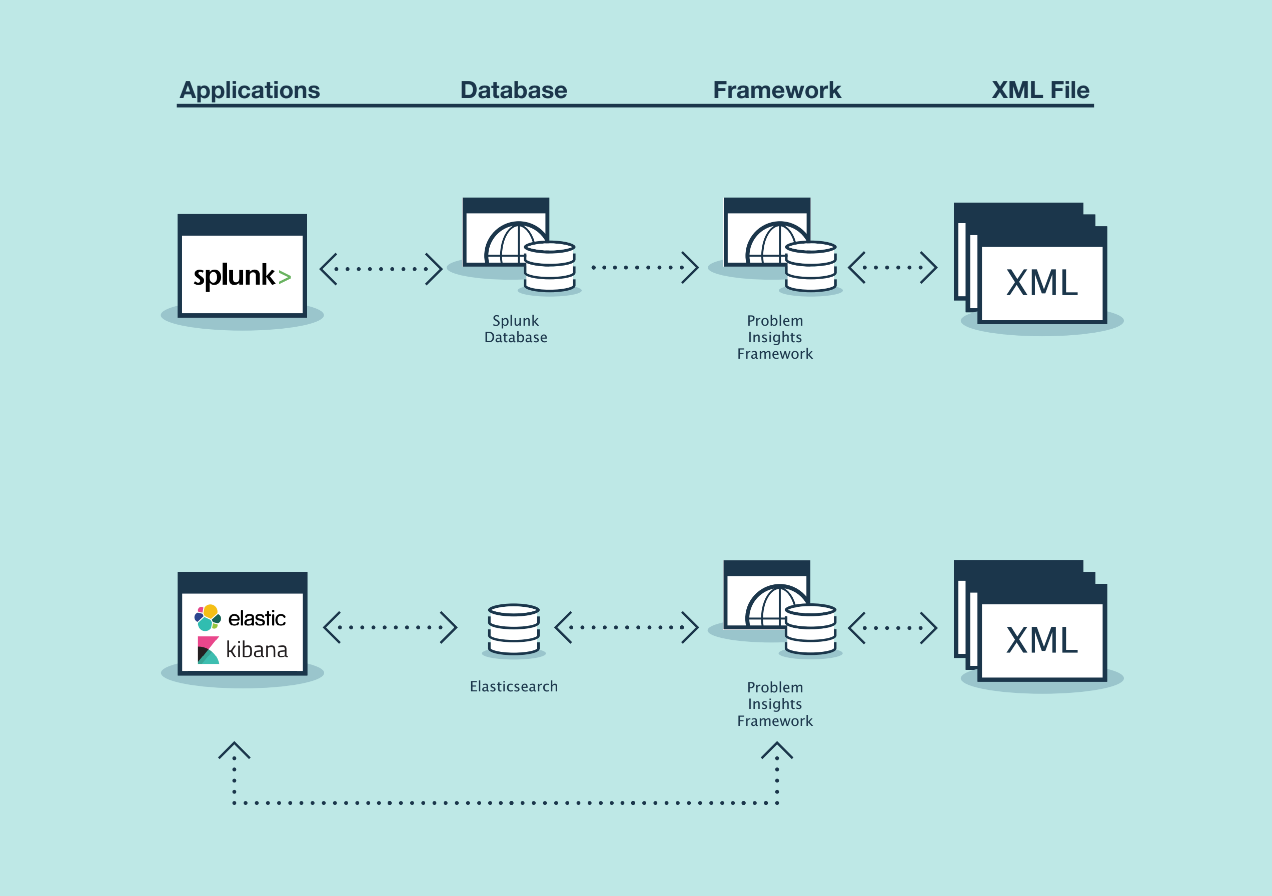

Figure 4. Data flow in the Problem Insights architecture

Installing the Problem Insights Framework

To install the IBM Z Operations Analytics Problem Insights Framework, you must extract the `pi_framework.zip` file from the installation media.

Before you begin

Review the following topics, and verify that all requirements are met:

- "System requirements for the Problem Insights Framework" on page 81
- "Deployment planning for the Problem Insights Framework" on page 81

Procedure

Extract the `pi_framework.zip` file from the installation media.
The following directories and files are extracted to your destination directory:

Directory or file	**Description**
`.appdata/*`	Workspace
• `bin/analysis.bat` • `bin/analysis.sh`	Scripts for running CLI commands
`config/*`	Configuration files

Directory or file	Description
`lib/*.jar`	Java archive (JAR) files
`license/*`	Licenses
`pi_framework`	Home directory for Problem Insights Framework
`usr/lang/*.xml`	IBM-provided Message Libraries
`wlp/*`	WebSphere Liberty Profile

Tip: If you ever need to uninstall the Problem Insights Framework, complete the following steps:

1. Stop the Problem Insights Framework, as described in "Command reference for the Problem Insights Framework" on page 97.
2. Delete the directory where the `pi_framework.zip` file was extracted.

Configuring the Problem Insights Framework

IBM Z Operations Analytics Problem Insights Framework must have access to the Elastic Stack or Splunk platform so that it can analyze operational data. Configure access to your platform by updating the configuration properties in the `config/PiSearch.config` file. You must also complete some additional configuration to make the Problem Insights Framework operational in your environment.

Procedure

1. To give the Problem Insights Framework access to your platform, update the configuration properties in the `config/PiSearch.config` file, as described in "Configuration property reference for the Problem Insights Framework" on page 98.
2. Complete the following additional configuration to make the Problem Insights Framework operational in your environment.

Platform	Additional configuration
Elastic Stack platform	Deploy the IBM Z Operations Analytics application, as described in "Deploying the Z Operations Analytics application on the Elastic Stack platform" on page 62. In step 8a on page 64 of "Deploying the Z Operations Analytics application on the Elastic Stack platform" on page 62, you must provide the value for the **`pi_ip`** parameter, which specifies the IP address where the Problem Insights Framework is bound.

Platform	Additional configuration
Splunk platform	On the Splunk platform, the Problem Insights Framework stores data match results locally. When a request is made from the Problem Insights dashboard in the Splunk UI, the UI pulls data from the Problem Insights Framework by using Splunk custom endpoints. The communication between the Splunk UI and the Problem Insights Framework is defined in the `$SPLUNK_HOME/etc/apps/ibm_zoa_insights/default/probleminsights.conf` file of the IBM Z Operations Analytics application. In this file, the value of the **`PI_BASE_URL`** parameter is `https://localhost:9446`. This value is valid only if the Problem Insights Framework is installed on the system where the Splunk custom endpoint is defined. If the Problem Insights Framework is installed on a remote host, complete the following steps: 1. In the `$SPLUNK_HOME/etc/apps/ibm_zoa_insights` directory, create a new directory named `local`. 2. Copy the `probleminsights.conf` file from the `default` directory to the `local` directory. 3. In the `$SPLUNK_HOME/etc/apps/ibm_zoa_insights/local/probleminsights.conf` file, update the value of the **`PI_BASE_URL`** parameter to replace `localhost` with the IP address or the fully qualified domain name of the system where the Problem Insights Framework is installed. 4. Restart the Splunk Enterprise server.

What to do next

IBM Z Operations Analytics Problem Insights Framework includes commands for operating the Framework, such as for starting and stopping the Framework and for importing and exporting Message Libraries. See the Command reference for the Problem Insights Framework for information about all the commands.

Configuring authentication for the Problem Insights Framework

For authentication between the IBM Z Operations Analytics Problem Insights Framework and your analytics platform, you can use the basic authentication that is provided, or you can configure mutual authentication. For example, you might want to use mutual authentication if the server where the Problem Insights Framework is installed is remote from the server for your analytics platform.

Configuring basic authentication for the Problem Insights Framework

IBM Z Operations Analytics Problem Insights Framework provides basic authentication between the Problem Insights Framework and your analytics platform by using HTTPS. If you want to change any information that is related to basic authentication, such as the host server, port, user name, or password, you must update all components that use that information.

Procedure

To configure basic authentication for the Problem Insights Framework, complete the following steps:

1. Stop the Problem Insights Framework, as described in "Command reference for the Problem Insights Framework" on page 97.
2. Update the following Problem Insights Framework files to reflect the change of host server, port, user name, or password:
   - `pi_framework/config/cli.properties`
   - `pi_framework/wlp/usr/servers/piFrameworkServer/server.xml`
3. Complete the following steps, depending on your platform:

Platform	Steps
Elastic Stack	From the **Management** tab in Kibana, click **Index Patterns**, select **pi-***, and update the URL of the **MESSAGEID** column.
Splunk	1. Copy the IBM Z Operations Analytics application file `$SPLUNK_HOME/etc/apps/ibm_zoa_insights/default/probleminsights.conf` into the `$SPLUNK_HOME/etc/apps/ibm_zoa_insights/local` directory. 2. Update the `$SPLUNK_HOME/etc/apps/ibm_zoa_insights/local/probleminsights.conf` file to reflect the change of host server, port, user name, or password.

Configuring mutual authentication for the Problem Insights Framework

You can configure mutual authentication between the IBM Z Operations Analytics Problem Insights Framework and your analytics platform.

About this task

The first step in configuring mutual authentication is to obtain certificates that are signed by a certificate authority (CA). If you do not have access to a CA-signed certificate, you might want to generate self-signed certificates. The steps in the following procedure are based on the use of self-signed certificates. An example script is provided to illustrate how to generate these certificates. For more information, see "Script for generating self-signed certificates" on page 88.

Also, for the Splunk platform, a sample scenario is provided to illustrate how to configure mutual authentication between the IBM Z Operations Analytics Problem Insights Framework and the Splunk server when the Problem Insights Framework is installed on a server that is remote from the Splunk server. To review this scenario, see "Scenario that illustrates how to configure mutual authentication" on page 92.

Procedure

To configure mutual authentication, complete the following steps:

1. Run the script to generate your certificates, as described in "Script for generating self-signed certificates" on page 88.
2. In the `pi_framework/wlp/usr/servers/piFrameworkServer/server.xml` file, uncomment the following code, and provide your values for the variables that are shown in italics.

```
<ssl id="mysslsettings" keyStoreRef="mykeystore" trustStoreRef="myTrustStore"
   clientAuthenticationSupported="true"/>
<keyStore id="mykeystore" location="SERVER KEYSTORE LOCATION"
   type="KEYSTORE TYPE" password="KEYSTORE PASSWORD"/>
<keyStore id="myTrustStore" location="SERVER TRUSTSTORE LOCATION"
   type="TRUSTSTORE TYPE" password="TRUSTSTORE PASSWORD"/>
<sslDefault sslRef="mysslsettings"/>
```

3. For the **`basicRegistry`** element, uncomment the 3rd line of the following code, and provide your values for the variables that are shown in italics.

```
<basicRegistry id="basic">
   <user name="izoa2018user" password="iZoa-2018pw"/>
   <user name="CN OF CLIENT CERT" password="PASSWORD"/>
</basicRegistry>
```

4. For the **`security-role`** element, uncomment the 3rd line of the following code, and provide your values for the variables that are shown in italics.

```
<security-role name="restricted">
   <user name="izoa2018user"/>
   <user name="CN OF CLIENT CERT"/>
</security-role>
```

5. Add the following parameters, with your parameter values, to the `pi_framework/config/cli.properties` file.

   Tips:

   - On Windows systems, you must use double backslashes in the path name, as shown in the following example:

     ```
     C:\\Users\\User1\\Documents
     ```

   - Java keystores and truststores are the only certificate types that are supported by the Problem Insights Framework.

```
#############################################
# Keystore Location
#############################################
keystoreLocation=CLIENT KEYSTORE LOCATION

#############################################
# Keystore Type (for example, JKS)
#############################################
keystoreType=CLIENT KEYSTORE TYPE

#############################################
# Keystore password
#############################################
# keystorePass=CLIENT KEYSTORE PASSWORD

#############################################
# Truststore Location
#############################################
truststoreLocation=CLIENT TRUSTSTORE LOCATION

#############################################
# Truststore Type (for example, JKS)
#############################################
truststoreType=CLIENT TRUSTSTORE TYPE

#############################################
# Truststore password
#############################################
truststorePass=CLIENT TRUSTSTORE PASSWORD

#############################################
# Certificate Alias
#############################################
keyAlias=client-cert
```

6. Depending on your platform, complete the following steps.

   Elastic Stack platform

   To import your client certificate, and to accept the Problem Insights Framework server certificate, complete the following steps. These steps are specific to a version of the Mozilla Firefox web browser, but you can apply similar steps in the web browser of your choice.

   a. From the main menu, click **Options**.

b. Click **Privacy and Security**, and in the resulting panel, scroll to the "Security" section, which has a "Certificates" subsection.

c. In the "Certificates" subsection, click **View Certificates**.

d. In the resulting Certificate Manager window, click the **Your Certificates** tab.

e. On the Your Certificates page, click **Import** to import the `clientkeystore.p12` file that was previously generated. You must also provide the keystore password.

f. If you are *not* using a certificate that is distributed by a trusted certificate authority (CA), also complete the following steps:

   1) In the same Certificate Manager window, click the **Servers** tab.

   2) On the Servers page, click **Add Exception**.

   3) Type the URL for a page on the Problem Insights Framework server. In the URL, use the fully qualified name of the server where the Problem Insights Framework is installed.

   4) Click **Get Certificate**.

Splunk platform

To update your server certificate, client certificate, and client key, complete the following steps.

a. In the `$SPLUNK_HOME/etc/apps/ibm_zoa_insights` directory, create the directory `local`.

b. Copy the `$SPLUNK_HOME/etc/apps/ibm_zoa_insights/default/probleminsights.conf` file into the `$SPLUNK_HOME/etc/apps/ibm_zoa_insights/local` directory.

c. In the `$SPLUNK_HOME/etc/apps/ibm_zoa_insights/local/probleminsights.conf` file, add the following parameters, with your parameter values:

```
SERVER_CERT=LOCATION OF CERTIFICATE FOR PROBLEM INSIGHTS FRAMEWORK SERVER
   (must have a file type of .pem)
CLIENT_KEY=CLIENT KEY LOCATION
   (must have a file type of .key)
CLIENT_CERT=LOCATION OF CLIENT CERTIFICATE
   (must have a file type of .pem)
```

Tip: Splunk supports only `.pem` as a certificate format, and only `.key` as a private key format.

d. Verify that the appropriate permissions are set for the `local` directory so that the directory is secure, and authorized users can make updates.

Script for generating self-signed certificates:

A script for Linux systems is provided "as is," without warranty of any kind, to illustrate how to generate the self-signed certificates that are required to configure mutual authentication. Although the use of self-signed certificates is not a best practice, this script can help you understand the minimum configuration that is required for mutual authentication.

"Script code" on page 89 provides the code for the script. The lines of code that are highlighted in bold text correlate to commands that require your input. For more information about the input that you must provide, see "Input that you must provide to the script" on page 90.

The commands in the script must be run in order because the files that each command generates might be needed as input to later commands.

On Windows systems: You can run the individual commands to generate the file that is needed. However, do not run the lines that start with echo or #. Also, verify that the following prerequisites are available in your Windows system:

- A version of OpenSSL
- The Java keytool utility

Script code

```
#!/bin/bash
#################################################################
# Disclaimer of Warranties:                                     #
# The following code is sample code created by IBM Corporation  #
# and is provided to you solely for instructional purposes for  #
# assisting you in the setup of certificates for secure         #
# connection between multiple products and third party          #
# components. The code is provided "AS IS" without warranty     #
# of any kind. IBM shall not be liable for any damages arising  #
# out of your use of this sample code, even if advised of the   #
# possibility of such damages. This sample code does not        #
# contain entitlements to an IBM program; entitlements to an    #
# IBM program are provided to you under a separate written      #
# agreement between you and IBM.                                #
#################################################################
echo Generate a private RSA key - SERVER
openssl genrsa -out myPIservercert.key 2048
echo Create a x509 certificate - SERVER
echo common name must match hostname from server.xml
echo all other inputs are optional
openssl req -x509 -new -nodes -key myPIservercert.key \
     -sha256 -days 1024 -out myPIservercert.pem
echo Create a PKCS12 keystore from private key and public certificate. - SERVER
openssl pkcs12 -export -name server-cert \
     -in myPIservercert.pem -inkey myPIservercert.key \
     -out serverkeystore.p12
echo Convert PKCS12 keystore into a JKS keystore - SERVER
keytool -importkeystore -destkeystore server.keystore \
     -srckeystore serverkeystore.p12 -srcstoretype pkcs12 \
     -alias server-cert
echo Generate a private key - CLIENT
openssl genrsa -out myPIclientcert.key 2048
echo Create a x509 certificate - CLIENT
echo common name must match username from server.xml
echo all other inputs are optional
openssl req -x509 -new -nodes -key myPIclientcert.key \
     -sha256 -days 1024 -out myPIclientcert.pem
echo Create PKCS12 keystore from private key and public certificate. - CLIENT
openssl pkcs12 -export -name client-cert \
     -in myPIclientcert.pem -inkey myPIclientcert.key \
     -out clientkeystore.p12
echo Convert a PKCS12 keystore into a JKS keystore - CLIENT
keytool -importkeystore -destkeystore client.keystore \
     -srckeystore clientkeystore.p12 -srcstoretype pkcs12 \
     -alias client-cert
echo IMPORTING
echo Import a client certificate to the server trust store.
keytool -import -alias client-cert \
     -file myPIclientcert.pem -keystore server.truststore
echo Import a server certificate to the server trust store.
keytool -import -alias server-cert \
     -file myPIservercert.pem -keystore server.truststore
echo Import a server certificate to the client trust store.
keytool -import -alias server-cert -file myPIservercert.pem \
     -keystore client.truststore
echo Import a client certificate to the client trust store.
keytool -import -alias client-cert -file myPIclientcert.pem \
     -keystore client.truststore
```

Input that you must provide to the script

Table 8 correlates line numbers in the script to descriptions of the input that you must provide to the script prompts.

Tips for using the script:

- Keep a record of the passwords that you use and the files for which you use them.
- If the script is run with the code that is shown, the alias for the client certificate is `client-cert`, which is specified in line 44.

Table 8. Correlation of line numbers in the script to descriptions of the input that you must provide to the script prompts

Line number	Input that you must provide

21	The command in line 21 prompts you for a Common Name (CN). The CN for the server must be the name of the server where the Problem Insights Framework is installed. If the server is local, the value must be `localhost`. Otherwise, specify the fully qualified domain name of the server.
	This CN is required for the `pi_framework/wlp/usr/servers/piFrameworkServer/server.xml` configuration file. For example, the following code shows an entry in the `server.xml` file, where the CN is `localhost`:
	`<httpEndpoint`
	`   httpsPort="9446"`
	`   host="`**`localhost`**`"`
	`   id="defaultHttpEndpoint"/>`
	The command output is used as input to later commands.

23	The command in line 23 prompts you for the PKCS12 keystore password on the server where the Problem Insights Framework is installed.
	This password is required for the `pi_framework/wlp/usr/servers/piFrameworkServer/server.xml` configuration file. For example, the following code shows an entry in the `server.xml` file, where the password is `pw_4_serverkeystore.p12`:
	`<keyStore id="mykeystore" location="`*`path`*`/serverkeystore.p12"`
	`   type="PKCS12" password="`**`pw_4_serverkeystore.p12`**`" />`
	If you want to use the Java Keystore (JKS) instead of the PKCS12 keystore, see the information in this table for line 25 of the script.
	The command output is used as input to later commands.

25	The command in line 25 prompts you for the JKS password on the server where the Problem Insights Framework is installed.
	If you want to use the JKS, you must update the **`keyStore`** element as shown in the following example. In this example, the value of **`password`** is `pw_4_server.keystore`.
	`<keyStore id="mykeystore" location="`*`path`*`/server.keystore"`
	`   type="jks" password="`**`pw_4_server.keystore`**`"/>`
	The command output is used as input to later commands.

31	The command in line 31 prompts you for the CN of the server where the Problem Insights Framework is installed. If the server is local, the value must be `localhost`. Otherwise, specify the fully qualified domain name of the server.

Table 8. Correlation of line numbers in the script to descriptions of the input that you must provide to the script prompts (continued)

Line number	Input that you must provide

33	The command in line 33 prompts you for the PKCS12 keystore password on the server where the Problem Insights Framework is installed.
	This password is required for the `pi_framework/config/cli.properties` configuration file. For example, the following code shows an entry in the `cli.properties` file, where the password is `pw_4_clientkeystore_p12`:
	`#############################################`
	`# Keystore Location.`
	`# On windows all paths must have double slashes`
	`#############################################`
	keystoreLocation=*path*/clientkeystore.p12
	
	`#############################################`
	`# Keystore Type`
	`#############################################`
	`keystoreType=PKCS12`
	
	`#############################################`
	`# Keystore password`
	`#############################################`
	keystorePass=**pw_4_clientkeystore_p12**
	The command output is used as input to later commands.

35	The command in line 35 prompts you for the JKS password on the server where the Problem Insights Framework is installed.
38	The command in line 38 prompts you for the `server.truststore` keystore password on the server where the Problem Insights Framework is installed. This password is required for the `pi_framework/wlp/usr/servers/piFrameworkServer/server.xml` configuration file. For example, the following code shows an entry in the `server.xml` file, where the password is `pw_4server.truststore`: <keyStore id="myTrustStore" location="*path*/server.truststore" type="jks" password="**pw_4server.truststore**"/> The command output is used as input to later commands.

Table 8. Correlation of line numbers in the script to descriptions of the input that you must provide to the script prompts (continued)

<table>
<tr><th>Line number</th><th>Input that you must provide</th></tr>
<tr><td>42</td><td>The command in line 42 prompts you for the <code>client.truststore</code> keystore password on the server where the Problem Insights Framework is installed.

This password is required for the <code>pi_framework/config/cli.properties</code> configuration file. For example, the following code shows an entry in the <code>cli.properties</code> file, where the password is <code>pw_4_client_truststore</code>:
<pre>
###
# Truststore Location
# On windows all paths must have double slashes
###
truststoreLocation=path/client.truststore

###
# Truststore Type
###
truststoreType=JKS

###
# Truststore password
###
truststorePass=pw_4_client_truststore
</pre>
The command output is used as input to later commands.</td></tr>
<tr><td>44</td><td>The command in line 44 passes the parameter value <code>client-cert</code>, which must also be in the <code>pi_framework/config/cli.properties</code> configuration file. For example, the following code shows an entry in the <code>cli.properties</code> file where this value is used:
<pre>
###
# Certificate Alias
###
keyAlias=client-cert
</pre></td></tr>
</table>

Scenario that illustrates how to configure mutual authentication:

This scenario illustrates how to configure mutual authentication between the IBM Z Operations Analytics Problem Insights Framework and the Splunk analytics platform when the Problem Insights Framework is installed on a server that is *remote from* the Splunk server.

"Commands" on page 93 shows the commands that are required to generate the self-signed certificates.

"Configuration example" on page 95 shows the configuration that must be done on the server where the Problem Insights Framework is installed and on the Splunk server.

Important notes about this scenario:

- The passwords that are used are only examples and must not be used in a production environment.
- The server where the Problem Insights Framework is installed has the host name `pi-server.mycompany.com`.
- The Splunk server has the host name `splunk-server.mycompany.com`.

Commands

The following commands show the generation of two client certificates (one for the CLI of the Problem Insights Framework, and one for the Splunk server) and one server certificate (for the server where the Problem Insights Framework is installed). After you run these commands, you must move the generated files to the server where the Problem Insights Framework is installed and to the Splunk server, as described in “Configuration example” on page 95.

Where the text is highlighted in bold, you must provide your own values (for example, for passwords and host names).

```
#  openssl genrsa -out myPIservercert.key 2048
Generating RSA private key, 2048 bit long modulus
.................................................+++
.............................................+++
e is 65537 (0x10001)
#  openssl req -x509 -new -nodes -key myPIservercert.key -sha256 -days 1024 -out myPIservercert.pem
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) []:US
State or Province Name (full name) []:NC
Locality Name (eg, city) []:
Organization Name (eg, company) []:IBM
Organizational Unit Name (eg, section) []:IZOA
Common Name (eg, fully qualified host name) []:pi-server.mycompany.com
Email Address []:
#  openssl pkcs12 -export -name server-cert -in myPIservercert.pem
      -inkey myPIservercert.key -out serverkeystore.p12
Enter Export Password:  serverkeystore_p12_pw
Verifying - Enter Export Password:   serverkeystore_p12_pw
#  keytool -importkeystore -destkeystore server.keystore -srckeystore
      serverkeystore.p12 -srcstoretype pkcs12 -alias server-cert
Importing keystore serverkeystore.p12 to server.keystore...
Enter destination keystore password:  server_keystore_pw
Re-enter new password:   server_keystore_pw
Enter source keystore password:  serverkeystore_p12_pw

Warning:
The JKS keystore uses a proprietary format. It is recommended to migrate to PKCS12
which is an industry standard format using
"keytool -importkeystore -srckeystore server.keystore -destkeystore server.keystore
   -deststoretype pkcs12".

#  openssl genrsa -out myPIclientcert.key 2048
Generating RSA private key, 2048 bit long modulus
................................................................................+++
...........................................+++
e is 65537 (0x10001)
#  openssl req -x509 -new -nodes -key myPIclientcert.key -sha256 -days 1024 -out myPIclientcert.pem
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) []:US
State or Province Name (full name) []:NC
Locality Name (eg, city) []:
Organization Name (eg, company) []:IBM
Organizational Unit Name (eg, section) []:IZOA
Common Name (eg, fully qualified host name) []:pi-server.mycompany.com
Email Address []:
#  openssl pkcs12 -export -name client-cert -in myPIclientcert.pem
      -inkey myPIclientcert.key -out clientkeystore.p12
Enter Export Password:  clientkeystore_p12_pw
Verifying - Enter Export Password:  clientkeystore_p12_pw
#  keytool -importkeystore -destkeystore client.keystore
      -srckeystore clientkeystore.p12 -srcstoretype pkcs12 -alias client-cert
Importing keystore clientkeystore.p12 to client.keystore...
```

```
Enter destination keystore password:   client_keystore_pw
Re-enter new password:   client_keystore_pw
Enter source keystore password:  clientkeystore_p12_pw

Warning:
The JKS keystore uses a proprietary format. It is recommended to migrate to PKCS12
which is an industry standard format using
"keytool -importkeystore -srckeystore client.keystore -destkeystore client.keystore
    -deststoretype pkcs12".

#  keytool -import -alias client-cert -file myPIclientcert.pem
      -keystore server.truststore
Enter keystore password:  server_truststore_pw
Re-enter new password:   server_truststore_pw
Owner: CN=pi-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Issuer: CN=pi-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Serial number: cf3255c1d4d5f29c
Valid from: Tue Nov 27 14:27:42 EST 2018 until: Thu Sep 16 15:27:42 EDT 2021
Certificate fingerprints:
         MD5:  89:1A:83:3B:A0:C8:9C:91:79:BC:69:31:31:C3:B6:8D
         SHA1: 58:94:8E:D9:35:87:3E:A4:1B:B1:0C:48:E8:30:97:28:BE:A6:38:12
         SHA256: D5:6C:6D:05:60:09:31:BE:11:BF:1B:66:DE:05:00:C4:FB:30:62:
            9F:10:41:BB:80:7A:05:82:C7:3A:1D:D1:C7
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: 2048-bit RSA key
Version: 1
Trust this certificate? [no]:  yes
Certificate was added to keystore
#  keytool -import -alias server-cert -file myPIservercert.pem
      -keystore server.truststore
Enter keystore password:  server_truststore_pw
Owner: CN=pi-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Issuer: CN=pi-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Serial number: c0abe229ab81d80a
Valid from: Tue Nov 27 14:23:56 EST 2018 until: Thu Sep 16 15:23:56 EDT 2021
Certificate fingerprints:
         MD5:  2F:5E:F2:4C:CB:05:5D:DF:2E:8C:3D:A7:E7:83:F1:64
         SHA1: 6C:BE:FC:6E:62:40:CF:A9:B2:02:F6:AC:6D:60:00:E7:C0:65:9E:B7
         SHA256: D6:3D:EE:E8:8E:AC:FB:DF:F7:1A:7A:B2:C9:A4:54:52:8C:C0:4E:A8:
            8A:0E:81:39:87:80:1E:E9:AB:A1:29:2D
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: 2048-bit RSA key
Version: 1
Trust this certificate? [no]:  yes
Certificate was added to keystore
#  keytool -import -alias server-cert -file myPIservercert.pem
      -keystore client.truststore
Enter keystore password:  client_truststore_pw
Re-enter new password:  client_truststore_pw
Owner: CN=pi-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Issuer: CN=pi-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Serial number: c0abe229ab81d80a
Valid from: Tue Nov 27 14:23:56 EST 2018 until: Thu Sep 16 15:23:56 EDT 2021
Certificate fingerprints:
         MD5:  2F:5E:F2:4C:CB:05:5D:DF:2E:8C:3D:A7:E7:83:F1:64
         SHA1: 6C:BE:FC:6E:62:40:CF:A9:B2:02:F6:AC:6D:60:00:E7:C0:65:9E:B7
         SHA256: D6:3D:EE:E8:8E:AC:FB:DF:F7:1A:7A:B2:C9:A4:54:52:8C:C0:4E:A8:8A:
            0E:81:39:87:80:1E:E9:AB:A1:29:2D
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: 2048-bit RSA key
Version: 1
Trust this certificate? [no]:  yes
Certificate was added to keystore
#  keytool -import -alias client-cert -file myPIclientcert.pem
      -keystore client.truststore
Enter keystore password:  client_truststore_pw
Owner: CN=pi-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Issuer: CN=pi-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Serial number: cf3255c1d4d5f29c
Valid from: Tue Nov 27 14:27:42 EST 2018 until: Thu Sep 16 15:27:42 EDT 2021
Certificate fingerprints:
         MD5:  89:1A:83:3B:A0:C8:9C:91:79:BC:69:31:31:C3:B6:8D
         SHA1: 58:94:8E:D9:35:87:3E:A4:1B:B1:0C:48:E8:30:97:28:BE:A6:38:12
         SHA256: D5:6C:6D:05:60:09:31:BE:11:BF:1B:66:DE:05:00:C4:FB:30:62:9F:10:
            41:BB:80:7A:05:82:C7:3A:1D:D1:C7
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: 2048-bit RSA key
Version: 1
Trust this certificate? [no]:  yes
Certificate was added to keystore
```

```
#  openssl genrsa -out mySplunkclientcert.key 2048
Generating RSA private key, 2048 bit long modulus
.....................................+++
...............................................+++
e is 65537 (0x10001)
#  openssl req -x509 -new -nodes -key mySplunkclientcert.key
      -sha256 -days 1024 -out mySplunkclientcert.pem
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) []:US
State or Province Name (full name) []:NC
Locality Name (eg, city) []:
Organization Name (eg, company) []:IBM
Organizational Unit Name (eg, section) []:IZOA
Common Name (eg, fully qualified host name) []:splunk-server.mycompany.com
Email Address []:
#  openssl pkcs12 -export -name splunkclient-cert -in mySplunkclientcert.pem
      -inkey mySplunkclientcert.key -out splunkclientkeystore.p12
Enter Export Password:   splunkclientkeystore_p12_pw
Verifying - Enter Export Password:  splunkclientkeystore_p12_pw
#  keytool -importkeystore -destkeystore splunkclient.keystore
      -srckeystore splunkclientkeystore.p12
#    -srcstoretype pkcs12 -alias splunkclient-cert
Importing keystore splunkclientkeystore.p12 to splunkclient.keystore...
Enter destination keystore password:  splunkclient_keystore_pw
Re-enter new password:  splunkclient_keystore_pw
Enter source keystore password:  splunkclientkeystore_p12_pw

Warning:
The JKS keystore uses a proprietary format. It is recommended to
migrate to PKCS12 which is an industry standard format using
"keytool -importkeystore -srckeystore splunkclient.keystore
-destkeystore splunkclient.keystore -deststoretype pkcs12".
#  keytool -import -alias splunkclient-cert -file mySplunkclientcert.pem
      -keystore server.truststore
Enter keystore password:   server_truststore_pw
Owner: CN=splunk-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Issuer: CN=splunk-server.mycompany.com, OU=IZOA, O=IBM, ST=NC, C=US
Serial number: ed6ee3c8c700f150
Valid from: Tue Nov 27 16:51:34 EST 2018 until: Thu Sep 16 17:51:34 EDT 2021
Certificate fingerprints:
         MD5:  8F:94:DC:B4:D9:48:D3:83:6E:C6:5F:F4:41:76:17:2E
         SHA1: BF:F6:D8:12:0E:F8:CD:A8:FD:4A:DD:22:EA:28:E1:03:0B:E5:83:D2
         SHA256: E5:93:02:3F:15:B6:B5:54:C6:2D:7B:08:CE:4D:90:1F:37:74:D8:AA:BF:
            94:81:02:B0:14:9A:E2:7F:1C:FC:3A
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: 2048-bit RSA key
Version: 1
Trust this certificate? [no]:  yes
Certificate was added to keystore
#
```

Configuration example

Code in file pi_framework/wlp/usr/servers/piFrameworkServer/server.xml

In the following example, the server keystore and truststore (for the server where the Problem Insights Framework is installed) that were generated by the commands are moved to the pi_framework/config directory:

```
<httpEndpoint
   httpsPort="9446"
   host="pi-server.mycompany.com"
   id="defaultHttpEndpoint"/>
.
.
.
<ssl id="mysslsettings" keyStoreRef="mykeystore" trustStoreRef="myTrustStore"
   clientAuthenticationSupported="true"/>
<keyStore id="mykeystore" location="/pi_framework/config/server.keystore"
   type="jks" password="server_keystore_pw">
      <keyEntry keyPassword="serverkeystore_p12_pw" name="server-cert"/>
</keyStore>
<keyStore id="myTrustStore" location="/pi_framework/config/server.truststore"
   type="jks" password="server_truststore_pw"/>
```

```
<sslDefault sslRef="mysslsettings"/>
.
.
.
<basicRegistry id="basic">
   <user
      name="izoa2018user"
      password="******"/>
   <user
      name="splunk-server.mycompany.com"
      password="******"/>
   <user
      name="pi-server.mycompany.com"
      password="******"/>
</basicRegistry>
.
.
.
<security-role name="restricted">
   <user name="izoa2018user"/>
   <user name="pi-server.mycompany.com"/>
   <user name="splunk-server.mycompany.com"/>
</security-role>
.
.
.
```

Code in file `pi_framework/config/cli.properties`

In the following example, the `client Keystore.p12` and client truststore (for the CLI of the Problem Insights Framework) that were generated by the commands are moved to the `pi_framework/config` directory:

```
#############################################
# PI host
#############################################
host=pi-server.mycompany.com
.
.
.
#############################################
# Keystore Location.
# On windows all paths must have double slashes
#############################################
keystoreLocation=/pi_framework/config/clientkeystore.p12

#############################################
# Keystore Type
#############################################
keystoreType=PKCS12

#############################################
# Keystore password
#############################################
keystorePass=clientkeystore_p12_pw

#############################################
# Truststore Location
# On windows all paths must have double slashes
#############################################
truststoreLocation=/pi_framework/config/client.truststore

#############################################
# Truststore Type
#############################################
truststoreType=JKS

#############################################
# Truststore password
#############################################
truststorePass=client_truststore_pw

#############################################
# Certificate Alias
#############################################
keyAlias=client-cert
```

Code in file `$SPLUNK_HOME/etc/apps/ibm_zoa_insights/local/probleminsights.conf`

In the following example, the `myPIservercert.pem`, `mySplunkclientcert.pem`, and `mySplunkclientcert.key` files (for the Splunk server) that were generated by the commands are moved to the `$SPLUNK_HOME/splunkadmin/ssl_cert` directory:

```
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
# ******************************************************************************
################################################################################
# This is NOT a default Splunk .conf file. This file is used by Problem
# Insights to specify security details related to interactions between Splunk and
# Problem Insights.
#
# To change or modify these values copy this file to this application's /local
# directory located at:
# $SPLUNK_HOME/etc/apps/ibm_zoa_insights/local
# and make the changes there
################################################################################
[Problem Insights Properties]
################################################################################
# URL of the Problem Insights server
################################################################################
PI_BASE_URL=https://pi-server.mycompany.com:9446
.
.
.
################################################################################
# Location of Problem Insights server certificate for 2-way
# authentication .pem only
################################################################################
SERVER_CERT=/home/splunkadmin/ssl_cert/myPIservercert.pem

################################################################################
# Location of Splunk's client certificate for 2-way authentication .pem only
################################################################################
CLIENT_CERT=/home/splunkadmin/ssl_cert/mySplunkclientcert.pem

################################################################################
# Location of client key file for 2-way authentication .key only
################################################################################
CLIENT_KEY=/home/splunkadmin/ssl_cert/mySplunkclientcert.key
```

Command reference for the Problem Insights Framework

This reference lists and describes the commands for operating the IBM Z Operations Analytics Problem Insights Framework.

Table 9. Commands for operating the Problem Insights Framework

Action	**Command**

Start the Framework.	**Linux system**
	`bin/analysis.sh start`
	Windows system
	`bin/analysis.bat start`
	This command loads the Message Libraries.

Stop the Framework.	**Linux system**
	`bin/analysis.sh stop`
	Windows system
	`bin/analysis.bat stop`

Import a Message Library.	**Linux system**
	`bin/analysis.sh importlibrary` *`xml_file_path`*
	Windows system
	`bin/analysis.bat importlibrary` *`xml_file_path`*

Table 9. Commands for operating the Problem Insights Framework (continued)

Action	Command

List the Message Libraries that were imported into the database.	**Linux system**
	`bin/analysis.sh getlibrarylist`
	Windows system
	`bin/analysis.bat getlibrarylist`
	This command gives you the Message Library ID and locale, which you must use in the commands for exporting or deleting a Message Library.

Export a Message Library.	**Linux system** `bin/analysis.sh exportlibrary` *`library_id library_locale`* **Windows system** `bin/analysis.bat exportlibrary` *`library_id library_locale`* **Tip:** To get the Message Library ID and locale, run the **`getlibrarylist`** command.
Delete a Message Library.	Use one of the following commands. The second command that is listed for each system includes the parameter **`f`**, which forces the deletion without input. **Linux system** • `bin/analysis.sh deletelibrary` *`library_id library_locale`* • `bin/analysis.sh deletelibrary -f` *`library_id library_locale`* **Windows system** • `bin/analysis.bat deletelibrary` *`library_id library_locale`* • `bin/analysis.bat deletelibrary -f` *`library_id library_locale`* **Tip:** To get the Message Library ID and locale, run the **`getlibrarylist`** command.
Delete data match results that are over 30 days old.	Use one of the following commands. The second command that is listed for each system includes the parameter **`f`**, which forces the deletion without input. **Linux system** • `bin/analysis.sh purgeResults` • `bin/analysis.sh purgeResults -f` **Windows system** • `bin/analysis.bat purgeResults` • `bin/analysis.bat purgeResults -f`

Encrypt text, such as a user name or password.	**Linux system**
	`bin/analysis.sh encrypt` *`text_to_encrypt`*
	Windows system
	`bin/analysis.bat encrypt` *`text_to_encrypt`*

Configuration property reference for the Problem Insights Framework

The IBM Z Operations Analytics Problem Insights Framework must have access to the Elastic Stack or Splunk platform so that it can analyze operational data. You must configure access to your platform by updating the configuration properties in the `config/PiSearch.config` file. This reference lists and describes the configuration properties in the `config/PiSearch.config`.

The configuration properties vary depending on your platform. The **`datasource.type`** property is the first property that you specify. It defines the platform that the Problem Insights Framework must communicate with. For more information about the properties, see one of the following topics, depending on your platform:

Elastic Stack platform
"Elastic Stack properties"

Splunk platform
"Splunk properties" on page 100

Elastic Stack properties

For the Elastic Stack platform, define the configuration properties for the IBM Z Operations Analytics Problem Insights Framework in the `config/PiSearch.config` file. You can configure either basic communication or secure communication between the Problem Insights Framework and the Elastic Stack platform.

Difference between configuring basic or secure communication

Table 10 shows the property values that vary depending on whether you configure basic or secure communication. "Properties reference for the Elastic Stack platform" provides descriptions and more information about each property.

Table 10. Property values that vary depending on whether you configure basic or secure communication between the Problem Insights Framework and the Elastic Stack platform

Property	Value for basic communication	Value for secure communication
`es.xpack.security.enabled`	`false`	`true`
`es.user`	Not applicable	Encrypted *username*
`es.password`	Not applicable	Encrypted *password*
`es.ssl.enabled`	`false`	`true`

Properties reference for the Elastic Stack platform

`datasource.type`
Specifies which data source to use during runtime. For the Elastic Stack platform, specify `Elastic` as the value.

`es.server`
Specifies the host name or IP address of the Elasticsearch master node, which must not contain the protocol information.

Default value
`localhost`

`es.port`
Specifies the Elasticsearch server port.

Default value
`9200`

`es.data.index.persist`
Specifies the storage index in which the results of the data matches (between incoming operational data and problem insights content) are stored.

Default value
`pi-zoa`

es.data.maxsize

Specifies the number of search result items to be returned. You can adjust this value according to the search interval and your data scale.

Default value

`5000`

es.xpack.security.enabled

Specifies whether X-Pack security is enabled for communication between the Problem Insights Framework and Elastic Stack. If X-Pack is installed in your environment, set this value to `true`.

Default value

`false`

es.user

Specifies the user name of the user who can access the Elasticsearch server.

"Command reference for the Problem Insights Framework" on page 97 includes the command that you can use to encrypt the user name.

es.password

Specifies the password of the user who can access the Elasticsearch server.

"Command reference for the Problem Insights Framework" on page 97 includes the command that you can use to encrypt the password.

es.ssl.enabled

Specifies whether SSL is enabled for communication between the Problem Insights Framework and Elastic Stack. If X-Pack is installed in your environment, and SSL is enabled, set this value to `true`, and verify that the `elasticsearch.xml` file contains the following fields:

- `xpack.security.http.ssl.enabled: true`
- `xpack.security.http.ssl.keystore.path:` *keystore*. For example, *keystore* might have a value of `certs/elastic-certificates.p12`.
- `xpack.security.http.ssl.truststore.path:` *truststore*. For example, *truststore* might have a value of `certs/elastic-certificates.p12`.

Default value

`false`

Tip: Transport Layer Security (TLS) is the cryptographic protocol that provides secure communications for your connections. Because the Secure Sockets Layer (SSL) protocol is the predecessor to TLS, the term *Secure Sockets Layer*, or *SSL*, is often used generically to refer to TLS encryption.

Splunk properties

For the Splunk platform, define the configuration properties for the IBM Z Operations Analytics Problem Insights Framework in the `config/PiSearch.config` file. You can configure either basic communication or secure communication between the Problem Insights Framework and the Splunk platform.

Difference between configuring basic or secure communication

The only difference between configuring basic or secure communication is the value that you set for the **splunk.scheme** property, as shown in the following table.

Value for basic communication	Value for secure communication
`splunk.scheme=http`	`splunk.scheme=https`

Properties reference for the Splunk platform

`datasource.type`
: Specifies which data source to use during runtime. For the Splunk platform, specify `Splunk` as the value.

`splunk.server`
: Specifies the host name or IP address of the Splunk Enterprise server, without the port for the protocol.

  Default value
  : `localhost`

`splunk.port`
: Specifies the Splunk Services port.

  Default value
  : `8089`

`splunk.user`
: Specifies the user name of the user who can access the Splunk Enterprise server. This user must have the user role or be authorized to send searches through the Java API.

  "Command reference for the Problem Insights Framework" on page 97 includes the command that you can use to encrypt the user name.

`splunk.password`
: Specifies the password of the user who can access the Splunk Enterprise server.

  "Command reference for the Problem Insights Framework" on page 97 includes the command that you can use to encrypt the password.

`splunk.index`
: Specifies the index for message data to search.

  Default value
  : `zosdex`

`splunk.scheme`
: Specifies the protocol to use for searches. Valid values are `http` or `https`.

  Default value
  : `http`

`splunk.sslVersion`
: Specifies the version of SSL that is used by Splunk. Valid values are `tls1`, `tls1.1`, and `tls1.2`. The preferred value is `tls1.2`.

  Default value
  : `tls1.2`

  Tip: Transport Layer Security (TLS) is the cryptographic protocol that provides secure communications for your connections. Because the Secure Sockets Layer (SSL) protocol is the predecessor to TLS, the term *Secure Sockets Layer*, or *SSL*, is often used generically to refer to TLS encryption.

`splunk.version`
: Specifies the version of Splunk that the Problem Insights Framework must communicate with.

  Default value
  : `7.0`

splunk.maxevents

Specifies the maximum number of events that are returned from Splunk. The preferred value is 5000.

Default value

5000

Message Library reference for the Problem Insights Framework

The Message Libraries in the IBM Z Operations Analytics Problem Insights Framework are defined by the XML Schema Definition (XSD) file messageLibrary.xsd. This reference includes, and describes, an example of a Message Library XML file. You can customize the Message Libraries to better represent problems that can occur in your environment.

Commands for operating Message Libraries

"Command reference for the Problem Insights Framework" on page 97 includes commands for operating (such as importing, exporting, or deleting) Message Libraries.

Key parts of a Message Library XML file

"Example of a Message Library" is an example of a Message Library. The following list describes some key parts of the file and how these parts correlate to the information in the Problem Insights dashboard of the Elastic Stack or Splunk UI:

- The domain (**domain** on the **message-library** element) specifies the domain of interest, such as CICS Transaction Server for z/OS, for the Message Library. In the UI, for each problem insight, the domain is shown in the Problem Insights table under the **Subsystem** column.
- The search interval (**search-interval** on the **message-library** element) specifies the time period between searches of the Problem Insights content in the Message Library.
- The message IDs (**id** on the **message** element) and source types (**source-types** element) are used as keys in searching for matches in the available operational data. Each source type correlates to the data source type specification in the IBM Common Data Provider for z Systems configuration of the associated data stream.
- For each message that is matched in the operational data, the associated suggested action (**suggested-action** element) and other resources (**other-resources** element) are shown in the Problem Insights table of the UI when you click the **View** link under the **Suggested action** column.

Example of a Message Library

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<message-library id="CICSforzOS" domain="CICS" version="3.2.0"
 country="US" language="en" search-interval="1minute">
    <description>Description of the Message Library<description>
    <source-types>
        <name>zOS-CICS-MSGUSR</name>
        <name>zOS-SYSLOG-Console</name>
    </source-types>
    <message id="BPXM023I" severity="2"
     domain="overwrites top level domain">
        <message-summary>
```

```
           A non-zero return code ...
        </message-summary>
        <suggested-action name="DFHSO0123_A">
            <action-summary>
               For a description of the return code,...
            </action-summary>
        </suggested-action>
        <other-resources>
            <resource>
               For a description of the return code,...
            </resource>
        </other-resources>
    </message>
    ...
</message-library>
```

Operational insights reference

IBM Z Operations Analytics provides operational insights for multiple domains of interest in your IT operations environment, including the z/OS system, databases, messaging, networks, security, transactions, and web servers. This reference correlates each domain of interest with the data sources that contribute to the insights for that domain of interest, and the configuration that must be done to send that data to IBM Z Operations Analytics. It also correlates the data sources with the dashboards that represent the operational data from those data sources and with the predefined searches for searching that operational data.

Operational insights are presented in different ways, depending on your platform, as described in Table 11.

Table 11. Presentation of operational insights depending on your platform

Platform	Presentation
IBM Operations Analytics - Log Analysis	• Problem Insights extension, which is focused on a defined set of potential problems that can occur in your IT environment and provides suggested actions for resolving these problems. For more information, see the following topics: – "Log Analysis extensions for z/OS Problem Insights and client-side Expert Advice" on page 13 – "Installing the z/OS Insight Packs, extensions, and IBM zAware data gatherer" on page 25 – "Configuring the Problem Insights extension" on page 46 – "Getting started with Problem Insights for z/OS" on page 56 • Dashboards and predefined searches in the user interface, which can help you identify, isolate, and resolve problems in your environment. For more information, see the following topics: – "Dashboards that represent the operational data" on page 141 – "Searches that are predefined for searching the operational data" on page 142

Table 11. Presentation of operational insights depending on your platform (continued)

Platform	Presentation
Elastic Stack and Splunk	• Problem Insights Framework, which is focused on a defined set of potential problems that can occur in your IT environment and provides suggested actions for resolving these problems. For more information, see “Problem Insights Framework” on page 80. • Dashboards and predefined searches in the user interface, which you can use to identify problems that might occur in your environment. For more information, see the following topics: – “Dashboards that represent the operational data” on page 141 – “Searches that are predefined for searching the operational data” on page 142

System insights

IBM Z Operations Analytics provides system insights that are based on data from the z/OS system.

Sources from which system data is retrieved

Insights are based on z/OS system data from the following sources:

- z/OS SYSLOG
- System Management Facilities (SMF) record type 30
- For the IBM Operations Analytics - Log Analysis platform only: zAware

Associated dashboards

See “Dashboards that represent the operational data” on page 141.

Associated searches

“z/OS system searches” on page 150

Database insights

IBM Z Operations Analytics provides database insights that are based on data from the Db2 for z/OS subsystem.

Sources from which database data is retrieved

Insights are based on Db2 for z/OS data from the following sources:

- z/OS SYSLOG
- For the Elastic Stack and Splunk platforms only: System Management Facilities (SMF) record type 100

Associated dashboards

See “Dashboards that represent the operational data” on page 141.

Associated searches

“Db2 for z/OS searches” on page 144

Messaging insights

IBM Z Operations Analytics provides messaging insights that are based on data from the MQ for z/OS subsystem.

Sources from which messaging data is retrieved

Insights are based on MQ for z/OS data from the z/OS SYSLOG.

Associated dashboards

See “Dashboards that represent the operational data” on page 141.

Associated searches

“MQ for z/OS searches” on page 146

Network insights

IBM Z Operations Analytics provides network insights that are based on data from, for example, the z/OS system, UNIX System Services, and the NetView for z/OS system.

Sources from which network data is retrieved

Insights are based on network data from the following sources:

- z/OS SYSLOG
- UNIX System Services system log (`syslogd`)
- NetView for z/OS program

Associated dashboards

See “Dashboards that represent the operational data” on page 141.

Associated searches

- “NetView for z/OS searches” on page 147
- “z/OS network searches” on page 150

Security insights

IBM Z Operations Analytics provides security insights that are based on data from, for example, the Resource Access Control Facility (RACF) and the Access Monitor component of IBM Security zSecure Admin.

Sources from which security data is retrieved

Insights are based on security data from the following sources:

- z/OS SYSLOG
- UNIX System Services system log (`syslogd`)
- Access Monitor component of IBM Security zSecure Admin
- System Management Facilities (SMF) record type 80

Associated dashboards

See "Dashboards that represent the operational data" on page 141.

Associated searches

- "Security searches: RACF" on page 147
- "Security searches: zSecure Access Monitor" on page 148

Transaction insights

IBM Z Operations Analytics provides transaction insights that are based on data from the CICS Transaction Server for z/OS and IMS for z/OS subsystems.

Sources from which transaction data is retrieved

Insights are based on transaction data from the following sources:

- CICS Transaction Server for z/OS data from the following sources:
  - z/OS SYSLOG
  - CICS Transaction Server for z/OS EYULOG and MSGUSR logs
  - System Management Facilities (SMF) record type 110
- IMS for z/OS data from the z/OS SYSLOG

Associated dashboards

See "Dashboards that represent the operational data" on page 141.

Associated searches

- "CICS Transaction Server for z/OS searches" on page 143
- "IMS for z/OS searches" on page 145

Web server insights

IBM Z Operations Analytics provides web server insights that are based on data from the WebSphere Application Server for z/OS subsystem.

Sources from which web server data is retrieved

Insights are based on WebSphere Application Server for z/OS data from the following sources:

- For the IBM Operations Analytics - Log Analysis platform only: WebSphere Application Server for z/OS High Performance Extensible Logging (HPEL)
- WebSphere Application Server for z/OS SYSOUT log
- WebSphere Application Server for z/OS SYSPRINT log
- System Management Facilities (SMF) record type 120

Associated dashboards

See "Dashboards that represent the operational data" on page 141.

Associated searches

"WebSphere Application Server for z/OS searches" on page 149

Data sources that contribute to the operational insights

For each type of source data that you want to gather to gain insights into your IT operations environment, this reference describes the configuration that must be done to send that data to IBM Z Operations Analytics.

Data that is provided by IBM Common Data Provider for z Systems

In the IBM Common Data Provider for z Systems Configuration Tool, you define data streams to gather source data and send that data to IBM Z Operations Analytics. For each type of source data that is provided by IBM Common Data Provider for z Systems, this reference describes how to define the associated data stream in a policy in the Configuration Tool.

Where necessary, this reference also describes how to enable the generation of the respective data at its source.

You can define data streams for the following types of source data:

- "CICS EYULOG and MSGUSR log data"
- "NetView message data" on page 110
- "SMF 30 data" on page 111
- "SMF 80 data" on page 112
- For the Elastic Stack and Splunk platforms only: "SMF 100 data" on page 122
- "SMF 110 data" on page 123
- "SMF 120 data" on page 129
- "SYSLOG data" on page 134
- "syslogd data" on page 135
- For the IBM Operations Analytics - Log Analysis platform only: "WebSphere HPEL data" on page 135
- "WebSphere SYSOUT data" on page 136
- "WebSphere SYSPRINT data" on page 137
- "zSecure data" on page 139

Data that is provided by IBM z Advanced Workload Analysis Reporter (IBM zAware)

"zAware interval anomaly data" on page 137 is applicable only to the IBM Operations Analytics - Log Analysis platform. To gather this data, you must configure the IBM zAware data gatherer, which is a component of IBM Z Operations Analytics on the Log Analysis platform.

CICS EYULOG and MSGUSR log data

CICS Transaction Server for z/OS EYULOG and MSGUSR log data includes information about the CICSPlex® System Manager (SM).

Data stream definition for CICS EYULOG and MSGUSR log data

Table 12. Data stream definition for CICS EYULOG and MSGUSR log data

Type of node in the policy	Required configuration value
Data Stream	For MSGUSR data, one or more of the following values: • **CICS User Messages**, with the default date format `MDY` • **CICS User Messages YMD**, with the date format `YMD` • **CICS User Messages DMY**, with the date format `DMY` For EYULOG data, one or more of the following values: • **CICS EYULOG**, with the default date format `MDY` • **CICS EYULOG YMD**, with the date format `YMD` • **CICS EYULOG DMY**, with the date format `DMY` **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **CICS Transaction Server**, and select the check box for the respective data stream.
Transcribe Transform	`UTF-8`
Splitter Transform	On the IBM Operations Analytics - Log Analysis platform, the Splitter Transform is not applicable because all data must be sent as unsplit. On the Elastic Stack and Splunk platforms, use the following values: • For **CICS MSGUSR**, the transform value is the corresponding CICS MSGUSR splitter. • For **CICS EYULOG**, the transform value is the corresponding CICS EYULOG splitter.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 140.

NetView message data

NetView message data includes network data from the IBM Tivoli NetView for z/OS program.

Data stream definition for NetView message data

Table 13 on page 111 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 13. Data stream definition for NetView message data

Type of node in the policy	Required configuration value
Data Stream	**NetView Netlog** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **Network** > **NetView**, and select the **NetView Netlog** check box.
Transcribe Transform	`UTF-8`
Splitter Transform	On the IBM Operations Analytics - Log Analysis platform, the Splitter Transform is not applicable because all data must be sent as unsplit. On the Elastic Stack and Splunk platforms, use `NetView Splitter`.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 140.

SMF 30 data

System Management Facilities (SMF) record type 30 data is job performance data (based on accounting data) for z/OS software.

- "SMF 30 data generation"
- "Data stream definition for SMF 30 data"
- "Annotated fields for SMF 30 data" on page 112

SMF 30 data generation

To enable the generation of SMF record type 30 data, you must include the SMF 30 record type in the single SMF log stream that the IBM Common Data Provider for z Systems System Data Engine processes.

Data stream definition for SMF 30 data

For prerequisite requirements for defining SMF data streams, see "Enabling support for SMF data destined for IBM Z Operations Analytics" on page 3.

Table 14 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 14. Data stream definition for SMF 30 data

Type of node in the policy	Required configuration value
Data Stream	**SMF30** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **z/OS** > **Address Space**, and select the **SMF30** check box.
Transcribe Transform	`UTF-8`
Splitter Transform	Not applicable
Filter Transform	Not required

Table 14. Data stream definition for SMF 30 data (continued)

Type of node in the policy	Required configuration value
Subscriber	See "Subscribers for each type of source data" on page 140.

Annotated fields for SMF 30 data

Table 15. Annotated fields for SMF 30 data

Field	Description
`CPU`	The CPU usage for the monitored task
`IORate`	The I/O rate for the monitored task
`JobName`	The 8-character name of the job on the z/OS system
`PagingRate`	The paging rate for the monitored task
`ProgName`	The name of the program that is running under the monitored task
`RecordType`	The type of SMF record
`SystemID`	The system identifier
`Task`	The job name for the task that issued the message
`WorkingSet`	The working set size for the monitored task

SMF 80 data

System Management Facilities (SMF) record type 80 data is produced during Resource Access Control Facility (RACF) processing.

- "SMF 80 data generation"
- "Data stream definition for SMF 80 data" on page 113
- "Annotated fields for SMF 80 data" on page 113

SMF 80 data generation

To enable the generation of SMF record type 80 data, you must include the SMF 80 record type in the single SMF log stream that the IBM Common Data Provider for z Systems System Data Engine processes. RACF must also be installed, active, and configured to protect resources.

For information about the subset of SMF record type 80 data that the System Data Engine collects, see "SMF type 80-related records that the System Data Engine creates" on page 115.

SMF also records information that is gathered by RACF auditing. By using various RACF options, you can regulate the granularity of SMF record type 80 data that is collected. In the IBM Knowledge Center, see the following information from the z/OS documentation:

- Information about the following options of the SETROPTS LOGOPTIONS command, through which you can control auditing:
  - `DIRSRCH`
  - `DIRACC`
  - `FSOBJ`
  - `FSSEC`

- Examples for setting audit controls by using SETROPTS

Before you enable RACF log options, consider the impact in your environment. For example, enabling RACF log options can result in the following consequences:

- An increase in the amount of disk space that is used for logging
- An increase in the network activity that is required to transmit SMF data

Data stream definition for SMF 80 data

For prerequisite requirements for defining SMF data streams, see “Enabling support for SMF data destined for IBM Z Operations Analytics” on page 3.

Table 16 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 16. Data stream definition for SMF 80 data

Type of node in the policy	Required configuration value
Data Stream	One or more of the following values: • **SMF80_COMMAND** • **SMF80_LOGON** • **SMF80_OMVS_RES_1** • **SMF80_OMVS_RES_2** • **SMF80_OMVS_SEC_1** • **SMF80_OMVS_SEC_2** • **SMF80_OPERATION** • **SMF80_RESOURCE** **To select this data stream in the Configuration Tool:** In the “Select data stream” window, click **IBM Z Operations Analytics** > **Security** > **RACF**, and select the check box for the respective data stream.
Transcribe Transform	`UTF-8`
Splitter Transform	Not applicable
Filter Transform	Not required
Subscriber	See “Subscribers for each type of source data” on page 140.

Annotated fields for SMF 80 data

In the following table, the column that is titled “Corresponding SMF field” indicates the name of the SMF field that corresponds to the field name in the annotation.

Table 17. Annotated fields for SMF 80 data

Field	Description	Corresponding SMF field
`AccessAllow`	Access authority allowed	SMF80DTA
`AccessReq`	Access authority requested	SMF80DTA

Table 17. Annotated fields for SMF 80 data (continued)

Field	Description	Corresponding SMF field
AccessType	Setting that is used in granting access. The following values are possible: • None • Owner • Group • Other	SMF80DA2
Application	Application name that is specified on the RACROUTE request	SMF80DTA
AuditDesc	Descriptive name of the operation that is audited	SMF80DA2
AuditName	Name of the operation that is audited	SMF80DA2
Auditor	AUDITOR attribute (Y/N)	SMF80ATH
AuditorExec	Auditor execute/search audit options	SMF80DA2
AuditorRead	Auditor read access audit options	SMF80DA2
AuditorUserExec	User execute/search audit options	SMF80DA2
AuditorUserRead	User read access audit options	SMF80DA2
AuditorUserWrite	User write access audit options	SMF80DA2
AuditorWrite	Auditor write access audit options	SMF80DA2
AuthorityFlags	Flags that indicate the authority checks that are made for the user who requested the action	SMF80ATH
CHOWNGroupID	z/OS UNIX group identifier (GID) input parameter	SMF80DA2
CHOWNUserID	z/OS UNIX user identifier (UID) input parameter	SMF80DA2
Class	The class entries that are supplied by IBM in the class descriptor table (ICHRRCDX)	SMF80DTA
Command	A string that is derived by using the SMF80EVT and SMF80EVQ values	SMF80EVT, SMF80EVQ
EffectiveGroup	User's effective GID setting	SMF80DA2
EffectiveUser	User's effective UID setting	SMF80DA2
Event	Short description of the event code and qualifier	SMF80EVT, SMF80EVQ
EventCode	Event code	SMF80EVT
EventDate	Date that the event occurred	SMF80DTE
EventDesc	Verbose description of the event code and qualifier	SMF80EVT
EventQual	Event code qualifier	SMF80EVQ
Failed	Event code qualifier is nonzero, which indicates a failed request (Y/N)	SMF80EVQ
Filename	File name of the file that is being checked	SMF80DA2
FileOwnerGroup	File owner's GID	SMF80DA2
FileOwnerUser	File owner's UID	SMF80DA2
Generic	Generic profile used (Y/N)	SMF80DTP
GroupExec	Group permissions bit: execute	SMF80DA2
GroupRead	Group permissions bit: read	SMF80DA2
GroupWrite	Group permissions bit: write	SMF80DA2
ISGID	Requested file mode: S_ISGID bit	SMF80DA2
ISUID	Requested file mode: S_ISUID bit	SMF80DA2
ISVTX	Requested file mode: S_ISVTX bit	SMF80DA2
OtherExec	Other permissions bit: execute	SMF80DA2

Table 17. Annotated fields for SMF 80 data (continued)

Field	Description	Corresponding SMF field
OtherRead	Other permissions bit: read	SMF80DA2
OtherWrite	Other permissions bit: write	SMF80DA2
OwnerExec	Owner permissions bit: execute	SMF80DA2
OwnerRead	Owner permissions bit: read	SMF80DA2
OwnerWrite	Owner permissions bit: write	SMF80DA2
Pathname	Full path name of the file that is being checked	SMF80DA2
ProfileName	Name of the Resource Access Control Facility (RACF) profile that is used to access the resource	SMF80DTA
RealGroup	User's real GID setting	SMF80DA2
RealUser	User's real UID setting	SMF80DA2
RecordType	Internal record type. The following values are possible: • SMF80_COMMAND • SMF80_LOGON • SMF80_OMVS_RES_1 • SMF80_OMVS_RES_2 • SMF80_OMVS_SEC_1 • SMF80_OMVS_SEC_2 • SMF80_OPERATION • SMF80_RESOURCE For information about these values, see the IBM Common Data Provider for z Systems documentation in the IBM Knowledge Center.	Set by the data provider
ResourceName	Resource name	SMF80DTA
SavedGroup	User's saved GID setting	SMF80DA2
SavedUser	User's saved UID setting	SMF80DA2
Special	SPECIAL attribute (Y/N)	SMF80ATH
SuperUser	z/OS UNIX superuser (Y/N)	SMF80AU2
SystemID	The system identifier from the SID parameter in the SMFPRMnn member	SMF80SID
TermID	Terminal ID of the foreground user (zero if not available)	SMF80TRM
UserID	Identifier of the user that is associated with this event. The value of JobName is used if the user is not defined to RACF.	SMF80USR

SMF type 80-related records that the System Data Engine creates

The IBM Common Data Provider for z Systems System Data Engine collects a subset of the SMF data that is generated by the Resource Access Control Facility (RACF). This reference describes the types of records that the System Data Engine creates as it extracts relevant data from SMF type 80 records.

The System Data Engine creates the following record types:

- SMF80_COMMAND
- SMF80_LOGON
- SMF80_OMVS_RES_1
- SMF80_OMVS_RES_2
- SMF80_OMVS_SEC_1

- SMF80_OMVS_SEC_2
- SMF80_OPERATION
- SMF80_RESOURCE

From each SMF type 80 record that it collects, the System Data Engine uses the following information to determine what data to extract:

- SMF event in the **SMF80EVT** field
- RACF event code qualifier in the **SMF80EVQ** field

The System Data Engine excludes SMF events that occur for hierarchical storage management (HSM), for example, events where the value of the user ID SMF80USR is HSM.

For more information about SMF record type 80 records, see the following topics from the z/OS documentation in the IBM Knowledge Center:

- SMF record type 80: RACF processing record
- Format of SMF record type 80 records
- SMF record type 80 event codes and event code qualifiers

SMF80_COMMAND record type

SMF record type 80 records for events 8 - 25 are created when RACF commands fail because the user who ran them does not have sufficient authority. Relevant fields from these SMF event records are stored in the SMF80_COMMAND records that are created by the System Data Engine.

Table 18 describes the event code qualifiers for events 8 - 25, which provide more information about why the command failed.

Table 18. SMF80_COMMAND record type: event code qualifiers for events 8 - 25

Event code qualifier	Description
1	Insufficient authority
2	Keyword violations detected
3	Successful listing of data sets
4	System error in listing of data sets

SMF80_LOGON record type

SMF record type 80 records for event 1 are created when RACF authentication fails because of incorrect user credentials, which prevents the user from accessing the system. Relevant fields from this SMF event record are stored in the SMF80_LOGON records that are created by the System Data Engine.

Table 19 describes the event code qualifiers for event 1, which provide more information about why the logon failed.

Table 19. SMF80_LOGON record type: event code qualifiers for event 1

Event code qualifier	Description
1	Invalid password
2	Invalid group
3	Invalid object identifier (OID) card
4	Invalid terminal/console

Table 19. SMF80_LOGON record type: event code qualifiers for event 1 (continued)

Event code qualifier	Description
5	Invalid application
6	Revoked user ID attempting access
7	User ID automatically revoked
9	Undefined user ID
10	Insufficient security label authority
11	Not authorized to security label
14	System now requires more authority
15	Remote job entry—job not authorized
16	Surrogate class is inactive
17	Submitter is not authorized by user
18	Submitter is not authorized to security label
19	User is not authorized to job
20	Warning—insufficient security label authority
21	Warning—security label missing from job, user, or profile
22	Warning—not authorized to security label
23	Security labels not compatible
24	Warning—security labels not compatible
25	Current password has expired
26	Invalid new password
27	Verification failed by installation
28	Group access has been revoked
29	Object identifier (OID) card is required
30	Network job entry—job not authorized
31	Warning—unknown user from trusted node propagated
32	Successful initiation using PassTicket
33	Attempted replay of PassTicket
34	Client security label not equivalent to servers
35	User automatically revoked due to inactivity
36	Passphrase is not valid
37	New passphrase is not valid
38	Current passphrase has expired
39	No RACF user ID found for distributed identity

SMF80_OMVS_RES record types

SMF record type 80 records for events 28 - 30 are created when the following z/OS UNIX operations occur: directory search, check access to directory, or check access to file. Relevant fields from these SMF event records are stored in the SMF80_OMVS_RES_1 and SMF80_OMVS_RES_2 records that are created by the System Data Engine.

Table 20 describes the event code qualifiers for events 28 - 30, which provide more information about the operation results.

Table 20. SMF80_OMVS_RES_1 and SMF80_OMVS_RES_2 record types: event code qualifiers for events 28 - 30

Event code qualifier	Description
0	Access allowed
1	Not authorized to search directory
2	Security label failure

SMF80_OMVS_SEC record types

SMF record type 80 records for events 31 and 33 - 35 are created when the z/OS UNIX commands **CHAUDIT**, **CHMOD**, or **CHOWN** are entered, or when the SETID bits for a file are cleared. Relevant fields from these SMF event records are stored in the SMF80_OMVS_SEC_1 and SMF80_OMVS_SEC_2 records that are created by the System Data Engine.

Table 21, Table 22, Table 23, and Table 24 on page 119 describe the event code qualifiers for events 31 and 33 - 35, which provide more information about the operation results.

Table 21. SMF80_OMVS_SEC_1 and SMF80_OMVS_SEC_2 record types: event code qualifiers for event 31

Event code qualifier	Description
0	File's audit options changed
1	Caller does not have authority to change user audit options of specified file
2	Caller does not have authority to change auditor audit options
3	Security label failure

Table 22. SMF80_OMVS_SEC_1 and SMF80_OMVS_SEC_2 record types: event code qualifiers for event 33

Event code qualifier	Description
0	File's mode changed
1	Caller does not have authority to change mode of specified file
2	Security label failure

Table 23. SMF80_OMVS_SEC_1 and SMF80_OMVS_SEC_2 record types: event code qualifiers for event 34

Event code qualifier	Description
0	File's owner or group owner changed
1	Caller does not have authority to change owner or group owner of specified file

Table 23. SMF80_OMVS_SEC_1 and SMF80_OMVS_SEC_2 record types: event code qualifiers for event 34 (continued)

Event code qualifier	Description
2	Security label failure

Table 24. SMF80_OMVS_SEC_1 and SMF80_OMVS_SEC_2 record types: event code qualifiers for event 35

Event code qualifier	Description
0	S_ISUID, S_ISGID, and S_ISVTX bits changed to zero (write).

SMF80_OPERATION record type

SMF record type 80 records for events 2 - 7 are created when a z/OS resource that is protected by RACF is updated, deleted, or accessed by a user that is defined to RACF with the SPECIAL attribute. Relevant fields from these SMF event records are stored in the SMF80_OPERATION records that are created by the System Data Engine.

Table 25, Table 26 on page 120, Table 27 on page 120, Table 28 on page 120, Table 29 on page 120, and Table 30 on page 121 describe the event code qualifiers for events 2 - 7, which provide more information about the operation results.

Table 25. SMF80_OPERATION record type: event code qualifiers for event 2

Event code qualifier	Description
0	Successful access
1	Insufficient authority
2	Profile not found—RACFIND specified on macro
3	Access permitted due to warning
4	Failed due to PROTECTALL SETROPTS
5	Warning issued due to PROTECTALL SETROPTS
6	Insufficient category/SECLEVEL
7	Insufficient security label authority
8	Security label missing from job, user, or profile
9	Warning—insufficient security label authority
10	Warning—data set not cataloged
11	Data set not cataloged
12	Profile not found—required for authority checking
13	Warning—insufficient category/SECLEVEL
14	Warning—non-main execution environment
15	Conditional access allowed via basic mode program

Table 26. `SMF80_OPERATION` record type: event code qualifiers for event 3

Event code qualifier	Description
0	Successful processing of new volume
1	Insufficient authority
2	Insufficient security label authority
3	Less specific profile exists with different security label

Table 27. `SMF80_OPERATION` record type: event code qualifiers for event 4

Event code qualifier	Description
0	Successful rename
1	Invalid group
2	User not in group
3	Insufficient authority
4	Resource name already defined
5	User not defined to RACF
6	Resource not protected SETROPTS
7	Warning——resource not protected SETROPTS
8	User in second qualifier is not RACF defined
9	Less specific profile exists with different security label
10	Insufficient security label authority
11	Resource not protected by security label
12	New name not protected by security label
13	New security label must dominate old security label
14	Insufficient security label authority
15	Warning—resource not protected by security label
16	Warning—new name not protected by security label
17	Warning—new security label must dominate old security label

Table 28. `SMF80_OPERATION` record type: event code qualifiers for event 5

Event code qualifier	Description
0	Successful scratch
1	Resource not found
2	Invalid volume

Table 29. `SMF80_OPERATION` record type: event code qualifiers for event 6

Event code qualifier	Description
0	Successful deletion

Table 30. SMF80_OPERATION record type: event code qualifiers for event 7

Event code qualifier	Description
0	Successful definition
1	Group undefined
2	User not in group
3	Insufficient authority
4	Resource name already defined
5	User not defined to RACF
6	Resource not protected
7	Warning—resource not protected
8	Warning—security label missing from job, user, or profile
9	Insufficient security label authority
10	User in second qualifier in not defined to RACF
11	Insufficient security label authority
12	Less specific profile exists with a different security label

SMF80_RESOURCE record type

SMF record type 80 records for event 2 are created when a z/OS resource that is protected by RACF is updated, deleted, or accessed by a user. Relevant fields from these SMF event records are stored in the SMF80_RESOURCE records that are created by the System Data Engine.

Table 31 describes the event code qualifiers for event 2, which provide more information about the operation results.

Table 31. SMF80_RESOURCE record type: event code qualifiers for event 2

Event code qualifier	Description
0	Successful access
1	Insufficient authority
2	Profile not found—RACFIND specified on macro
3	Access permitted due to warning
4	Failed due to PROTECTALL SETROPTS
5	Warning issued due to PROTECTALL SETROPTS
6	Insufficient category/SECLEVEL
7	Insufficient security label authority
8	Security label missing from job, user, or profile
9	Warning—insufficient security label authority
10	Warning—data set not cataloged
11	Data set not cataloged

Table 31. SMF80_RESOURCE record type: event code qualifiers for event 2 (continued)

Event code qualifier	Description
12	Profile not found—required for authority checking
13	Warning—insufficient category/SECLEVEL
14	Warning—non-main execution environment
15	Conditional access allowed via basic mode program

SMF 100 data

System Management Facilities (SMF) record type 100 data is generated by Db2 for z/OS.

- "SMF 100 data generation"
- "Data stream definition for SMF 100_1 data"
- "Annotated fields for SMF 100_1 data" on page 123

SMF 100 data generation

When it sends data to IBM Z Operations Analytics, the IBM Common Data Provider for z Systems System Data Engine collects only a subset of the SMF record type 100 data that is generated by Db2 for z/OS. It collects data from SMF type 100 subtype 1. To enable the generation of this data, you must enable the following trace options in Db2 for z/OS:

```
START TRACE(STAT) DEST(SMF) CLASS(1,2,3)
```

Data stream definition for SMF 100_1 data

Tip: This data stream can be defined only for the Elastic Stack and Splunk platforms.

For prerequisite requirements for defining SMF data streams, see "Enabling support for SMF data destined for IBM Z Operations Analytics" on page 3.

Table 32 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 32. Data stream definition for SMF 100_1 data

Type of node in the policy	Required configuration value
Data Stream	**SMF100_1** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **Database** > **Db2** > **Statistical**, and select the **SMF100_1** check box.
Transcribe Transform	UTF-8
Splitter Transform	Not applicable
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 140.

Annotated fields for SMF 100_1 data

Table 33. Annotated fields for SMF 100_1 data

Field	Description
sysplex	Sysplex name
system	System name
hostname	Host name
sourcename	Source name
timezone	Time zone offset
UPDATE_NAME	This value is always IZOA, which represents "IBM Z Operations Analytics."
UPDATE_VERSION	This value is always 3.2, which represents Version 3.2 of IBM Z Operations Analytics.
TIMESTAMP	Time stamp
MVS_SYSTEM_ID	MVS™ system ID, which is also the SMF system ID
DB2_SYSTEM_ID	Db2 system ID, which is also the SMF subsystem ID
SSID	Subsystem ID
DEADLOCKS	The number of times that deadlocks were detected
TOTAL_SUSPENDS	The sum of the values of the following fields: • LOCK_SUSPENDS • LATCH_SUSPENDS • OTHER_SUSPENDS
LOCK_SUSPENDS	The number of times that a lock cannot be obtained, and the unit of work is suspended
LATCH_SUSPENDS	The number of latch suspensions
OTHER_SUSPENDS	The number of suspensions that are caused by something other than lock or latch
LOCK_TIMEOUTS	The number of times that a unit of work was suspended for a period of time that exceeds the timeout value

SMF 110 data

System Management Facilities (SMF) record type 110 data is generated by CICS Transaction Server for z/OS.

- "SMF 110 data generation"
- "SMF110_E record type for monitoring exceptions data" on page 124
- "SMF110_S_10 for global transaction manager statistics data" on page 126

SMF 110 data generation

When it sends data to IBM Z Operations Analytics, the IBM Common Data Provider for z Systems System Data Engine collects only a subset of the SMF record type 110 data that is generated by CICS Transaction Server for z/OS. It collects the following data from SMF record type 110:

- Monitoring exceptions data for CICS Transaction Server for z/OS from SMF type 110 subtype 1 records, with a class where data = 4
- Global transaction manager statistics data for CICS Transaction Server for z/OS from SMF type 110 subtype 2 records, with a class where STID = 10

To enable the generation of SMF record type 110 data, you must include the SMF 110 record type in the single SMF log stream that the System Data Engine processes. You must also define the following CICS Transaction Server for z/OS initialization parameters in the SYSIN data set of the CICS startup job stream:

```
STATRCD=ON,             Interval statistics recording
STATINT=001000,         Interval definition
MN=ON,                  Turn monitoring on or off
MNEXC=ON,               Exceptions monitoring
MNRES=ON,               Resource monitoring
```

For more information about enabling the generation of SMF record type 110 data, see Specifying system initialization parameters before startup in the CICS Transaction Server for z/OS Version 5.3 documentation.

The System Data Engine creates the following record types as it extracts the relevant data from SMF type 110 records:

- `zOS-SMF110_E` for monitoring exceptions data
- `zOS-SMF110_S_10` for global transaction manager statistics data

The monitoring exceptions records contain information about CICS Transaction Server for z/OS resource shortages that occur during a transaction, such as queuing for file strings and waiting for temporary storage. This data highlights possible problems in CICS system operation. It can help you identify system constraints that affect the performance of your transactions. CICS writes one exception record for each exception condition that occurs.

The global transaction manager records contain transactions summary information for CICS Transaction Server for z/OS. This data can give you a more holistic view of the CICS region, including a comparison among the current and peak numbers of transactions that are running in the region, and the maximum number of allowed transactions.

`SMF110_E` record type for monitoring exceptions data

`SMF110_E` records contain information about CICS Transaction Server for z/OS resource shortages that occur during a transaction, such as queuing for file strings and waiting for temporary storage. This data highlights possible problems in CICS system operation. It can help you identify system constraints that affect the performance of your transactions. CICS writes one exception record for each exception condition that occurs.

- "Data stream definition for SMF 110_E data"
- "Annotated fields for SMF 110_E data" on page 125

Data stream definition for SMF 110_E data

For prerequisite requirements for defining SMF data streams, see "Enabling support for SMF data destined for IBM Z Operations Analytics" on page 3.

Table 34 on page 125 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 34. Data stream definition for SMF 110_E data

Type of node in the policy	Required configuration value
Data Stream	**SMF110_E** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **CICS Transaction Server**, and select the **SMF110_E** check box.
Transcribe Transform	`UTF-8`
Splitter Transform	Not applicable
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 140.

Annotated fields for SMF 110_E data

In the following table, the column that is titled "Corresponding SMF field" indicates the name of the SMF field that corresponds to the field name in the annotation.

Table 35. Annotated fields for SMF 110_E data

Field	Description	Corresponding SMF field
`ApplID`	The product name (Generic APPLID)	SMFMNPRN
`ApplIDSpec`	The product name (Specific APPLID)	SMFMNSPN
`BridgeTransID`	The bridge transaction ID	EXCMNBTR
`CICSTrans`	The transaction identification	EXCMNTRN
`ExceptionEnd`	The exception stop time	EXCMNSTO
`ExceptionID`	The exception ID	EXCMNRIX
`ExceptionID2`	The extended exception ID	EXCMNRIX
`ExceptionLen`	The exception resource ID length	EXCMNRIL
`ExceptionNumber`	The exception sequence number for the task	EXCMNEXN
`ExceptionStart`	The exception start time	EXCMNSTA
`ExceptionType`	The exception type	EXCMNTYP
`JobName`	The 8-character name of the job on the z/OS system	SMFMNJBN
`LU`	The real logical unit on the z/OS system	EXCMNRLU
`LUName`	The logical unit on the z/OS system	EXCMNLUN
`NetID`	The `NETID` if a network qualified name was received from z/OS Communications Server. For a z/OS Communications Server resource where the network qualified name was not yet received, `NETID` is eight blanks. In all other cases, this field is null.	EXCMNNID
`ProgName`	The name of the currently running program for the user task when the exception condition occurred	EXCMNCPN

Table 35. Annotated fields for SMF 110_E data (continued)

Field	Description	Corresponding SMF field
RecordType	The internal record type, which is SMF110_E	Set by the data provider
RecordVersion	The record version in CICS Transaction Server for z/OS	SMFMNRVN
ReportClass	The report class name	EXCMNRPT
ResourceID	The exception resource identification	EXCMNRID
ResourceType	The exception resource type	EXCMNRTY
ServiceClass	The service class name	EXCMNSRV
SubsystemID	The subsystem identification	SMFMNSSI
SystemID	The system identifier from the **SID** parameter in the SMFPRMnn member	SMFMNSID
TerminalID	The terminal identification	EXCMNTER
TranClassName	The transaction class name	EXCMNTCN
TransFacName	The transaction facility name	EXCMNFCN
TransFlags	The transaction flags. For more information about these flags, see the description of the 8-byte TRANFLAG field at offset 164 in Performance data in group DFHTASK in the CICS Transaction Server for z/OS Version 5.3 documentation.	EXCMNTRF
TransNum	The transaction identification number	EXCMNTNO
TransPriority	The transaction priority	EXCMNTPR
UORID	Resource management services (RRMS) MVS unit of recovery identification	EXCMNURI
UOWName	The network unit-of-work suffix	EXCMNNSX
UserID	The user identification at task creation. This identifier can also be the remote user identifier for a task that is created as the result of receiving an ATTACH request across a multiregion operation (MRO) or Advanced Program-to-Program Communication (APPC) link with attach-time security enabled.	EXCMNUSR
zCSName	The network unit-of-work prefix	EXCMNNPX

SMF110_S_10 for global transaction manager statistics data

SMF110_S_10 records contain transactions summary information for CICS Transaction Server for z/OS. This data can give you a more holistic view of the CICS region, including a comparison among the current and peak numbers of transactions that are running in the region, and the maximum number of allowed transactions.

- “Data stream definition for SMF110_S_10 data” on page 127
- “Annotated fields for SMF110_S_10 data” on page 127

Data stream definition for SMF110_S_10 data

For prerequisite requirements for defining SMF data streams, see “Enabling support for SMF data destined for IBM Z Operations Analytics” on page 3.

Table 36 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 36. Data stream definition for SMF110_S_10 data

Type of node in the policy	Required configuration value
Data Stream	**SMF110_S_10** **To select this data stream in the Configuration Tool:** In the “Select data stream” window, click **IBM Z Operations Analytics** > **CICS Transaction Server**, and select the **SMF110_S_10** check box.
Transcribe Transform	`UTF-8`
Splitter Transform	Not applicable
Filter Transform	Not required
Subscriber	See “Subscribers for each type of source data” on page 140.

Annotated fields for SMF110_S_10 data

In the following table, the column that is titled “Corresponding SMF field” indicates the name of the SMF field that corresponds to the field name in the annotation.

Table 37. Annotated fields for SMF110_S_10 data

Field	Description	Corresponding SMF field
`ApplID`	The product name (Generic APPLID)	SMFSTPRN
`AtsMxt`	An indicator of the limit for the number of concurrent tasks	XMGATMXT
`GmtsLast_TxnAttch`	The time when the last transaction was attached	XMGGTAT
`GmtsMxtReached`	According to Greenwich mean time (GMT), the time when the task limit (the value of `MAXTASKS`) was met	XMGGAMXT
`GmtsMxtSet`	According to Greenwich mean time (GMT), the time when the task limit (the value of `MAXTASKS`) was set	XMGGSMXT
`IntervalDuration`	For a status type (`StatsType`) of `INT`, the interval duration, which is represented in the time format `HHMMSS`	SMFSTINT
`LclsLast_TxnAttch`	The date and time when the last transaction was attached	XMGLTAT

Table 37. Annotated fields for SMF110_S_10 data (continued)

Field	Description	Corresponding SMF field
LclsMxtReached	The local time when the task limit (the value of MAXTASKS) was met	XMGLAMXT
LclsMxtSet	The local time when the task limit (the value of MAXTASKS) was set	XMGLSMXT
MAXTASKS	The limit for the number of concurrent tasks	XMGMXT
RecordIncomplete	An indicator that is set to YES if incomplete data is recorded	SMFSTICD
RecordType	The internal record type, which is SMF110_S_10	Set by the data provider
RecordVersion	The record version in the following format: x'0vrm'	SMFSTRVN
StatsArea	The status area	Set by the data provider
StatsType	The status type. For example, one of the following types: • EOD • INT • REQ • RRT • USS	SMFSTRQT
SystemID	The system identifier from the **SID** parameter in the SMFPRMnn member	SMFMNSID
TransCount	The number of user and system transactions that are attached	XMGNUM
TransCurrentActiveUser	At the present time, the number of active user transactions in the system	XMGCAT
TransCurrent_QSec	At the present time, the number of seconds that transactions are queued because the task limit (the value of MAXTASKS) was met	W_CUR_Q_TIME
TransPeakActiveUser	The highest number of active user transactions	XMGPAT
TransPeakQueued	The highest number of queued user transactions	XMGPQT
TransQueuedUser	The number of queued user transactions in the system	XMGCQT
TransTimesAtMAXTASKS	The number of times that the task limit (the value of MAXTASKS) was met	XMGTAMXT
TransTotalActive	For a specified time interval, the number of active user transactions in the system	XMGTAT

Table 37. Annotated fields for SMF110_S_10 data (continued)

Field	Description	Corresponding SMF field
TransTotalDelayed	For a specified time interval, the number of user transactions that were delayed because the task limit (the value of MAXTASKS) was met	XMGTDT
TransTotal_QSec	For a specified time interval, the number of seconds that transactions were queued because the task limit (the value of MAXTASKS) was met	W_TOT_Q_TIME
TransTotalTasks	At the time of the last reset, the number of transactions in the system	XMGTNUM

SMF 120 data

System Management Facilities (SMF) record type 120 data is generated by WebSphere Application Server for z/OS.

- "SMF record type 120 data generation"
- "Data stream definition for SMF 120 data" on page 130
- "Annotated fields for SMF 120 data" on page 130

SMF record type 120 data generation

When it sends data to IBM Z Operations Analytics, the IBM Common Data Provider for z Systems System Data Engine collects only a subset of the SMF record type 120 data that is generated by WebSphere Application Server for z/OS. It collects performance data from SMF record type 120 subtype 9. The default SMF type 120 subtype 9 record contains information for properly monitoring the performance of your EJB components and web applications.

Restriction: This performance data does not include data for the WebSphere Liberty server.

To enable the generation of SMF record type 120 data, you must include the SMF 120 record type in the single SMF log stream that the IBM Common Data Provider for z Systems System Data Engine processes. Also, for each application server instance that you want to monitor, you must specify properties for SMF data collection by setting WebSphere Application Server for z/OS environment variables from the WebSphere Application Server Administrative Console. For more information about enabling the generation of SMF record type 120 data, see Using the administrative console to enable properties for specific SMF record types in the WebSphere Application Server for z/OS Version 9.0 documentation.

The System Data Engine creates the following record types as it extracts the performance data from SMF type 120 subtype 9 records:

- SMF120_REQAPPL for WebSphere application records
- SMF120_REQCONT for WebSphere controller records

The SMF type 120 subtype 9 record contains information about the activity of the WebSphere server and the hosted applications. This record is produced whenever a server receives a request. When you do capacity planning, consider the costs that are involved in running requests and the number of requests that you process during a specific time. You can use the SMF type 120 subtype 9 record to monitor which requests are associated with which applications, the number of requests that occur, and the amount of resource that each request uses. You can also use this record to identify the applications that are involved and the amount of CPU time that the requests use.

As part of planning to collect SMF 120 data, consider the disk space requirements for storing the data and the increase in network activity that is required to transmit SMF data.

To reduce any system performance degradation due to data collection and to improve the usability of the data, the System Data Engine aggregates the SMF activity records in 1-minute collection intervals by default. Ensure that the collection interval is an integral factor of the SMF global recording interval, as measured in minutes, so that data collection is synchronized. For example, a 1-, 3-, or 5-minute collection interval is an integral factor of a typical 15-minute SMF global recording interval, but a 4-minute collection interval is not. The SMF global recording interval `INTERVAL(nn)` is defined in the `SMFPRMxx` member of `SYS1.PARMLIB` (or its equivalent).

Data stream definition for SMF 120 data

For prerequisite requirements for defining SMF data streams, see "Enabling support for SMF data destined for IBM Z Operations Analytics" on page 3.

"SMF 120 data" on page 129 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 38. Data stream definition for SMF 120 data

Type of node in the policy	Required configuration value
Data Stream	One or more of the following values: • **SMF120_REQAPPL** • **SMF120_REQCONT** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **Web Servers** > **WebSphere Application Server**, and select the check box for the respective data stream.
Transcribe Transform	`UTF-8`
Splitter Transform	Not applicable
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 140.

Annotated fields for SMF 120 data

In the following table, the column that is titled "Corresponding SMF field" indicates the name of the SMF field that corresponds to the field name in the annotation.

Table 39. Annotated fields for SMF 120 data

Field	Description	Corresponding SMF field
`Application`	The application name	SM1209EO
`ControllerJobname`	The job name for the controller	SM1209BT
`DeleteServiceCPUActiveCount`	The count of samples when the enclave delete CPU service time was non-zero. Time is accumulated by the enclave as reported by the **CPUSERVICE** parameter of the IWM4EDEL API. A value of `0` indicates that the enclave was not deleted.	SM1209DN count
`DispatchCPU`	The amount of CPU time, in microseconds, that is used by dispatch TCB.	SM1209CI
`EnclaveCPU`	The amount of CPU time that was used by the enclave as reported by the **CPUTIME** parameter of the IWM4EDEL API.	SM1209DH
`EnclaveServiceDeleteCPU`	The enclave delete CPU service that is accumulated by the enclave as reported by the **CPUSERVICE** parameter of the IWM4EDEL API. A value of `0` indicates that the enclave was not deleted.	SM1209DN
`RecordType`	Internal record type. The following values are possible: • `SMF120_REQAPPL`, which indicates a WebSphere application record • `SMF120_REQCONT`, which indicates a WebSphere controller record	Set by the data provider
`RequestCount`	Request count	Set by the data provider
`RequestEnclaveCPU`	The enclave CPU time at the end of the dispatch of this request, as reported by the **CPUTIME** parameter of the IWMEQTME API. The units are in TOD format.	SM1209DA
`RequestTime`	The time that the request was received, or the time that the WebSphere application or controller completed processing of the request response.	SM1209CM, SM1209CQ

Table 39. Annotated fields for SMF 120 data (continued)

Field	Description	Corresponding SMF field
RequestType	The type of request that was processed. The following values are possible: • HTTP • HTTPS • IIOP • INTERNAL • MBEAN • MDB-A • MDB-B • MDB-C • NOTKNOWN • OTS • SIP • SIPS • UNKNOWN	SM1209CK
SpecialtyCPU	The amount of CPU time that was spent on non-standard CPs, such as the z Systems Application Assist Processor (zAAP) and z Systems Integrated Information Processor (zIIP). This value is obtained from the TIMEUSED API.	SM1209CX
SpecialtyCPUActiveCount	The count of samples when the amount of CPU time that was spent on non-standard CPs, such as the zAAP and zIIP, was non-zero. The CPU utilization value is obtained from the TIMEUSED API.	SM1209CX count
SystemID	The system identifier	SM120SID
zAAPCPUActiveCount	The count of samples when the delete zAAP CPU enclave time was non-zero. A value of 0 indicates that the enclave was not deleted or not normalized. This CPU time is obtained from the ZAAPTIME field in the IWM4EDEL macro.	SM1209DI count
zAAPEligibleCPU	The amount of CPU time at the end of the dispatch of this request that is spent on a regular CP that could have been run on a zAAP, but the zAAP was not available. This value is obtained from the ZAAPONCPTIME field in the IWMEQTME macro.	SM1209DC
zAAPEnclaveCPUNormalized	The enclave zAAP CPU time at the end of the dispatch of this request, as reported by the **ZAAPTIME** parameter of the IWMEQTME API. This utilization is adjusted by the zAAP normalization factor at the end of the dispatch of this request. The normalization factor is obtained from the **ZAAPNFACTOR** parameter of the IWMEQTME API.	SM1209DG, SM1209DB

Table 39. Annotated fields for SMF 120 data (continued)

Field	Description	Corresponding SMF field
`zAAPEnclaveDeleteCPU`	The delete zAAP CPU enclave. A value of `0` indicates that the enclave was not deleted or not normalized. This value is obtained from the ZAAPTIME field in the IWM4EDEL macro. This value is normalized by the enclave delete zAAP normalization factor as reported by the **`ZAAPNFACTOR`** parameter of the IWM4EDEL API.	SM1209DJ, SM1209DI
`zAAPEnclaveServiceDeleteCPU`	The enclave delete zAAP Service that is accumulated by the enclave as reported by the **`ZAAPSERVICE`** parameter of the IWM4EDEL API. A value of `0` indicates that the enclave was not deleted.	SM1209DM
`zAAPServiceCPUActiveCount`	The count of samples when the enclave delete zAAP service time was non-zero. Time is accumulated by the enclave as reported by the **`ZAAPSERVICE`** parameter of the IWM4EDEL API. A value of `0` indicates that the enclave was not deleted.	SM1209DM count
`zIIPCPUActiveCount`	The count of samples when the enclave delete zIIP time was non-zero. Time is accumulated by the enclave as reported by the **`ZIIPTIME`** parameter of the IWM4EDEL API. A value of `0` indicates that the enclave was not deleted.	SM1209DK count
`zIIPEligibleCPUEnclave`	The eligible zIIP enclave that is on the CPU at the end of the dispatch of this request. This value is obtained from the ZIIPTIME field in the IWMEQTME macro.	SM1209DF
`zIIPEnclaveCPU`	The zIIP enclave that is on the CPU at the end of the dispatch of this request. This value is obtained from the ZIIPONCPTIME field in the IWMEQTME macro.	SM1209DD
`zIIPEnclaveDeleteCPU`	The enclave delete zIIP time that is accumulated by the enclave as reported by the **`ZIIPTIME`** parameter of the IWM4EDEL API. A value of `0` indicates that the enclave was not deleted.	SM1209DK
`zIIPEnclaveQualityCPU`	The zIIP Quality Time enclave that was on the CPU at the end of the dispatch of this request. This value is obtained from the ZIIPQUALTIME field in the IWMEQTME macro.	SM1209DE
`zIIPEnclaveServiceDeleteCPU`	The enclave delete zIIP service that is accumulated by the enclave as reported by the **`ZIIPSERVICE`** parameter of the IWM4EDEL API. A value of `0` indicates that the enclave was not deleted or not normalized.	SM1209DL

Table 39. Annotated fields for SMF 120 data (continued)

Field	Description	Corresponding SMF field
`zIIPServiceCPUActiveCount`	The count of samples when the enclave delete zIIP service time was non-zero. Time is accumulated by the enclave as reported by the **ZIIPSERVICE** parameter of the IWM4EDEL API. A value of `0` indicates that the enclave was not deleted or not normalized.	SM1209DL count

SYSLOG data

z/OS system log (z/OS SYSLOG) data can originate either from the z/OS user exits or the operations log (OPERLOG).

Tip: For the IBM Operations Analytics - Log Analysis platform only, you can also gather z/OS SYSLOG data from a static print file in System Display and Search Facility (SDSF) format. The data source type for this log data is `zOS-SYSLOG-SDSF`. This data is rendered by SDSF and can be ingested in batch mode by using the IBM Operations Analytics - Log Analysis Data Collector client.

Data stream definition for SYSLOG data

Table 40 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 40. Data stream definition for SYSLOG data

Type of node in the policy	Required configuration value
Data Stream	One of the following values: • **z/OS SYSLOG** (for data from the z/OS user exits) • **z/OS SYSLOG from OPERLOG** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **z/OS** > **Logs**, and select the check box for the respective data stream. **Tip:** You cannot define both a **z/OS SYSLOG** and a **z/OS SYSLOG from OPERLOG** data stream in the same policy.
Transcribe Transform	`UTF-8`
Splitter Transform	On the IBM Operations Analytics - Log Analysis platform, the Splitter Transform is not applicable because all data must be sent as unsplit. On the Elastic Stack and Splunk platforms, use the following values: • For a **z/OS SYSLOG** data stream, the transform value is `SYSLOG Splitter`. • For a **z/OS SYSLOG from OPERLOG** data stream, this transform is not applicable.
Filter Transform	Not required

Table 40. Data stream definition for SYSLOG data (continued)

Type of node in the policy	Required configuration value
Subscriber	See "Subscribers for each type of source data" on page 140.

syslogd data

Syslogd data is network data from the UNIX System Services system log (`syslogd`). The abbreviation syslogd represents the term *syslog daemon*.

Data stream definition for syslogd data

Table 41 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 41. Data stream definition for syslogd data

Type of node in the policy	Required configuration value
Data Stream	One or more of the following values: • **USS Syslogd Admin** • **USS Syslogd Debug** • **USS Syslogd Error** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **Network** > **UNIX System Services**, and select the check box for the respective data stream.
Splitter Transform	On the IBM Operations Analytics - Log Analysis platform, the Splitter Transform is not applicable because all data must be sent as unsplit. On the Elastic Stack and Splunk platforms, use `SyslogD Splitter`.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 140.

WebSphere HPEL data

WebSphere Application Server for z/OS High Performance Extensible Logging (HPEL) data is log data from an HPEL repository.

Data stream definition for HPEL data

Tip: This data stream can be defined only for the IBM Operations Analytics - Log Analysis platform.

Table 42 on page 136 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 42. Data stream definition for HPEL data

Type of node in the policy	Required configuration value
Data Stream	**WebSphere HPEL** **To select this data stream in the Configuration Tool:** In the “Select data stream” window, click **IBM Z Operations Analytics** > **Web Servers** > **WebSphere Application Server**, and select the **WebSphere HPEL** check box. In the “Configure Log Forwarder data stream” window for this data stream, for **File Path**, use the value `/u/WebSphere/V8R5/bbocell/bbonode/AppServer/profiles/default/logs/Server`.
Transcribe Transform	`UTF-8`
Splitter Transform	Not applicable
Filter Transform	Not required
Subscriber	See “Subscribers for each type of source data” on page 140.

WebSphere SYSOUT data

WebSphere Application Server for z/OS SYSOUT data is from the SYSOUT job log.

Data stream definition for SYSOUT data

Table 43 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 43. Data stream definition for SYSOUT data

Type of node in the policy	Required configuration value
Data Stream	**WebSphere SYSOUT** **To select this data stream in the Configuration Tool:** In the “Select data stream” window, click **IBM Z Operations Analytics** > **Web Servers** > **WebSphere Application Server**, and select the **WebSphere SYSOUT** check box.
Transcribe Transform	`UTF-8`
Splitter Transform	On the IBM Operations Analytics - Log Analysis platform, the Splitter Transform is not applicable because all data must be sent as unsplit. On the Elastic Stack and Splunk platforms, use `WAS for zOS SYSOUT Splitter`.
Filter Transform	Not required
Subscriber	See “Subscribers for each type of source data” on page 140.

WebSphere SYSPRINT data

WebSphere Application Server for z/OS SYSPRINT data is from the SYSPRINT job log.

Data stream definition for SYSPRINT data

Table 44 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 44. Data stream definition for SYSPRINT data

Type of node in the policy	Required configuration value
Data Stream	One or more of the following values: • **WebSphere SYSPRINT** • **WebSphere USS Sysprint** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **Web Servers** > **WebSphere Application Server**, and select the check box for the respective data stream.
Transcribe Transform	`UTF-8`
Splitter Transform	On the IBM Operations Analytics - Log Analysis platform, the Splitter Transform is not applicable because all data must be sent as unsplit. On the Elastic Stack and Splunk platforms, use `WAS for zOS SYSPRINT Splitter`.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 140.

zAware interval anomaly data

zAware interval anomaly data is applicable only on the IBM Operations Analytics - Log Analysis platform. The IBM zAware data gatherer, a component of IBM Z Operations Analytics on the Log Analysis platform, gathers this data from IBM z Advanced Workload Analysis Reporter (IBM zAware) and sends it to IBM Z Operations Analytics.

zAware interval anomaly data is provided as a z/OS SYSLOG data source of type `zOS-Anomaly-Interval`.

Annotated fields for zAware interval anomaly data

Table 45. Annotated fields for zAware interval anomaly data

Field	Description	Data type
`IntervalAnomaly`	A double value that indicates the anomaly score for the interval. The score is the percentile of the sum of each anomaly score for individual message IDs within the interval.	Double

Table 45. Annotated fields for zAware interval anomaly data (continued)

Field	Description	Data type
`IntervalEndTime`	The time, based on Coordinated Universal Time (UTC), that indicates the end of an interval for which the log messages that are produced are used to generate the anomaly record. The format is `YYYY-MM-DDTHH:mm:ss.sssZ`.	Date
`IntervalIndex`	An integer that indicates the sequence number of this interval within the specified date. Each index represents a 10-minute period.	Long
`IntervalStartTime`	The time, based on UTC, that indicates the start of an interval for which log messages that are produced are used to generate the anomaly record. The format is `YYYY-MM-DDTHH:mm:ss.sssZ`.	Date
`LimitedModelStatus`	An indication of whether the model that is used to calculate the anomaly score for this interval is a limited model. The following values are valid: • `YES` • `NO` • `UNKNOWN`	Text
`ModelGroupName`	The name of an analysis group. Each analysis group is associated with one or more systems from which the logs are used to create a single model.	Text
`NumMessagesNeverSeenBefore`	An integer that indicates the number of message IDs that were issued during this analysis interval for the first time but were never seen in any previous analysis interval or in the current model.	Long
`NumMessagesNotInModelFirstReported`	An integer that indicates the number of message IDs that are not in the model and were issued during this analysis interval for the first time.	Long
`NumMessagesUnique`	An integer that indicates the number of unique message IDs that were issued during this analysis interval.	Long
`SysplexName`	The sysplex name	Text
`SystemName`	The system name	Text

Table 45. Annotated fields for zAware interval anomaly data (continued)

Field	Description	Data type
`timestamp`	The time, based on UTC, that indicates the end of the interval record. This time is equivalent to the value for the `IntervalEndTime` field. When you search for interval anomaly scores that are based on a time stamp, ensure that you search for the end time of the interval record. The format is `YYYY-MM-DDTHH:mm:ss.sssZ`.	Date
`zAwareServer`	The hostname or IP address of the IBM z Advanced Workload Analysis Reporter (IBM zAware) server from which the interval anomaly data is retrieved.	Text

zSecure data

zSecure data is data from the Access Monitor component of IBM Security zSecure Admin. This data includes information about security events.

- “Data generation”
- “Data stream definition for zSecure data”

Data generation

The Access Monitor component of IBM Security zSecure Admin generates security events that the IBM Common Data Provider for z Systems sends to IBM Z Operations Analytics. These events include the following data:

- Successful and unsuccessful attempts to log on to applications
- Successful and unsuccessful attempts to access system resources, such as data sets and the z/OS file system (zFS)
- Successful and unsuccessful commands that are issued

The Access Monitor generates data transfer files on the UNIX System Services file system. For IBM Z Operations Analytics to use the Access Monitor data, IBM Common Data Provider for z Systems must be configured to read these data transfer files from the hierarchical file system (HFS) or the zFS, and send the file to IBM Z Operations Analytics by using the generic zFS file type.

Data stream definition for zSecure data

Table 46 on page 140 indicates the configuration values to use in defining this data stream in the IBM Common Data Provider for z Systems Configuration Tool.

Table 46. Data stream definition for zSecure data

Type of node in the policy	Required configuration value
Data Stream	**zSecure Access Monitor** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **Security** > **zSecure**, and select the **zSecure Access Monitor** check box. In the "Configure Log Forwarder data stream" window for this data stream, for **File Path**, use the zFS directory that contains the data transfer files for the Access Monitor component of IBM Security zSecure Admin.
Transcribe Transform	`UTF-8`
Splitter Transform	On the IBM Operations Analytics - Log Analysis platform, the Splitter Transform is not applicable because all data must be sent as unsplit. On the Elastic Stack and Splunk platforms, use `CRLF Splitter`.
Filter Transform	Not applicable
Subscriber	See "Subscribers for each type of source data."

Subscribers for each type of source data

In the policy that you define IBM Common Data Provider for z Systems Configuration Tool, each data stream must have a subscriber. The subscriber values vary depending on your IBM Z Operations Analytics platform. This reference lists the subscriber values that you must use in the policy, based on your platform.

IBM Operations Analytics - Log Analysis platform

In the "Add subscriber" or "Configure subscriber" window, use the following values:

- For **Protocol**, use either `CDP Logstash` or `CDP Logstash SSL`.
- For **Send As**, use `Unsplit`.

Elastic Stack platform

In the "Add subscriber" or "Configure subscriber" window, use the following values:

- For **Protocol**, use either `CDP Logstash` or `CDP Logstash SSL`.
- For **Send As**, use `Split`.

Splunk platform

In the "Add subscriber" or "Configure subscriber" window, use one of the following values for **Protocol**:

- `CDP Data Receiver`
- `CDP Data Receiver SSL`

Dashboards that represent the operational data

For each IBM Z Operations Analytics platform, IBM Z Operations Analytics provides dashboards in the user interface to help you troubleshoot problems in your IT operations environment. Each main troubleshooting dashboard contains visual representations of the data that is generated by the associated dashboard application. This reference lists the available dashboards for each platform.

The content of the dashboards can vary depending on the platform.

IBM Operations Analytics - Log Analysis platform

The following dashboard applications are provided in the z/OS Insight Packs. These dashboard applications also contain “Information links” dashboards, which link to troubleshooting information in the respective software documentation, including message explanations.

- WebSphere Application Server for z/OS dashboard applications, which include dashboards that represent data from WebSphere Application Server for z/OS
- z/OS Network dashboard applications, which include dashboards that represent data from the following software:
  - NetView for z/OS
  - TCP/IP
  - UNIX System Services system log (`syslogd`)
- z/OS SMF dashboard applications, which include dashboards that represent SMF data from the following software:
  - CICS Transaction Server for z/OS
  - Db2 for z/OS
  - IMS for z/OS
  - MQ for z/OS
  - Security for z/OS
  - WebSphere Application Server for z/OS
- z/OS SYSLOG dashboard applications, which include dashboards that represent data from the following software:
  - z/OS SYSLOG
  - CICS Transaction Server for z/OS
  - Db2 for z/OS
  - IMS for z/OS
  - MQ for z/OS
  - Security for z/OS

Elastic Stack platform

The following dashboards are provided for the Elastic Stack platform.

- CICS Transaction Server for z/OS Enterprise Dashboard by Region
- CICS Transaction Server for z/OS Enterprise Dashboard by System
- CICS Transaction Server for z/OS System Dashboard
- CICS Transaction Server for z/OS Region Dashboard
- CICS Transaction Server for z/OS Transaction Dashboard
- CICS Transaction Server for z/OS Job Dashboard

- Db2 for z/OS Enterprise Dashboard by Subsystem
- Db2 for z/OS Enterprise Dashboard by System
- Db2 for z/OS System Dashboard
- Db2 for z/OS Subsystem Dashboard
- Db2 for z/OS Job Dashboard
- IMS for z/OS Job Dashboard
- MQ for z/OS Job Dashboard
- Saved Searches Dashboard
- Systems Dashboard
- Welcome Dashboard
- z/OS Job Dashboard
- z/OS Security Server RACF Dashboard
- zSecure Access Monitor Dashboard

Splunk platform

The following dashboards are provided for the Splunk platform.

- CICS Transaction Server for z/OS Enterprise Dashboard by Region
- CICS Transaction Server for z/OS Enterprise Dashboard by System
- CICS Transaction Server for z/OS System Dashboard
- CICS Transaction Server for z/OS Region Dashboard
- CICS Transaction Server for z/OS Transaction Dashboard
- CICS Transaction Server for z/OS Job Dashboard
- Db2 for z/OS Enterprise Dashboard by Subsystem
- Db2 for z/OS Enterprise Dashboard by System
- Db2 for z/OS System Dashboard
- Db2 for z/OS Subsystem Dashboard
- Db2 for z/OS Job Dashboard
- IMS for z/OS Job Dashboard
- MQ for z/OS Job Dashboard
- Saved Searches Dashboard
- Systems Dashboard
- Welcome Dashboard
- z/OS Job Dashboard
- z/OS Security Server RACF Dashboard
- zSecure Access Monitor Dashboard

Searches that are predefined for searching the operational data

IBM Z Operations Analytics provides predefined searches (sometimes also known as *sample searches*, *saved searches*, or *Quick Search samples*) that can be accessed from the user interface to search operational data. This reference lists and describes the available searches.

These searches include queries of key annotated fields for z/OS systems and subsystems. These fields can contain important information that contributes to the operational insights.

CICS Transaction Server for z/OS searches

The name for each CICS Transaction Server for z/OS sample search is shown with a brief description of what the associated query looks for.

CICS Transaction Server Abend or Severe Messages

Searches for CICS Transaction Server messages that have the format `DFH`*`ccxxxx`*, where *cc* represents a component identifier (such as `SM` for Storage Manager), and *xxxx* is either `0001` or `0002` (which indicates an abend or severe error in the specified component).

For example: This sample would search for `DFHSM0001` but not for `DFH0001`.

CICS Action, Decision, or Error Messages

Searches for CICS messages that indicate any of the following situations:

- Immediate action is required.
- A decision is required.
- An error occurred.

The search is based on the CICS message IDs and on an action code of `A`, `D`, `E`, `S`, or `U`.

CICS Transaction Server Key Messages

Searches for a set of predefined message numbers to determine whether any of the messages occurred.

CICS Transaction Server Messages

Searches for CICS Transaction Server messages, which start with the prefix `DFH` or `EYU`.

CICS Transaction Server Short on Storage Messages

Searches for CICS Transaction Server for z/OS messages that indicate that a storage shortage occurred.

CICS Transaction Server Start Stop Messages

Searches for CICS Transaction Server for z/OS messages that are written to the system log while the CICS Transaction Server for z/OS is started or stopped. Messages with the following numbers are examples:

- `EYUXL0010I`
- `DFHPA1101`

CICS Transaction Server Storage Violations

Searches for CICS Transaction Server for z/OS messages that indicate that a storage violation occurred.

List of CICS Transaction Server for z/OS searches that are based on System Management Facilities (SMF) data

To obtain results from the following searches, CICS Transaction Server for z/OS must be active and configured to create SMF 110 records. For more information, see "SMF 110 data generation" on page 123.

CICS Job Performance

Searches for records that have a program name of `DFHSIP` or `EYU9XECS`.

CICS Transaction Server Exceptions

Searches for CICS Transaction Server for z/OS exceptions that occurred.

CICS Transaction Server Policy Exceptions
Searches for CICS Transaction Server for z/OS SMF policy-based exceptions that occurred.

CICS Transaction Server Summary
Searches for CICS Transaction Server for z/OS transaction summary interval records that occurred.

CICS Transaction Server Summary End-of-Day
Searches for CICS Transaction Server for z/OS end-of-day transaction summary records that occurred.

CICS Transaction Server Task Limit Met
Searches for CICS Transaction Server for z/OS transaction records where the number of active user transactions equaled the specified maximum allowed number of user transactions.

CICS Transaction Server Wait on Storage Exceptions
Searches for CICS storage manager messages and CICS Transaction Server for z/OS SMF `Wait on Storage` exceptions.

Db2 for z/OS searches

The name for each Db2 for z/OS sample search is shown with a brief description of what the associated query looks for.

Db2 Action, Decision, or Error Messages
Searches for Db2 messages that indicate any of the following situations:

- Immediate action is required.
- A decision is required.
- An error occurred.

Db2 Data Set Messages
Searches for Db2 messages that indicate any of the following situations:

- Failure of a data set definition
- Failure of a data set extend
- Impending space shortage

Db2 Data Sharing Messages
Searches for internal resource lock manager (IRLM) messages that were issued to Db2 and that indicate at least one of the following situations:

- The percentage of available lock structure capacity is low.
- An error occurred when IRLM used the specified z/OS automatic restart manager (ARM) function.

Db2 Job Performance
Searches for records that have a program name of `DSNYASCP` or `DSNADMT0`.

Db2 Lock Conflict Messages
Searches for Db2 messages that indicate that a plan was denied an IRLM lock due to a detected deadlock or timeout.

Db2 Log Data Set Messages
Searches for messages that indicate that Db2 log data sets are full, are becoming full, or could not be allocated.

Db2 Log Frequency Messages
Searches for Db2 messages that indicate that log archives were offloaded or are waiting to be offloaded.

Db2 Messages
Searches for Db2 messages, which start with the prefix DSN.

Db2 Pool Shortage Messages
Searches for Db2 messages that indicate that the amount of storage in the group buffer pool (GBP) coupling facility structure that is available for writing new pages is low or critically low.

IMS for z/OS searches

The name for each IMS for z/OS sample search is shown with a brief description of what the associated query looks for.

IMS Abend Messages
Searches for messages that indicate abends were detected.

IMS Action, Decision, or Error Messages
Searches for IMS messages that indicate any of the following situations:

- Immediate action is required.
- A decision is required.
- An error occurred.

The search is based on the IMS message IDs and on an action code of A, E, W, or X.

IMS Common Queue Server Messages
Searches for IMS Common Queue Server component messages, which start with the prefix CQS.

IMS Connect Messages
Searches for IMS Connect component messages, which start with the prefix HWS.

IMS Database Recovery Control Errors
Searches for Database Recovery Control component error messages, which start with the prefix DSP.

IMS Job Performance
Searches for records that have a program name of DFSAMVRC0, DFSRRC00, or DXRRLM00.

IMS Locking Messages
Searches for messages that indicate which IMS resources are locked.

IMS Log Messages
Searches for messages that indicate how often IMS logs are rolled.

IMS Messages
Searches for IMS messages, which start with any of the following prefixes:

BPE, CQS, CSL, DFS, DSP, DXR, ELX, FRP, HWS, MDA, PCB, PGE, SEG, or SFL

IMS Pool Issues
Searches for messages that indicate IMS pool-related issues.

IMS Resources in Waiting Errors
Searches for error messages that indicate a resource is waiting on other resources to become available.

IMS Security Violations
Searches for error messages that indicate security violations were detected.

IMS Stopped Resources
Searches for messages that indicate IMS and related components are no longer running.

IMS Terminal Related Messages
Searches for messages that indicate IMS terminal-related issues, including terminals that are no longer receiving messages.

MQ for z/OS searches

The name for each MQ for z/OS sample search is shown with a brief description of what the associated query looks for.

MQ Action, Decision, or Error Messages
Searches for MQ messages that indicate any of the following situations:

- Immediate action is required.
- A decision is required.
- An error occurred.

The search is based on the MQ message IDs and on an action code of `A`, `D`, or `E` .

MQ Buffer Pool Errors
Searches for error messages that indicate the occurrence of MQ buffer pool errors.

MQ Channel Errors
Searches for error messages that indicate the occurrence of MQ channel errors.

MQ Channel Initiator Errors
Searches for error messages that indicate the occurrence of MQ channel initiator errors.

MQ Interesting Informational Messages
Searches for a set of predefined informational message numbers to determine whether any of the corresponding messages occurred.

MQ Job Performance
Searches for records that have a program name of `CSQXJST` or `CSQYASCP`.

MQ Key Messages
Searches for a set of predefined message numbers to determine whether any of the corresponding messages occurred.

MQ Logs Start and Stop Messages
Searches for messages that are related to the starting, stopping, and flushing of the MQ log data sets.

MQ Messages
Searches for MQ messages, which start with the prefix `CSQ`.

MQ Queue Manager Storage Messages
Searches for messages that indicate whether MQ queue manager required more storage.

MQ Start Stop Messages
Searches for messages that are written to the system log while the MQ queue manager or channel initiator is started or stopped. Messages with the following numbers are examples:

- `CSQY000I`
- `CSQY001I`

NetView for z/OS searches

The name for each NetView for z/OS sample search is shown with a brief description of what the associated query looks for.

NetView Action, Decision, or Error Messages
Searches for NetView for z/OS messages that indicate any of the following situations:

- Immediate action is required.
- A decision is required.
- An error occurred.

NetView Automation
Searches for a set of predefined NetView for z/OS messages that indicate possible automation table violations.

NetView Command Authorization
Searches for a set of predefined NetView for z/OS messages that indicate possible command authorization table violations.

NetView Messages
Searches for NetView for z/OS messages.

NetView Resource Limits
Searches for a set of predefined NetView for z/OS messages that indicate that resource limits or storage thresholds might have been exceeded.

NetView Security Messages
Searches for a set of predefined NetView for z/OS messages that indicate insufficient access authority or security environment violations.

Security searches: RACF

The name for each Resource Access Control Facility (RACF) sample search is shown with a brief description of what the associated query looks for.

Security RACF Action, Decision, or Error Messages
Searches for RACF messages that indicate any of the following situations:

- Immediate action is required.
- A decision is required.
- An error occurred.

Security RACF Insufficient Access Messages
Searches for RACF messages that indicate insufficient access authority.

Security RACF Insufficient Authority Messages
Searches for RACF messages that indicate insufficient authority.

Security RACF Invalid Logon Attempt Messages
Searches for RACF messages that indicate invalid logon attempts.

Security RACF Messages
Searches for RACF messages, which start with either of the following prefixes:

- ICH
- IRR

List of RACF searches that are based on System Management Facilities (SMF) data
To obtain results from the following searches, RACF must be active and protecting the resources or commands that are the subject of each search:

Security RACF Accesses of Configuration Files
Searches for any accesses of files with the extension `.config`.

Security RACF Activity for Operations
Searches for any events that were caused by a user with the RACF `OPERATIONS` attribute.

Security RACF CHOWN, CHGRP, CHMOD Commands
Searches for occurrences of the UNIX commands CHOWN, CHGRP, and CHMOD that were issued.

Security RACF Data Set Access Successes
Searches for successful attempts to access data sets.

Security RACF Failed Access Attempts
Searches for unsuccessful attempts to access data sets.

Security RACF Logons and Commands
Searches for logons and commands that were issued from a specific terminal ID (`TermID` field). The default value for the `TermID` field is non-blank.

Security RACF SETROPTS Commands Issued
Searches for SETROPTS commands that were issued.

Security searches: zSecure Access Monitor

The name for each sample search for the Access Monitor component of IBM Security zSecure Admin is shown with a brief description of what the associated query looks for.

zSecure Access Monitor All Records
Searches for all records that are created by the Access Monitor.

zSecure Access Monitor Authorization Nonzero Result
Searches for records with the following characteristics:

- Are based on the RACF `AUTH` definition
- Have a non-zero return code
- Are created by the Access Monitor

zSecure Access Monitor Authorization Records
Searches for records with the following characteristics:

- Are based on the RACF `AUTH` definition
- Are created by the Access Monitor

zSecure Access Monitor CICS Authorization Nonzero Result
Searches for CICS transaction-related records with a non-zero return code that are created by the Access Monitor.

zSecure Access Monitor CICS Transactions
Searches for all CICS transaction-related records that are created by the Access Monitor.

zSecure Access Monitor Command Nonzero Result
Searches for records with the following characteristics:

- Are based on the use of the RACF **`DEFINE`** command to add or remove a profile in the RACF database
- Have a non-zero return code
- Are created by the Access Monitor

zSecure Access Monitor Command Records

Searches for records with the following characteristics:

- Are based on the use of the RACF **DEFINE** command to add or remove a profile in the RACF database
- Are created by the Access Monitor

zSecure Access Monitor Define Nonzero Result

Searches for records with the following characteristics:

- Are based on the RACF `DEFINE` definition
- Have a non-zero return code
- Are created by the Access Monitor

zSecure Access Monitor Define Records

Searches for records with the following characteristics:

- Are based on the RACF `DEFINE` definition
- Are created by the Access Monitor

zSecure Access Monitor Fast Nonzero Result

Searches for records with the following characteristics:

- Are based on the RACF `FASTAUTH` definition
- Have a non-zero return code
- Are created by the Access Monitor

zSecure Access Monitor Fast Records

Searches for records with the following characteristics:

- Are based on the RACF `FASTAUTH` definition
- Are created by the Access Monitor

zSecure Access Monitor Verify Nonzero Result

Searches for records with the following characteristics:

- Are based on the RACF `VERIFY` definition
- Have a non-zero return code
- Are created by the Access Monitor

zSecure Access Monitor Verify Records

Searches for records with the following characteristics:

- Are based on the RACF `VERIFY` definition
- Are created by the Access Monitor

WebSphere Application Server for z/OS searches

The name for each WebSphere Application Server for z/OS sample search is shown with a brief description of what the associated query looks for.

WebSphere Error Messages

Searches for WebSphere Application Server for z/OS messages that indicate an error.

WebSphere Exceptions

Searches for occurrences of Java exceptions in the WebSphere Application Logs.

List of WebSphere Application Server for z/OS searches that are based on System Management Facilities (SMF) data

To obtain results from the following searches, WebSphere Application Server for z/OS must be active and configured to create SMF 120 subtype 9 records:

WebSphere Activity for All Applications
Searches for the requests for processing that are attributed to WebSphere Application Server for z/OS applications.

WebSphere Applications with Nonzero Dispatch TCB
Searches for the requests for processing that are attributed to WebSphere Application Server for z/OS applications with nonzero dispatch Task Control Block (TCB) time.

WebSphere Controller Managed JavaBeans
Searches for the managed JavaBeans requests that are processed by the WebSphere Application Server Controller.

WebSphere Controller Requests Non-Internal
Searches for the requests for controller processing that are not attributed to internal WebSphere processing.

z/OS network searches

The name for each z/OS network sample search is shown. These samples look for common network errors.

Searches for common network errors

The following z/OS network sample searches are provided:

- Network ATTLS Error Messages
- Network CSSMTP Error Messages
- Network Device Error Messages
- Network FTP Error Messages
- Network IKED Error Messages
- Network IPSEC Error Messages
- Network OMPROUTE Error Messages
- Network PAGENT Error Messages
- Network Storage Error Messages
- Network syslogd FTPD Messages
- Network syslogd Messages
- Network syslogd SSHD Messages
- Network syslogd TELNETD Messages
- Network TCPIP Error Messages
- Network TN3270 Telnet Error Messages
- Network VTAM® Connection Error Messages
- Network VTAM CSM Error Messages
- Network VTAM Storage Error Messages

z/OS system searches

The name for each sample search of the z/OS system is shown with a brief description of what the associated query looks for.

Job Performance
Searches for records that have an assigned program name.

Troubleshooting Z Operations Analytics

Depending on the IBM Z Operations Analytics platform that you are using, this reference lists known problems that you might experience in using IBM Z Operations Analytics and describes known solutions. The IBM Common Data Provider for z Systems V1.1.0 documentation also contains troubleshooting information that might be helpful.

About this task

Some of the known problems apply to multiple IBM Z Operations Analytics platforms and are listed under *Problems that can occur on multiple platforms*. The known problems that apply to a specific platform are listed by platform.

Troubleshooting Z Operations Analytics problems that apply to multiple platforms

This reference lists known problems that you might experience in using IBM Z Operations Analytics on at least two of the three platforms.

APPLID values for CICS Transaction Server might not be correct in the user interface

This problem applies to all platforms. For CICS Transaction Server for z/OS, the application identifier (APPLID) values might not be correct in the IBM Z Operations Analytics user interface.

Symptom

APPLID values for CICS Transaction Server for z/OS are expected to be provided in the user interface. However, if the APPLID value is not present in the CICS Transaction Server for z/OS message text, the first word of the message text is incorrectly used as the APPLID.

Cause

CICS Transaction Server for z/OS typically includes the APPLID as the first word of the message. However, when CICS Transaction Server for z/OS messages do not include the APPLID as the first word in the message, IBM Z Operations Analytics incorrectly assumes that the first word of the message is an APPLID.

Solution

No workaround is available.

Db2 or MQ command prefix values might not be correct in the user interface

This problem applies to all platforms. For Db2 for z/OS and MQ for z/OS, the command prefix values might not be correct in the IBM Z Operations Analytics user interface.

Symptom

Command prefix values for Db2 for z/OS and MQ for z/OS are expected to be provided in the user interface. However, if the command prefix value is not present in the Db2 for z/OS or MQ for z/OS message text, the first word of the message text is incorrectly used as the command prefix.

Cause

Db2 for z/OS and MQ for z/OS typically include the command prefix as the first word of the message. However, when Db2 for z/OS or MQ for z/OS messages do not include the command prefix as the first word in the message, IBM Z Operations Analytics incorrectly assumes that the first word of the message is a command prefix.

Solution

No workaround is available.

Duplicate entries are shown for SMF data streams in the Configuration Tool

This problem applies to all platforms. In the IBM Common Data Provider for z Systems Configuration Tool, you see duplicate entries for SMF data streams under the **IBM Z Operations Analytics** category (such as two entries for **SMF30** or two entries for **SMF80**).

Cause

The IBM Common Data Provider for z Systems Configuration Tool is not using the appropriate configuration file for SMF data streams that are destined for IBM Z Operations Analytics.

For each IBM Z Operations Analytics platform, IBM Z Operations Analytics provides configuration files for the IBM Common Data Provider for z Systems Configuration Tool.

Depending on your environment, you must use only *one* of the following configuration files:

If you are using only the IBM Operations Analytics - Log Analysis platform
Use the `glasmf.streams.json` file.

If you are using the Elastic Stack or Splunk platform
Use the `glaELKSplunk.streams.json` file.

If you are using the IBM Operations Analytics - Log Analysis platform *and* **any of the other platforms**
Use the `glaELKSplunk.streams.json` file.

Solution

Verify that only one of the IBM Z Operations Analytics configuration files is being used by the IBM Common Data Provider for z Systems Configuration Tool.

For more information, see “Enabling support for SMF data destined for IBM Z Operations Analytics” on page 3.

During the import of a Message Library, an invalid character is found

This problem applies only to the Elastic Stack and Splunk platforms. In the IBM Z Operations Analytics Problem Insights Framework, when you import a Message Library, an invalid character is found.

Symptom

The following message is shown:

```
[javax.xml.stream.XMLStreamException:
An invalid XML character (Unicode: 0xffffffff) was found
in the element content of the document.]
```

Cause

During the parsing of the Message Library XML file, the XML parser detected an invalid character that it cannot process. The cause might be the presence of a binary character or the use of an apostrophe rather than a single quotation mark.

Solution

To determine the cause, load the XML file in a browser (for example, in Google Chrome), which can show the line and column of the invalid character.

Troubleshooting Z Operations Analytics on the Log Analysis platform

This reference lists known problems that you might experience in using IBM Z Operations Analytics on the IBM Operations Analytics - Log Analysis platform. The IBM Operations Analytics - Log Analysis V1.3.5 documentation also contains troubleshooting information that might be helpful.

Log files

Troubleshooting information is available in log files from IBM Common Data Provider for z Systems, IBM Operations Analytics - Log Analysis, Logstash, and the `ioaz` Logstash output plugin.

IBM Common Data Provider for z Systems log files
The following log files provide information about the IBM Common Data Provider for z Systems:

Data Streamer log files
Logging information (and trace information, if trace is enabled) is sent to the STDOUT data set on the HBODS001 job. Some Data Streamer messages are also written to the console.

Log Forwarder log files
Logging information (and trace information, if trace is enabled) is sent to the STDERR data set on the GLAPROC job. Significant Log Forwarder messages are also written to the console.

System Data Engine log files
Logging information is sent to the HBOOUT data set on the HBOSMF job. The following information is also included if you specify in the logging configuration that this information must be sent:

- A copy of the input stream is included with the information that is sent to the HBOOUT data set on the HBOSMF job.
- A report about the records that are processed for each processing interval is sent to the HBODUMP data set on the HBOSMF job.

Log Analysis log files

The following log files, which are located on the Log Analysis server, provide information about the processing of z/OS log data:

LA_INSTALL_DIR/**logs/GenericReceiver.log**

Contains information about the ingestion of log records and Insight Pack processing.

LA_INSTALL_DIR/**logs/UnityApplication.log**

Contains information about searches.

Logstash log file

The log file is *LOGSTASH_INSTALL_DIR*/logs/logstash-ioaz.log.

For information about using the following options when you start Logstash, see the Logstash documentation:

- --debug
- --debug-config

ioaz Logstash output plugin log file

The log file is ioaz-logstash.*n*.log, where *n* is a number from 0 to 19, and where 0 indicates the most recent file, and 19 indicates the oldest file. The log file location is specified by the user at installation time. The location is the value of the log_path option in the Logstash configuration file.

The Logstash configuration file also contains a log_level option that you can use to gather more information.

Enabling tracing for the ioaz Logstash output plugin

For the ioaz Logstash output plugin, you can enable tracing by using the log_level option in the Logstash configuration file.

About this task

The Logstash configuration file is *LOGSTASH_INSTALL_DIR*/config/logstash-ioaz.conf. The value of the log_level option in the Logstash configuration file is applied each time that Logstash is started.

Procedure

To enable tracing for the ioaz Logstash output plugin, complete the following steps:

1. Edit the Logstash configuration file, and change the value of the log_level option to debug or trace. A value of debug results in limited additional information, and a value of trace results in more detailed information.
2. Restart the ioaz Logstash output plugin by using the *LOGSTASH_INSTALL_DIR*/bin/logstash_util.sh script.

What to do next

When you no longer need the trace settings, change the value of the log_level option to info (the default value).

Log record skipped and not available in Log Analysis

Individual log records are not retained by the IBM Operations Analytics - Log Analysis server.

Symptom

If the Log Analysis server is shut down, the server might not retain the last transmitted log record for each data source.

Cause

When the Log Analysis server shuts down, the log records that the server is processing might be lost because the server does not cache in-process log records.

Solution

No workaround is available.

Search results do not include the expected z/OS data

In the IBM Operations Analytics - Log Analysis user interface, search results do not include the expected z/OS data.

If data that is issued in a z/OS logical partition (LPAR) is not shown in the IBM Operations Analytics - Log Analysis user interface, the following steps can help you determine possible causes.

Step 1: Verify that IBM Common Data Provider for z Systems is running

Verify that IBM Common Data Provider for z Systems is running. If it is running, determine whether it logged any error or warning messages.

For more information about any error or warning messages that you find, see the IBM Common Data Provider for z Systems V1.1.0 documentation.

Step 2: Check the `logstash-ioaz.log` and `ioaz-logstash.`*`n`*`.log` files for error or warning messages

If no IBM Common Data Provider for z Systems error or warning messages are logged, review the *`LOGSTASH_INSTALL_DIR`*`/logs/logstash-ioaz.log` and *`LOGSTASH_INSTALL_DIR`*`/logs/ioaz-logstash.`*`n`*`.log` files on the Logstash server, where *n* represents a number from 0 to 19, with 0 indicating the most recent file and 19 indicating the oldest file.

Look for messages with severity `ERROR` or `WARN`.

Step 3: If no Logstash or `ioaz` Logstash output plugin error or warning messages are logged, check the `GenericReceiver.log` file for error or warning messages

If no Logstash or `ioaz` Logstash output plugin error or warning messages are logged, review the *`LA_INSTALL_DIR`*`/logs/GenericReceiver.log` file on the IBM Operations Analytics - Log Analysis server.

Look for messages with severity `ERROR` or `WARN`.

For more information about any error or warning messages that you find, see Troubleshooting in the Log Analysis documentation.

Step 4: Check the `GenericReceiver.log` file for ingestion status

The *LA_INSTALL_DIR*/logs/GenericReceiver.log file also specifies the number of log records that are processed in each batch of log data, as shown in the following example:

```
03/06/14 15:58:25:660 EST [Default Executor-thread-4] INFO  - DataCollectorRestServlet :
Batch of Size 9 processed and encountered 0 failures
```

In this example, nine log records were successfully ingested, and no log records failed to be ingested.

You can determine the data source name by looking for a data source ingestion record before the batch status record in the GenericReceiver.log. In the following example, the data source name is my_sysprint:

```
03/06/14 15:58:25:650 EST [Default Executor-thread-4] INFO  - UnityFlowController :
Adding data source under ingestion, logsource = my_sysprint threadId = 96
```

Determine whether the status information in the GenericReceiver.log indicates one of the following three situations:

All batch status records for a data source have a size value of 0

Log data is being ingested into IBM Operations Analytics - Log Analysis, but one of the following issues is preventing any log records from being identified:

- The data source type that is specified for the data source that is configured in IBM Operations Analytics - Log Analysis is incorrect.

  Examples:

  - The data source type is zOS-CICS-MSGUSR, but the log data is being sent from an EYULOG data set. Therefore, the data source type must be specified as zOS-CICS-EYULOG.
  - The data source type is zOS-WAS-SYSPRINT, and the log data is being sent from a SYSPRINT data set, but the SYSPRINT data set contains log data in the distributed format. Therefore, the data source type must be specified as WASSystemOut.
- Although the data source type is correct, the log data does not contain any records in the format that is expected by the data source type.

  Examples:

  - The SYSPRINT data contains only WebSphere Application Server internal trace records. It does not contain records that are produced by the Java RAS component. For SYSPRINT, only records that are produced by the Java RAS component are ingested.
  - The SYSPRINT data contains only unstructured data, such as data that is produced by System.out.println() calls from Java code. For SYSPRINT, only records that are produced by the Java RAS component are ingested.
  - The SYSOUT data contains Java garbage collector trace data. Java garbage collector data is not supported by the zOS-WAS-SYSOUT data source type.

Batch status records are present for the data source, the batch size is nonzero, and the number of failures is zero

Log records are being ingested successfully. Determine whether you have one of the following situations:

- The search criteria is incorrect or not broad enough to include the expected log data.

  Example: If the search time filter is set from 2:00 PM until 2:05 PM, a record that was logged at 2:05:01 PM is not flagged.

  Verify the search filters. You might want to start with a broad search, and refine it as needed.
- The time zone information is incorrectly configured in the IBM Common Data Provider for z Systems. This issue can cause the time stamps of the ingested records to be 1 - 12 hours earlier or later than they should be.
- If only the most recent log record is missing, the record might be held in buffer by IBM Operations Analytics - Log Analysis. IBM Operations Analytics - Log Analysis cannot confirm whether a log record is complete until the next log record is received for that data source. Therefore, for each data source, the last log record that is sent to IBM Operations Analytics - Log Analysis is typically held (and not ingested) until the next log record is received.

No batch status records exist for the data source

Determine whether you have one of the following situations:

- IBM Common Data Provider for z Systems and Logstash are started, but no log records are generated for the data source. No log records are ingested because no log data exists to ingest.
- IBM Common Data Provider for z Systems does not have the appropriate access to files or directories.

  Example: For example, the user ID might not have read access to one of the following items:
  - A z/OS UNIX log file
  - The High Performance Extensible Logging (HPEL) log or trace directory.

  Change the file permissions to give IBM Common Data Provider for z Systems the appropriate access.
- Because the IBM Common Data Provider for z Systems is incorrectly configured, it cannot find the log data.

After upgrade, interval anomaly data is not visible in user interface

After you upgrade to IBM Z Operations Analytics Version 3.2.0, you do not see any information about interval anomalies on the Problem Insights page of the Log Analysis user interface. For example, when you look at the data for a sysplex, you do not see the new Interval Score column in the table, and you do not see a bar chart for each system within the selected sysplex.

Cause

The most probable cause is that the IBM zAware data gatherer is not configured. If the data gather is configured, the most probable cause is that the browser cache needs to be cleared.

Solution

Complete the following steps:

1. Verify that the IBM zAware data gatherer is configured.
2. Complete one of the following steps:
   - If the data gatherer is ***not*** configured, configure it.
   - If the data gatherer ***is*** configured, clear the browser cache.

Search error after installing Problem Insights extension

When you search in the IBM Operations Analytics - Log Analysis interface for the first time after you install the IBM Z Operations Analytics Problem Insights extension, the message CTGLA2005E indicates that an unexpected error occurred.

Solution

No action is required. This problem occurs only during the initial search after you install the Problem Insights extension, and it resolves quickly.

Data cache for Problem Insights extension is corrupted

If the data cache for the Problem Insights extension is corrupted, you can use a utility to reset the cache.

Solution

To reset the data cache, run the following command from the directory *LA_INSTALL_DIR*/utilities/cacheUtility:

```
./cacheUtility.sh -reset
```

For help information, run the following command:

```
./cacheUtility.sh -help
```

Impact of SSL certificate verification changes in Python 2.7.9

Secure Sockets Layer (SSL) certificate verification changes in Python 2.7.9 or later can impact IBM Z Operations Analytics dashboards and the IBM zAware data gatherer. Python Enhancement Proposal (PEP) 476 changes the default behavior for HTTPS certificate verification in Python clients.

Symptom

For Python clients where PEP 476 is applied, the verification of self-signed certificates is usually unsuccessful, which prevents search results from showing in the web browser. The following message is shown:

```
CTGLA0630E : Application execution failed due to unknown
error. An error occurred while executing GET /CSRFToken.
```

Before the application of PEP 476, Python clients that were using HTTPS did not present errors if the verification of self-signed certificates was unsuccessful.

For more information, see the following sources:

- From the Python Software Foundation: PEP 476 -- Enabling certificate verification by default for stdlib http clients
- For Red Hat Enterprise Linux: Certificate verification in Python standard library HTTP clients

Solution for use of the IBM zAware data gatherer

The IBM zAware data gatherer establishes HTTPS sessions with both the IBM zAware and IBM Operations Analytics - Log Analysis servers. By default, the data gatherer does not present an error if the verification of self-signed certificates is unsuccessful.

The environment variable *PYTHONHTTPSVERIFY* controls certificate verification. Before you run the `zAwareDataGatherer.py` script to enable certificate verification in the IBM zAware data gatherer, complete the following steps:

1. Set the value of *PYTHONHTTPSVERIFY* to `1`, which specifies that, if certificate verification is unsuccessful, an error message is recorded in the log file, and the `zAwareDataGatherer.py` script ends.

   If the value of *PYTHONHTTPSVERIFY* is not set, or is set to `0` (the default value), certificate verification is disabled for both the IBM zAware and IBM Operations Analytics - Log Analysis servers.
2. Add the IBM zAware and IBM Operations Analytics - Log Analysis certificates to the Python certificate store.

Solution for use of the IBM Z Operations Analytics dashboards

To simplify the base configuration of IBM Z Operations Analytics, SSL certificate verification is disabled.

If you want to enable SSL certificate verification for the IBM Z Operations Analytics dashboards, you can purchase an SSL certificate from a certificate authority (CA), and deploy it to the IBM Operations Analytics - Log Analysis keystore.

To enable certificate verification for the dashboards, complete the following steps:

1. In the IBM Operations Analytics - Log Analysis keystore, install the SSL certificate that you purchased from the CA.

   For more information, see Configuring CA certificates for SSL in the Log Analysis documentation.
2. In the Python script `CommonAppMod.py` that is in each of the following four directories, set the value of the environment variable *PYTHONHTTPSVERIFY* to 1.
   - `LA_INSTALL_DIR/AppFramework/Apps/WASforzOSInsightPack_v3.2.0.0/CommonAppMod.py`
   - `LA_INSTALL_DIR/AppFramework/Apps/zOSNetworkInsightPack_v3.2.0.0/CommonAppMod.py`
   - `LA_INSTALL_DIR/AppFramework/Apps/SMFforzOSInsightPack_v3.2.0.0/CommonAppMod.py`
   - `LA_INSTALL_DIR/AppFramework/Apps/SYSLOGforzOSInsightPack_v3.2.0.0/CommonAppMod.py`

   This example shows how this value must be set:

   ```
   os.environ["PYTHONHTTPSVERIFY"] = "1"
   ```

Troubleshooting Z Operations Analytics on the Elastic Stack platform

This reference lists known problems that you might experience in using IBM Z Operations Analytics on the Elastic Stack platform.

Search results do not include the expected z/OS data

In Kibana, the search results do not include the expected z/OS data.

If data that is issued in a z/OS logical partition (LPAR) is not shown in Kibana, the following steps can help you determine possible causes.

Step 1: Verify that Logstash is running

Run the following commands to verify that Logstash is using the default port that is specified in the `B_cdpz.conf` file:

On Linux systems

```
netstat -an | grep 8080
```

If the type of your Logstash image is `.deb` or `.rpm`, you can also use the following command:

```
service logstash status
```

On Windows systems

```
netstat -an | find "8080"
```

Step 2: Verify that Elasticsearch has no errors

Check the `elasticsearch.log` file for any errors that indicate problems with the network connection, data ingestion, incorrect mapping, or an incorrect template.

On Linux systems

If the type of your Logstash image is `.tar.gz` or `.zip`, run the following command to open the `elasticsearch.log` file:

```
cat YOUR_EXTRACTION_PATH/logs/elasticsearch.log
```

If the type of your Logstash image is `.deb` or `.rpm`, run the following command to open the `elasticsearch.log` file:

```
cat /var/logs/elasticsearch/elasticsearch.log
```

On Windows systems

Use Notepad to open the *YOUR_EXTRACTION_PATH*`/logs/elasticsearch.log` file.

Step 3: Verify that data is being received from IBM Common Data Provider for z Systems

Verify that data is being received by Logstash and written to Elasticsearch.

On Linux systems

To verify that data is being received by Logstash and written to Elasticsearch, run the following command:

```
curl ELASTICSEARCH_HOST/IP>:9200/_cat/indices
```

On Windows systems

To verify that data is being received by Logstash and written to Elasticsearch, open the following URL in a web browser:

```
http://ELASTICSEARCH_HOST/IP>:9200/_cat/indices
```

The output includes index names. Check for indices that start with `zoa` and end with a recent time stamp.

Complete the following steps, depending on whether data is being written to Elasticsearch:

If data is not being written to Elasticsearch
Check for Logstash error logs.

If data is being written to Elasticsearch

1. Verify that IBM Z Operations Analytics is installed on the Elastic Stack platform.
2. In Kibana, verify that you are searching within a valid time range.
3. Verify that your index pattern is `zoa-*` and that this pattern is created automatically.
4. In Kibana, verify that no warnings are indicated.

No correlation between sysplexes and systems in user interface filters for Problem Insights dashboard

In the IBM Z Operations Analytics Problem Insights dashboard in Kibana, the user interface filters might not be correctly correlated. For example, after you select a sysplex filter to restrict data to a specific sysplex, you want to select a system filter for a system in that sysplex. However, when you start to select a system filter, you see systems that are not part of that sysplex in the selection list.

Cause

In Elastic Stack 6.1, the filter functions in Kibana cannot recognize relationships among filters. Therefore, your choice for the system filter is unrelated to your choice for the sysplex filter.

Solution

In later versions of Elastic Stack, filter relationships are enabled in Kibana. Upgrade to Elastic Stack 6.3.

Troubleshooting Z Operations Analytics on the Splunk platform

This reference lists known problems that you might experience in using IBM Z Operations Analytics on the Splunk platform.

Search results do not include the expected z/OS data

In the Splunk user interface, the search results do not include the expected z/OS data.

If data that is issued in a z/OS logical partition (LPAR) is not shown in the Splunk user interface, the following steps can help you determine possible causes.

Step 1: Verify that the IBM Common Data Provider for z Systems Data Receiver is running

To verify that the Data Receiver is running, log on to your Data Receiver system, and run the following commands with the port that is specified in the `cdpdr.properties` file:

On Linux systems

```
netstat -an | grep 8989
```

On Windows systems

```
netstat -an | find "8989"
```

Step 2: Verify that the *CDPDR_PATH* environment variable is set

The path that is specified by the *CDPDR_PATH* environment variable must be available to the Data Receiver and to the Splunk service that is ingesting data.

Step 3: Verify that data is being received from IBM Common Data Provider for z Systems

Verify that data is being received by the Data Receiver and written to disk. To view the data files that are written to disk, run the following command. The most recent files are shown at the bottom of the list.

On Linux systems

```
ls -alrt $CDPDR_PATH
```

On Windows systems

```
dir /od %CDPDR_PATH%
```

Complete the following steps, depending on whether data is being written to disk:

If data is not being written to disk

To troubleshoot this problem, see Subscriber is not receiving data in the IBM Common Data Provider for z Systems V1.1.0 documentation.

If data is being written to disk

1. Verify that both the IBM Common Data Provider for z Systems Buffered Splunk Ingestion App and IBM Z Operations Analytics are installed on the Splunk platform.

   Restriction: If you installed IBM Z Operations Analytics Version 3.2.1 or later, you must install Version 1.1.3 of the IBM Common Data Provider for z Systems Buffered Splunk Ingestion App.
2. Verify that the IBM Common Data Provider for z Systems Buffered Splunk Ingestion App is enabled in Splunk.
3. Verify that you are searching the appropriate indexes.

   If you are running with the default index names, run the following search:

   ```
   index=zos*
   ```

   Otherwise, you might want to run the following search:

   ```
   index=* sourcetype=zOS*
   ```
4. Verify that the correct data streams were defined in the policy so that the correct data is sent.

   In the IBM Common Data Provider for z Systems Configuration Tool, the IBM Z Operations Analytics data streams that you can define are under the category **IBM Z Operations Analytics**. The IBM Z Operations Analytics dashboards and searches are based on the data from only these data streams.

When you open the Problem Insights dashboard, a message indicates no server connection is found

When you open the IBM Z Operations Analytics Problem Insights dashboard in the Splunk user interface, a message indicates that no server connection is found, and no data is shown in the tables.

Cause

The Splunk user interface cannot communicate with the IBM Z Operations Analytics Problem Insights Framework. The Problem Insights Framework might not be active, or a problem might exist in the communication between the user interface and the Problem Insights Framework.

Solution

1. Verify that your configuration of the Problem Insights Framework is correct.
2. Verify that the Problem Insights Framework is active.
3. Verify that the Problem Insights Framework is accessible from the Splunk Enterprise system. For example, complete the following steps:
   a. Verify that you can ping the Problem Insights Framework server from the Splunk Enterprise system.
   b. Verify that no firewalls are blocking communication between the Problem Insights Framework server and the Splunk Enterprise system.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Printed in USA