IBM Z Operations Analytics
4.1

User Guide

Figures

Tables

Contents

Chapter 1. Z Operations Analytics overview

IBM® Z Operations Analytics provides two primary capabilities: 1) the capability to search, visualize, and analyze large amounts of structured and unstructured operational data across Z systems IT environments, and 2) the capability to identify and understand patterns of unusual activity in Z systems IT environments, and provide actionable insight into problems.

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.

- "Primary capabilities of Z Operations Analytics" on page 1
- "Flow of data among the primary components" on page 2
- "Component descriptions" on page 2
- "Other key concepts" on page 3

Primary capabilities of Z Operations Analytics

Capability to search, visualize, and analyze large amounts of structured and unstructured operational data across Z systems IT environments

This capability helps you more quickly identify workload issues and discover previously unknown problems in operational data, and perform root cause analysis. Operational data includes log, event, and service request data and performance metrics.

This capability is available on the following analytics platforms, but the GUI and available functions can vary depending on the analytics platform:

- IBM Operations Analytics - Log Analysis
- Elastic Stack
- Splunk

Capability to identify and understand patterns of unusual activity in Z systems IT environments, and provide actionable insight into problems

This capability can consolidate evidence of unusual activity and provide actionable recommendations to correct the condition. With the help of advanced machine learning technology, it can also identify abnormal behavior and highlight behavioral trends to help prevent system and application outages.

This capability includes its own GUI, which is called the *Problem Insights GUI*.

Table 1 on page 1 correlates each primary product capability with its associated use case and with the typical end user roles that the capability supports.

For each use case, Chapter 2, "Road map for using Z Operations Analytics," on page 5 summarizes the setup tasks (planning, installation, and configuration tasks) that the administrator must do to set up the analytics environment.

Table 1. Mapping of product capability to associated use case and end user role

Use case	End user role	Product capability
Review dashboards, or use predefined searches, to understand the status of systems and subsystems.	• Application owner • Application developer	Capability to search, visualize, and analyze large amounts of structured and unstructured operational data across Z systems IT environments

Table 1. Mapping of product capability to associated use case and end user role (continued)

Use case	End user role	Product capability
Use Problem Insights to gain actionable insights into current problems.	- Z operator - Subject matter expert (SME)	Capability to identify and understand patterns of unusual activity in Z systems IT environments, and provide actionable insight into problems

Flow of data among the primary components

Figure 1 on page 2 illustrates the flow of data among IBM Z Operations Analytics components. The following topics include more detailed illustrations and explanations of these components:

- Flow of data among metric-based machine learning components
- Flow of source data among Z Operations Analytics components on the Elastic Stack and Splunk platforms
- Flow of source data among Z Operations Analytics components on the Log Analysis platform

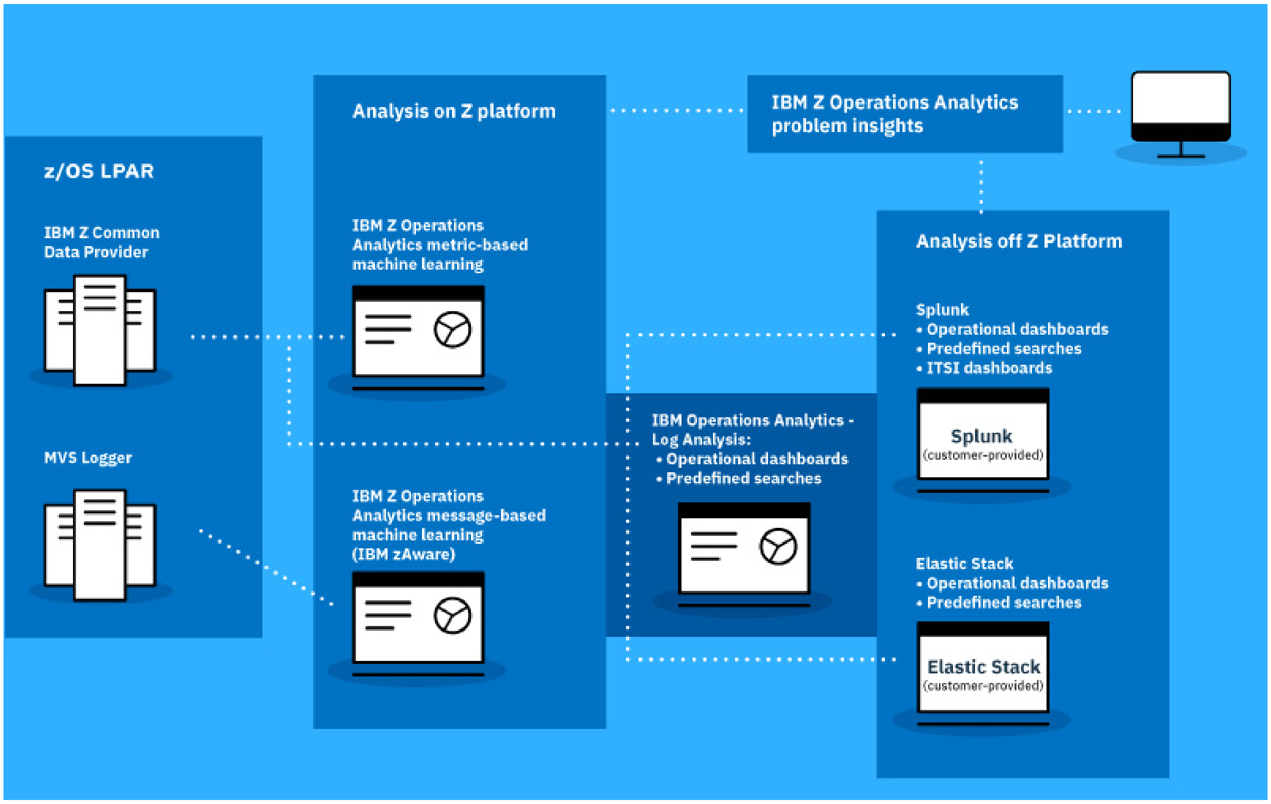

z/OS LPAR
IBM Z Common Data Provider
MVS Logger
Analysis on Z platform
IBM Z Operations Analytics metric-based machine learning
IBM Z Operations Analytics message-based machine learning (IBM zAware)
IBM Z Operations Analytics problem insights
IBM Operations Analytics - Log Analysis:
• Operational dashboards
• Predefined searches
Analysis off Z Platform
Splunk
• Operational dashboards
• Predefined searches
• ITSI dashboards
Splunk (customer-provided)
Elastic Stack
• Operational dashboards
• Predefined searches
Elastic Stack (customer-provided)

Figure 1. Flow of data among IBM Z Operations Analytics components

Component descriptions

The following definitions provide more information about some key components of IBM Z Operations Analytics. Also, see "Other key concepts" on page 3 for explanations of other key concepts in IBM Z Operations Analytics.

Insight Pack

In the IBM Z Operations Analytics machine learning system, a grouping of problem insights and supporting artifacts that provide the following capabilities:

- Enable data collection and analysis

- Enable the presentation of resulting insights in the IBM Z Operations Analytics Problem Insights GUI
- Provide subsystem-specific machine learning capabilities

machine learning system
The IBM Z Operations Analytics software that provides the machine learning capability. It is based on IBM Watson Machine Learning for z/OS.

machine learning system CLI
The command-line interface (CLI) that is used to manage the workflow of the IBM Z Operations Analytics machine learning system.

message broker
The component that facilitates communication between the scoring agents of the IBM Z Operations Analytics machine learning system and the rules engine of the IBM Z Operations Analytics Problem Insights server.

Problem Insights CLI
The command-line interface (CLI) for the IBM Z Operations Analytics Problem Insights server.

Problem Insights GUI
The graphical user interface (GUI) for the IBM Z Operations Analytics Problem Insights server. Users can visualize and administer problem insights in this GUI.

Problem Insights server
The IBM Z Operations Analytics software that provides the capability to process, visualize, and administer problem insights.

rules engine
The component that applies analytics logic (called *analysis routines*) to incoming machine learning results data, based on configurable criteria such as event notification rules.

scoring agent
A long-running process that is used by the IBM Z Operations Analytics machine learning system in continuous scoring mode to score the incoming operational data. The scoring agent analyzes newly collected SMF data at time intervals that are configurable. When you start the scoring process, a unique scoring agent is launched for each subsystem type (for example Db2® or CICS®).

IBM WebSphere® Application Server for z/OS Liberty
The embedded application server in IBM Z Operations Analytics.

Other key concepts

These terms represent key concepts in IBM Z Operations Analytics.

analysis model
A model of normal system behavior that is computed by the IBM Z Operations Analytics machine learning system, based on analysis of historical system data. The machine learning system uses this model as a baseline for detecting anomalies in your Z environment.

enterprise data warehouse
For an instance of the IBM Z Operations Analytics machine learning system, the repository for operational data that is being sent to the analytics platform and for the data that results from the analysis of this operational data.

event management system
Software that collects relevant events (such as the completion or failure of an operation, a user action, or the change in state of a process) from its associated computer systems and manages the processing of these events. An event management system typically includes an event gathering system, an event database, monitoring agents, and an event reporting system that can provide event information in multiple formats.

You can configure the IBM Z Operations Analytics Problem Insights server to send events to the IBM Tivoli® Netcool®/OMNIbus event management system.

problem insight

Near real-time insight that is provided for a defined set of potential problems in your IT environment, with suggested actions to help resolve the problem.

Chapter 2. Road map for using Z Operations Analytics

Depending on your use cases for IBM Z Operations Analytics, the planning, installation, and configuration tasks that the administrator must do to set up the analytics environment vary. This road map summarizes the setup tasks for each use case.

Determine which use case is relevant to your organization

Determine which of the following use cases is relevant to your organization:

- Review dashboards, or use predefined searches, to understand the status of systems and subsystems.
- Use problem insights to gain actionable insights into current problems.

Install and configure IBM Z Common Data Provider

First, install and configure IBM Z Common Data Provider. This configuration includes defining the data streams for streaming the operational data to IBM Z Operations Analytics.

For more information, see Chapter 3, "Z Common Data Provider: Planning for installation and configuration," on page 7.

Based on your use case, complete other setup tasks

Table 2 on page 5 summarizes other planning, installation, and configuration tasks that the administrator must do to set up the analytics environment, depending on the use case.

Remember: Before you start these other tasks, install and configure IBM Z Common Data Provider.

Table 2. Summary of planning, installation, and configuration tasks that the administrator must do to set up the analytics environment, depending on the use case

Use case	Administrator tasks for setup
Review dashboards, or use predefined searches, to understand the status of systems and subsystems.	Configure the analytics platform to set up the IBM Z Operations Analytics dashboards and predefined searches for the application owner and application developer. See the following information, depending on your platform: • #unique_7 • "Z Log and Data Analytics on the Elastic Stack and Splunk platforms" on page 105
Use problem insights to gain actionable insights into current problems.	Complete the following tasks to set up the IBM Z Operations Analytics problem insights for the Z operator and SME: 1. Install and configure the IBM Z Operations Analytics Problem Insights server. See Chapter 4, "Deploying the Problem Insights server," on page 15. 2. If you want to use the problem insights that are based on z/OS SYSLOG messages, also configure integration of the Problem Insights server with either the Elastic Stack or Splunk analytics platform. See "Z Log and Data Analytics on the Elastic Stack and Splunk platforms" on page 105. 3. If you want to use machine learning to gain insight into behavioral trends and patterns of normal and anomalous activity, also install and configure the IBM Z Operations Analytics machine learning system. See Chapter 5, "Deploying the machine learning system," on page 47.

Chapter 3. Z Common Data Provider: Planning for installation and configuration

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

About this task

IBM Z Operations Analytics includes IBM Z Common Data Provider V2.1.0.

IBM Z Common Data Provider 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 information about the IBM Z Common Data Provider system requirements, see Planning to use IBM Z® Common Data Provider in the IBM Z Common Data Provider documentation.

Preparing the Configuration Tool to support SMF record types for Z Operations Analytics

If you plan to send System Management Facilities (SMF) data to IBM Z Operations Analytics, you must prepare the IBM Z Common Data Provider Configuration Tool to support the SMF record types that are destined for IBM Z Operations Analytics. You must complete this task before you create any policies that define SMF data streams with IBM Z Operations Analytics as the target destination.

About this task

For more information about the definition of SMF data streams, see "Definition of your SMF data streams in the Configuration Tool" on page 10.

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 Z Common Data Provider 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 Z Common Data Provider 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/zoa/V4R1M0/zcdp/streams/glaELKSplunk.streams.json
      config_tool_workdir
   ```

   Important: If you are using the file `glaELKSplunk.streams.json`, the file `glasmf.streams.json` must not be in the working directory for the IBM Z Common Data Provider Configuration Tool.

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

- SMF100_1
- SMF101_SUMMARY
- SMF110_1_SUMMARY

Table 3 on page 8 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 3. 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/zoa/V4R1M0/zcdp/streams/glasmf.streams.json
   config_tool_workdir
```

Important: If you are using the file glasmf.streams.json, the file glaELKSplunk.streams.json must not be in the working directory for the IBM Z Common Data Provider Configuration Tool.

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/zoa/V4R1M0/zcdp/streams/glaELKSplunk.streams.json
   config_tool_workdir
```

Important: If you are using the file glaELKSplunk.streams.json, the file glasmf.streams.json must not be in the working directory for the IBM Z Common Data Provider Configuration Tool.

If you are sending SMF data to the IBM Z Operations Analytics machine learning system

The following files must be in the working directory for the IBM Z Common Data Provider Configuration Tool:

`glaDb2_ML.streams.json`

This file includes required data for configuring IBM Z Operations Analytics machine learning system streams that contain IBM Db2 for z/OS data.

Tip: From UNIX System Services, use the following sample command, with the appropriate values for your environment, to copy the `glaDb2_ML.streams.json` file:

```
cp /usr/lpp/IBM/zoa/V4R1M0/zcdp/streams/glaDb2_ML.streams.json
   config_tool_workdir
```

`glaCICS_ML.streams.json`

This file includes required data for configuring IBM Z Operations Analytics machine learning system streams that contain IBM CICS Transaction Server for z/OS data.

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

```
cp /usr/lpp/IBM/zoa/V4R1M0/zcdp/streams/glaCICS_ML.streams.json
   config_tool_workdir
```

`izoa.protocols.json`

This file includes required data for configuring the IBM Z Operations Analytics enterprise data warehouse subscribers to receive data from the IBM Z Operations Analytics machine learning system.

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

```
cp /usr/lpp/IBM/zoa/V4R1M0/zcdp/protocols/izoa.protocols.json
   config_tool_workdir
```

2. Verify that the `concats.json` file is in the working directory for the IBM Z Common Data Provider 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 Z Common Data Provider. The value in quotation marks (in this example, `CDP.SHBODEFS`) must be the data set name for your installation.
`"IZOA" : "ZOA.V4R1M0.SGLADEFS"`	A reference to the SGLADEFS data set that is installed with IBM Z Operations Analytics. The value in quotation marks (in this example, `ZOA.V4R1M0.SGLADEFS`) must be the data set name for your installation.

Results

In the IBM Z Common Data Provider Configuration Tool, you can now create policies that define SMF data streams with IBM Z Operations Analytics as the target destination, as described in "Definition of your SMF data streams in the Configuration Tool" on page 10.

Definition of your SMF data streams in the Configuration Tool

In the IBM Z Common Data Provider Configuration Tool, different System Management Facilities (SMF) data streams are listed under different categories. For the SMF record types that are destined for IBM Z Operations Analytics, ensure that you define the SMF data streams from the appropriate categories.

Important: Before you create any policies that define these SMF data streams, you must prepare the Configuration Tool to support the associated SMF record types, as described in "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

SMF data stream categories

In the Configuration Tool, different SMF data streams are listed under the categories that are described in Table 4 on page 10.

The SMF data streams in each category are similar, but not the same. For example, the **IBM Z Common Data Provider** 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 two **IBM Z Operations Analytics** categories include 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.

Table 4. SMF data stream categories in IBM Z Common Data Provider Configuration Tool

Category in Configuration Tool	More information
IBM Z Operations Analytics	To use the dashboards and predefined searches that are provided by IBM Z Operations Analytics, configure the SMF data streams in this category.
IBM Z Operations Analytics - Machine Learning	To use the machine learning Insight Packs that are provided by IBM Z Operations Analytics, configure the SMF data streams in this category. For more information about these Insight Packs, see "Insight Packs" on page 20. **Important:** When you configure the IBM Z Operations Analytics machine learning system subscriber (as described in "Subscribers for each type of source data" on page 10), send only the SMF data streams that are in the **IBM Z Operations Analytics - Machine Learning** category. Do not send SMF data streams from the other categories to the machine learning system subscriber. **Restriction:** The SMF data streams in this category are not supported by the IBM Operations Analytics - Log Analysis, Elastic Stack, or Splunk platforms.
IBM Z Common Data Provider	IBM Z Operations Analytics does not use data from the SMF data streams in this category.

Subscribers for each type of source data

In the policy that you define in the IBM Z Common Data Provider Configuration Tool, each data stream must have a subscriber. The subscriber values vary depending on your IBM Z Operations Analytics environment, including your analytics platform. This reference lists the subscriber values that you must use in the policy, depending on your environment.

Subscriber values in the policy

When you configure a subscriber in a policy, use the following values, depending on your environment.

Important: If you use an SSL value for **Protocol** (to secure communications between the IBM Z Common Data Provider Data Streamer and the subscriber), you must also complete the relevant configuration steps that are described in Securing communications between the Data Streamer and its subscribers.

IBM Operations Analytics - Log Analysis platform

Use one of the following values for **Protocol**:

- `IZOA on IOA-LA via Logstash`
- `IZOA on IOA-LA via Logstash SSL`

Elastic Stack platform

Use one of the following values for **Protocol**:

- `IZOA on Elasticsearch via Logstash`
- `IZOA on Elasticsearch via Logstash SSL`

Splunk platform

Use one of the following values for **Protocol**:

- `IZOA on Splunk via Data Receiver`
- `IZOA on Splunk via Data Receiver SSL`
- `IZOA on Splunk via HEC via HTTP`
- `IZOA on Splunk via HEC via HTTPS`

IBM Z Operations Analytics machine learning system

To send data to the IBM Z Operations Analytics enterprise data warehouse for the machine learning system, use one of the following values for **Protocol**:

- `IZOA Enterprise Data Warehouse (Db2)`
- `IZOA Enterprise Data Warehouse (Db2) - SSL`

  If you use this SSL value, also complete the configuration steps that are described in More subscriber configuration information for the IBM Z Operations Analytics machine learning system.

For these protocols, you must provide the following configuration values:

Db2 host

The IP Address or URL for the system where the enterprise data warehouse is installed.

Db2 port

The number for the port that the enterprise data warehouse uses to accept external Distributed Relational Database Architecture™ (DRDA) connections.

Db2 subsystem

The IBM Db2 for z/OS subsystem that is configured with the appropriate tables for ingesting IBM Z Operations Analytics data.

Db2 schema

The schema that is used to configure the IBM Z Operations Analytics tables.

Db2 user ID

A RACF® user ID with permission to write to the enterprise data warehouse tables in the respective IBM Db2 for z/OS subsystem.

Db2 password

The RACF password that is associated with the Db2 user ID.

Send metadata

Select `true` to receive optional metadata fields. Otherwise, select `false`.

More subscriber configuration information for the IBM Z Operations Analytics machine learning system

For the machine learning system, if you use the value `ZOA Enterprise Data Warehouse (Db2) - SSL` for **Protocol** (to secure communications between the IBM Z Common Data Provider Data Streamer and the enterprise data warehouse subscriber), you must also complete the following subscriber configuration steps:

1. Complete the following steps on the z/OS host where the enterprise data warehouse is installed:

   a. Configure the IBM Communications Server Policy Agent, and create a policy to control incoming communications to the configured Db2 secure port. For more information, see the topic "Configuring DB2® for z/OS as a server with TLS/SSL support" in the IBM Redpaper DB2 for z/OS: Configuring TLS/SSL for Secure Client/Server Communications. The commands that you enter must use the same ID that is running the Db2 subsystem that receives the data.

   b. After the key ring is configured, use the following sample z/OS UNIX System Services command to export the certificate from the certificate authority (CA) to the working directory for the IBM Z Common Data Provider Data Streamer:

   ```
   cp "//'EXAMPLE.CERTS.DBCA'" /u/cdp/work/dbCA.cer
   ```

2. Transfer the public certificate to IBM Z Common Data Provider, as described in step 5 of Securing communications between the Data Streamer and its subscribers.
3. Complete the configuration steps that are described in steps 1, 2, and 6 of Securing communications between the Data Streamer and its subscribers.

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 Z Common Data Provider 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.

List of WebSphere Application Server for z/OS regions with focus area of each region

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.
   List of WebSphere Application Server for z/OS regions with typical job name of each region

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.
   List of WebSphere Application Server for z/OS regions with focus area of each region

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 Z Common Data Provider documentation.
5. Ensure that the user ID that is associated with the IBM Z Common Data Provider 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.

Chapter 4. Deploying the Problem Insights server

For the Elastic Stack and Splunk analytics platforms, and for metric-based machine learning, the IBM Z Operations Analytics Problem Insights server provides insights for a defined set of potential problems in your IT environment. For metric-based machine learning, the Problem Insights server also provides the capability to browse patterns and trends of anomalous behavior for monitored z/OS subsystems. The Problem Insights server can also be integrated with event management systems so that you can manage alerts for key events that are detected by IBM Z Operations Analytics.

Before you begin

Restriction: The Problem Insights server is not available on the IBM Operations Analytics - Log Analysis platform. A limited subset of the Problem Insights functionality is available by using the IBM Operations Analytics - Log Analysis Problem Insights extensions. For more information about these extensions, see the following topics:

- extensions for z/OS Problem Insights and client-side Expert Advice
- Installing the , extensions, and data gatherer

About this task

The following steps are the basic deployment steps:

1. Plan for deployment.

   See the following information:

   - “Planning for deployment of the Problem Insights server” on page 15
   - “System requirements for the Problem Insights server” on page 16
2. Install the Problem Insights server.

   See the following information, depending on the operating system:

   - “Installing the Problem Insights server on IBM z/OS UNIX System Services” on page 25
   - “Installing the Problem Insights server on a Linux system” on page 27
   - “Installing the Problem Insights server on a Windows system” on page 29
3. Configure the Problem Insights server.

   See “Configuring the Problem Insights server” on page 30.

Planning for deployment of the Problem Insights server

The IBM Z Operations Analytics Problem Insights server can be installed on IBM z/OS UNIX System Services, Linux, or Windows systems. However, the metric-based machine learning use cases are supported only when the Problem Insights server is installed on IBM z/OS UNIX System Services.

Before you begin

If you want to integrate the Problem Insights server with the IBM Z Operations Analytics machine learning system, install it on the z/OS system where the machine learning system and its prerequisite software are installed. For system requirements and other planning information for the machine learning system, see “Planning for deployment of the machine learning system” on page 49.

About this task

For more information about system requirements, see “System requirements for the Problem Insights server” on page 16.

When you install the Problem Insights server, the following directory structure is created:

runtime directory

The runtime directory contains the Problem Insights server. When the Problem Insight server is installed together with the machine learning system, this directory structure is identified by the environment variable *IZOA_HOME*.

Use the following guidelines in deciding which directory to use as the runtime directory:

- The directory must not be on a read-only file system.
- The directory must have at least 0.5 GB of storage space available.
- To avoid possible conflicts, do not use the SMP/E installation directory if you are installing the Problem Insights server on IBM z/OS UNIX System Services.

Important: With the exception of the Problem Insights server configuration files, do not update, delete, or move the files in the runtime directory.

For information about the Problem Insights server configuration files, see "Problem Insights server: configuration file reference" on page 85.

If you want to integrate the Problem Insights server with the IBM Z Operations Analytics machine learning system, install it on the z/OS system where the machine learning system and its prerequisite software are installed. For system requirements and other planning information for the machine learning system, see "Planning for deployment of the machine learning system" on page 49.

System requirements for the Problem Insights server

Ensure that your environment meets the system requirements for deploying the IBM Z Operations Analytics Problem Insights server.

- "Operating system requirements" on page 16
- "Prerequisite software" on page 16
- "Port requirements" on page 17
- "Data storage requirements on your analytics platform" on page 17
- "Memory requirements" on page 17
- "Screen resolution" on page 18

Operating system requirements

Table 5. Operating systems on which the Problem Insights server can run

Operating system	**Version**
IBM z/OS UNIX System Services	IBM z/OS 2.2 or 2.3, with APAR OA56570, or IBM z/OS 2.4
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 Z Systems	11 or 12
SUSE Linux Enterprise Server: for x86	11 or 12
Windows Server	2012, 2016, or 2019

Prerequisite software

Table 6 on page 17 indicates the software that must be operational for the Problem Insights server to run.

Table 6. Runtime requirements for the Problem Insights server	
Runtime requirement	**Version**
64-bit Java Runtime Environment (JRE)	8

Port requirements

The following ports must be available for the Problems Insights server:

Port for HTTPS communication with the Problem Insights server

The default port is 9446.

This value is specified in the following configuration files on the following configuration properties:

Configuration file	Configuration property
`analysisroutine.config`	`Application.serverPort`
`cli.config`	`port`

Ports for communication with the message broker

For the Apache ActiveMQ message broker, which is used by the rules engine, the following ports must be available:

Port for HTTPS communication with the message broker

The default port is 8162.

This value is specified in the following configuration files on the following configuration properties:

Configuration file	Configuration property
`amq.config`	`AMQ_HTTPS_PORT`
`izoaml.config`	`IZOA_RULES_PORT`

Port for secure OpenWire communication with the message broker

The default port is 61617.

This value is specified in the following configuration files on the following configuration properties:

Configuration file	Configuration property
`amq.config`	`AMQ_OPENWIRE_SSL_PORT`
`analysisroutine.config`	`izoa.activemq.broker.url`

Tip: For more information about these configuration files, see the following topics:

- “Problem Insights server: configuration file reference” on page 85
- “Machine learning system: configuration file reference” on page 103

Data storage requirements on your analytics platform

The Problem Insights server pulls data from the Elastic Stack platform, the Splunk platform, and IBM Db2 for z/OS. The results of data matches (between incoming operational data and problem insights content) are stored locally. Therefore, the system where the Problem Insights server is installed must have enough storage space to hold data for an extended period of time.

Memory requirements

For the Problem Insights server, the minimum memory requirement is 4 GB.

To run the Problem Insights server as a z/OS started task, plan for an additional 2 GB of memory.

Screen resolution

For the Problem Insights GUI to be rendered correctly, a minimum screen resolution of 1920 x 1200 is required. For example, the use of a lower screen resolution might cause truncation of the subsystem scorecard view.

Configuration planning for the Problem Insights server

During the installation of IBM Z Operations Analytics, the setup script `izoa-setup.sh` creates the configuration file `cli.config` in the *IZOA_HOME*`/config` directory and populates it with values for some configuration parameters. Depending on the choices that you make during installation, additional configuration files might also be created. All parameters in the configuration files must have a specified value, unless otherwise stated in a comment in the file. For some parameters, default values are specified, but you must verify that these values are appropriate for your environment, and change them if they are not.

After you verify the parameter values in the configuration files, you must run the `izoa-setup.sh` script to apply the configuration settings to the Problem Insights server.

Tip: The configuration files that are referenced in the following sections are also described in "Problem Insights server: configuration file reference" on page 85.

- "Configuration related to the machine learning system" on page 18
- "Configuration with other analytics platforms" on page 19
- "Configuration with an event management system or service, or with IBM Z ChatOps" on page 19

Configuration related to the machine learning system

As part of the configuration planning for the Problem Insights server, you must plan for configuration of the following components that are related to the IBM Z Operations Analytics machine learning system:

Enterprise data warehouse configuration

The Problem Insights server can communicate with only one IBM Db2 for z/OS enterprise data warehouse database at a time. During its configuration step, the `izoa-setup.sh` script reads the database connection information from the machine learning system configuration file `izoaml.config` (if this file is available), or it prompts you for the necessary input. The `izoa-setup.sh` script then automatically populates the following Problem Insights server configuration files with this information as needed:

- `analysisroutine.config`
- `db2zos.config`
- `microservices.config`

If you want to change the database configuration, you must rerun the configuration step of the `izoa-setup.sh` script.

For more information about the configuration of the enterprise data warehouse, see "Configuration planning for the machine learning system" on page 53.

Message broker configuration

The Problem Insights server uses Apache ActiveMQ as a message broker to facilitate communication between the scoring agents of the machine learning system and the rules engine of the Problem Insights server. The message broker is installed and configured with the Problem Insights server, but it must be started as a separate process.

The Problem Insights server configuration file for the message broker is `amq.config`.

Rules engine configuration

The Problem Insights server includes a rules engine that applies analytics logic (called *analysis routines*) to incoming machine learning results data, based on configurable criteria such as event notification rules. If the rules engine discovers an anomaly pattern in the data that matches the configured criteria, it notifies the Problem Insights server and an external event management system (if the event management system is configured).

The Problem Insights server configuration file for the rules engine is `analysisroutine.config`.

Configuration with other analytics platforms

You can configure the Problem Insights server to display actionable insights from operational data that is stored in Elastic Stack or Splunk. The Problem Insights server does not support both Elastic Stack and Splunk at the same time. It can, however, be configured to communicate concurrently with both the machine learning system and either Elastic Stack or Splunk.

If you want to integrate the Problem Insights server with the Elastic Stack or Splunk analytics platform, use the following deployment guidelines to minimize network latency:

Elastic Stack platform

Install the Problem Insights server on a system that is physically located near the system where the Elasticsearch component is installed.

The configuration file for this integration is `elk.config`.

Splunk platform

Install the Problem Insights server on a system that is physically located near the system where the Splunk search head component is installed.

The configuration file for this integration is `splunk.config`.

Configuration with an event management system or service, or with IBM Z ChatOps

The Problem Insights server can be integrated with a separately installed event management system, a cloud-based event management service, or IBM Z ChatOps. Event notifications can be forwarded for metric-based machine learning insights and for insights from the Elastic Stack or Splunk analytics platform. For more information, see “Sending events from the Problem Insights server to other applications” on page 34.

The following event management systems or services are supported:

- IBM Tivoli Netcool/OMNIbus

  If you want to forward metric-based machine learning events to IBM Tivoli Netcool/OMNIbus, complete the steps in “Sending events to IBM Tivoli Netcool/OMNIbus” on page 34.
- Cloud-based event management services:
  - IBM Cloud® Event Management
  - IBM Netcool Operations Insight®
  - IBM Watson AIOps Event Manager

  If you want to forward metric-based machine learning events to these cloud-based event management services, complete the steps in “Sending events to cloud-based event management services” on page 40.

IBM Z ChatOps (which represents "chat-based operations") is a collaboration model that connects people, processes, tools, and automation in a seamless and transparent way through a chat platform and extensive use of specialized chatbots. It can help IT operations teams improve service reliability, service recovery time, and collaboration efficiency.

The Problem Insights server can be integrated with IBM Z ChatOps. Events can be forwarded for metric-based machine learning insights and for insights from the Elastic Stack or Splunk analytics platform. For more information, see “Sending events to IBM Z ChatOps” on page 43.

The Problem Insights server configuration file for integration with any of these applications is `event.config`.

Insight Packs

IBM Z Operations Analytics includes Insight Packs for IBM CICS Transaction Server for z/OS and IBM Db2 for z/OS. Depending on the insights that you want, you must enable the appropriate Insight Pack in 1) the Problem Insights server so that it can provide visualizations of the subsystem-specific insights, and 2) the machine learning system so that it can provide subsystem-specific machine learning capabilities.

Key performance indicators (KPIs)

Each Insight Pack analyzes key performance indicators (KPIs) for the subsystem to which it applies. For information about these KPIs, see the following topics:

- "KPIs that are analyzed by the IBM CICS Transaction Server for z/OS Insight Pack" on page 20
- "KPIs that are analyzed by the IBM Db2 for z/OS Insight Pack" on page 22

Installation of Insights Packs

For information about installing or updating Insight Packs in the Problem Insights server, see "Installing or updating Insight Packs" on page 33.

KPIs that are analyzed by the IBM CICS Transaction Server for z/OS Insight Pack

This reference lists the key performance indicators (KPIs) that are analyzed by the IBM CICS Transaction Server for z/OS Insight Pack in IBM Z Operations Analytics. These KPIs are based on the System Management Facilities (SMF) record type 110 data is generated by that is generated by IBM CICS Transaction Server for z/OS.

For the IBM Z Operations Analytics machine learning system, SMF type 110 records must be generated in 1-minute intervals.

The KPIs are listed by the KPI group and KPI name that is provided by IBM Z Operations Analytics.

Table 7. KPIs that are analyzed by the CICS Insight Pack

KPI group	KPI name
CPU Time	• QRCPUT_CLOCK • MSCPUT_CLOCK • L8CPUT_CLOCK • J8CPUT_CLOCK • S8CPUT_CLOCK • ROCPUT_CLOCK • KY8CPU_CLOCK • KY9CPU_CLOCK • J9CPUT_CLOCK • L9CPUT_CLOCK • X8CPUT_CLOCK • X9CPUT_CLOCK • T8CPUT_CLOCK

Table 7. KPIs that are analyzed by the CICS Insight Pack (continued)

KPI group	KPI name
Dispatch Time	• QRDISPT_CLOCK • MSDISPT_CLOCK • RODISP_CLOCK • KY8DIS_CLOCK • KY9DIS_CLOCK
Response Time	• CPU_SEC • RESPONSE_SEC • DISPATCH_SEC • DISPATCH_WAIT_SEC • SUSPEND_SEC • FC_IO_WAIT_SEC • JC_IO_WAIT_SEC • EXCEPTION_WAIT_SEC
Storage	• STORAGE_UDSA_MAX • STORAGE_OCC_UDSA • STORAGE_OCC_EUDSA • STORAGE_PGM_TOTAL • STORAGE_PGM_B16M • STORAGE_TIOA • STORAGE_EUDSA_MAX • STORAGE_CDSA_MAX • STORAGE_ECDSA_MAX • STORAGE_OCC_CDSA • STORAGE_OCC_ECDSA • STORAGE_PGM_RDSA • STORAGE_PGM_ERDSA • STORAGE_PGM_SDSA • STORAGE_PGM_ESDSA • STORAGE_OCC_DSA • STORAGE_OCC_EDSA • STORAGE_PGM_UDSA
Transactions per second	• RECORDS • TRANSACTION_COUNT • TPS • DISPATCH_COUNT

KPIs that are analyzed by the IBM Db2 for z/OS Insight Pack

This reference lists the key performance indicators (KPIs) that are analyzed by the IBM Db2 for z/OS Insight Pack in IBM Z Operations Analytics. These KPIs are based on the System Management Facilities (SMF) record type 100 data that is generated by IBM Db2 for z/OS.

For the IBM Z Operations Analytics machine learning system, SMF type 100 records must be generated in 1-minute intervals.

The KPIs are listed by the KPI group and KPI name that is provided by IBM Z Operations Analytics.

Table 8. KPIs that are analyzed by the Db2 Insight Pack

KPI group	KPI name
CPU time	• DDF_SRB_TIME • DDF_TCB_TIME • DSAS_SRB_TIME • DSAS_TCB_TIME • IRLM_SRB_TIME • IRLM_TCB_TIME • SSAS_SRB_TIME • SSAS_TCB_TIME
Datasets	• CUR_OPEN_DATASETS
DDF	• ACC_QU_INAC_THR_T2 • ACTIVE_CI_CREATED • ACTIVE_DBATS • CONV_DEALLOC • CUR_QU_INAC_THR_T2 • DBAT_QUEUED • DBATS_INACTIVE • DBATS_NOT_USED • INACT_DBATS_T2 • TERM_COUNT

Table 8. KPIs that are analyzed by the Db2 Insight Pack (continued)

KPI group	KPI name
Latch	• LATCH_COUNTER_01 • LATCH_COUNTER_02 • LATCH_COUNTER_03 • LATCH_COUNTER_04 • LATCH_COUNTER_05 • LATCH_COUNTER_06 • LATCH_COUNTER_07 • LATCH_COUNTER_08 • LATCH_COUNTER_09 • LATCH_COUNTER_10 • LATCH_COUNTER_11 • LATCH_COUNTER_12 • LATCH_COUNTER_13 • LATCH_COUNTER_14 • LATCH_COUNTER_15 • LATCH_COUNTER_16 • LATCH_COUNTER_17 • LATCH_COUNTER_18 • LATCH_COUNTER_19 • LATCH_COUNTER_20 • LATCH_COUNTER_21 • LATCH_COUNTER_22 • LATCH_COUNTER_23 • LATCH_COUNTER_24 • LATCH_COUNTER_25 • LATCH_COUNTER_26 • LATCH_COUNTER_27 • LATCH_COUNTER_28 • LATCH_COUNTER_29 • LATCH_COUNTER_30 • LATCH_COUNTER_31 • LATCH_COUNTER_32 • LATCH_COUNTER_254

Table 8. KPIs that are analyzed by the Db2 Insight Pack (continued)

KPI group	**KPI name**
Local Locking	• CHANGE_REQ • CLAIM_REQUESTS • CLAIMS_FAILED • DEADLOCK • DRAIN_REQUESTS • DRAINS_FAILED • LOCK_REQ • LOCK_REQUESTS_TOTAL • OTHER_REQ • QUERY_REQ • SUSPENSION_LOCK • SUSPENSION_OTHER • SUSP_IRLM_LATCH • TIMEOUT
Logs	• LOG_CI_WRITTEN • LOG_WRITE_IO_REQ
Parallel Groups	• PAR_GROUPS_EXEC

Table 8. KPIs that are analyzed by the Db2 Insight Pack (continued)

KPI group	KPI name
Storage	• AUX_STORAGE_SLOT • A2GB_AUX_SLOT • A2GB_AUX_SLOT_TS • A2GB_COMMON_AUXS • A2GB_COMMON_REALF • A2GB_FIXED_STORAGE • A2GB_GETM_STORAGE • A2GB_REAL_FRAME • A2GB_REAL_FRAME_TS • A2GB_SHR_AUXS_STK • A2GB_SHR_AUXS_TS • A2GB_SHR_REALF_STK • A2GB_SHR_REALF_TS • A2GB_VAR_STORAGE • BIT_EXT_HGH_PRI_31 • BIT_EXT_LOW_PRI_31 • COMMON_FIXED_STOR • COMMON_GETM_STOR • COMMON_VAR_STOR • DIST_AUX_SLOT • DIST_REAL_FRAME • REAL_STORAGE_FRAME • TOT_FIXED_STORAGE • TOT_GETM_STCK_STOR • TOT_GETM_STORAGE

Installing the Problem Insights server on IBM z/OS UNIX System Services

To install the IBM Z Operations Analytics Problem Insights server on IBM z/OS UNIX System Services, complete these steps.

Procedure

1. Complete the installation steps in the *Program Directory for IBM z/OS UNIX System Services*, which is available at https://www.ibm.com/support/knowledgecenter/SS55JD_4.1.0/com.ibm.zosla.doc/pdf.html.
2. Start a z/OS UNIX System Services session by using one of the following methods:
   - From the Interactive System Productivity Facility (ISPF), by using the **omvs** command
   - By using Telnet
   - If a Secure Shell (SSH) daemon is available on the z/OS system, by using an SSH client. Typically, for z/OS UNIX System Services operations, the use of an SSH client provides the best user experience and a high level of security.

3. Go to the file system location where hierarchical file system (HFS) artifacts were installed during SMP/E processing.

   The default location is `/usr/lpp/IBM/zoa/V4R1M0`, but your system administrator might use a different location.

4. Set and export the following environment variables:

   WLPHOME

   Set the value of this variable to the installation directory for IBM WebSphere Application Server for z/OS Liberty.

   Tip: Although IBM WebSphere Application Server for z/OS Liberty is provided as part of IBM Z Common Data Provider FMID HHBO21L, any other copy of IBM WebSphere Application Server for z/OS Liberty Version 19, or later, can be used.

   JAVA_HOME

   Set the value of this variable to the installation directory for the Java Runtime Environment (JRE) for the Problem Insights server. For information about the runtime requirements, see "System requirements for the Problem Insights server" on page 16.

   PATH

   Add the following information to the value of this variable:

   ```
   $JAVA_HOME/bin
   ```

   Tips for z/OS UNIX System Services environment variables:

   - To set and export an environment variable, use the following command:

     ```
     export name="value"
     ```

     Example

     ```
     export JAVA_HOME="/Java/J8-0_64"
     ```

   - To append information to a path environment variable, use the following command:

     ```
     export name="$name:new_directory"
     ```

     Example

     ```
     export PATH="$PATH:$JAVA_HOME/bin:$PYTHON_HOME/bin
     ```

5. Verify that the following programs are available in the system path ($*PATH*).

   These programs are provided by either the operating system or one of the prerequisite software packages.

   - `chtag`
   - `gunzip`
   - `iconv`
   - `jar`
   - `keytool`
   - `pax`
   - `sed`
   - `tar`

   Tip for verifying that the programs are in the path: To determine the path name or command that the shell uses to call a program, use the following command:

```
command -v command-name
```

Example

```
command -v gunzip
```

If no path name or command is returned, the program is not available in the system path.

6. To start the IBM Z Operations Analytics setup, issue the setup command, as indicated in the following description.

   The setup command presents a set of menus to help you navigate the setup process. The menus prompt you for input. After you provide the requested input, press **Enter** to proceed to the next step.

   To launch the setup utility, issue the following setup command from the installation directory:

   ```
   ./bin/izoa-setup.sh
   ```

   The main menu of the setup utility is shown with the following options:

   ```
   ************************************************************
   *           Z Operations Analytics Setup Menu              *
   ************************************************************
   Select one of the following processing options:
   1       Perform prerequisite check
   2       Create runtime environment
   3       Configure runtime environment and machine learning system instance
   4       Remove machine learning system instance
   5       Remove runtime environment
   6       Print configuration settings
   7       Deploy machine learning support into Problem Insights server
   8       Remove machine learning support from Problem Insights server
   9       Manage rules engine
   10      Change passwords
   11      Collect logs
   12      Quit
   ```

   Tip: You can process only one setup menu option each time that you run the setup command. Therefore, to process multiple options, you must run the setup command multiple times.

7. Create the IBM Z Operations Analytics runtime environment, as described in the following steps:

   a) From the main menu of the setup utility, select option 2, which is `Create runtime environment`.

   b) In the secondary menu, specify that you want to set up a runtime environment for the IBM Z Operations Analytics Problem Insights server (suboption 2).

   c) Provide your target directory for the runtime environment.

8. Exit the setup utility.

What to do next

Complete the steps in "Configuring the Problem Insights server" on page 30.

Installing the Problem Insights server on a Linux system

To install the IBM Z Operations Analytics Problem Insights server on a Linux system, complete these steps.

Procedure

1. Either mount the DVD that contains the IBM Z Operations Analytics Insight Pack media (LCD7-7516), or extract the IBM Z Operations Analytics Insight Pack electronic image (LCD7-7517) into a temporary directory on the target system.

Tip: In the remaining steps, the term *installation directory* refers to this directory.

2. Start a command shell on the target system, or remotely log in to a terminal session by using a Secure Shell (SSH) client.
3. Verify that the following programs are available in the system path ($*PATH*).

   These programs are provided by the operating system.

   - `gunzip`
   - `keytool`
   - `sed`
   - `tar`
   - `unzip`
4. In the command shell, go to the installation directory.
5. Verify that the directory for the appropriate Java Runtime Environment (JRE) for the Problem Insights server is in the system path ($*PATH*).

   For information about the runtime requirements, see "System requirements for the Problem Insights server" on page 16.
6. To start the IBM Z Operations Analytics setup, issue the setup command, as indicated in the following description.

   The setup command presents a set of menus to help you navigate the setup process. The menus prompt you for input. After you provide the requested input, press **Enter** to proceed to the next step.

   To launch the setup utility, issue the following commands from the installation directory:

   ```
   cd piserver
   ./bin/izoa-setup.sh
   ```

   The main menu of the setup utility is shown with the following options:

   ```
   ************************************************************
   *           Z Operations Analytics Setup Menu              *
   ************************************************************
   Select one of the following processing options:
   1) Perform prerequisite check     5) Change passwords
   2) Create runtime environment     6) Collect logs
   3) Configure runtime environment  7) Quit
   4) Remove runtime environment

   Your selection:
   ```
7. Create the IBM Z Operations Analytics runtime environment, as described in the following steps:

   a) From the main menu of the setup utility, select option 2, which is `Create runtime environment`.

   b) In the secondary menu, specify that you want to set up a runtime environment for the IBM Z Operations Analytics Problem Insights server (suboption 2).

   c) Provide your target directory for the runtime environment.
8. Exit the setup utility.

What to do next

Complete the steps in "Configuring the Problem Insights server" on page 30.

Installing the Problem Insights server on a Windows system

To install the IBM Z Operations Analytics Problem Insights server on a Windows system, complete these steps.

Procedure

1. Start a command window on the target system.
2. Mount the DVD that contains the IBM Z Operations Analytics Insight Pack media (LCD7-7516), or extract the IBM Z Operations Analytics Insight Pack electronic image (LCD7-7517) into a temporary directory on the target system.

   For the remainder of this section, the term *installation directory* refers to this directory.
3. In the command window, go to the installation directory.
4. Verify that the directory for the appropriate Java Runtime Environment (JRE) for the Problem Insights server is in the system path (*%PATH%*).

   For information about the runtime requirements, see "System requirements for the Problem Insights server" on page 16.
5. To start the IBM Z Operations Analytics setup, issue the setup command, as indicated in the following description.

   The setup command presents a set of menus to help you navigate the setup process. The menus prompt you for input. After you provide the requested input, press **Enter** to proceed to the next step.

   To launch the setup utility, issue the following commands from the installation directory:

   ```
   cd piserver
   bin\izoa-setup.bat
   ```

   The main menu of the setup utility is shown with the following options:

   ```
   ===================== Z Operations Analytics Setup Menu =====================
   Select one of the following processing options:
   1. Perform prerequisite check
   2. Create runtime environment
   3. Configure runtime environment
   4. Remove runtime environment
   5. Change passwords
   6. Collect logs
   Q. Quit
   Select one of the following processing options:
   ```
6. Create the IBM Z Operations Analytics runtime environment, as described in the following steps:

   a) From the main menu of the setup utility, select option 2, which is `Create runtime environment`.

   b) In the secondary menu, specify that you want to set up a runtime environment for the IBM Z Operations Analytics Problem Insights server (suboption 2).

   c) Provide your target directory for the runtime environment.
7. Exit the setup utility.

What to do next

Complete the steps in "Configuring the Problem Insights server" on page 30.

Configuring the Problem Insights server

You can configure the IBM Z Operations Analytics Problem Insights server to interact with the following systems: the Elastic Stack or Splunk platform for message-based analysis, IBM Db2 for z/OS for metric-based analysis by machine learning, and event management systems for forwarding events.

About this task

The following topics contain important reference information:

- "Problem Insights server: command reference" on page 83
- "Problem Insights server: configuration file reference" on page 85
- "Problem Insights server: message library reference" on page 86

Depending on how you want to use the Problem Insights server, the configuration files that are described in "Problem Insights server: configuration file reference" on page 85 must be copied from the `samples` directory to the `config` directory. You must then update the files according to the comments in the files, and restart the Problem Insights server.

Important:

On IBM z/OS UNIX System Services, the following configuration files are automatically copied from the `samples` directory to the `config` directory and customized by the setup command `izoa-setup.sh` during runtime ***creation***:

- `amq.config`
- `cli.config`

Also on IBM z/OS UNIX System Services, the following configuration files are automatically copied from the `samples` directory to the `config` directory and customized by the setup command `izoa-setup.sh` during runtime ***configuration***:

- `analysisroutine.config`
- `db2zos.config`
- `microservices.config`

On Linux and Windows systems, only the `cli.config` file is automatically copied.

Important: On IBM z/OS UNIX System Services, the configuration files are encoded in ISO8859-1 and are tagged in the HFS accordingly. To view the tagging of the files, run the following command from the runtime (*IZOA_HOME*) directory:

```
ls -T *.config
```

To update these files, you must use a file editor that is capable of reading and writing ISO8859-1 data. If you plan to use the vi editor, or a similar command line editor on z/OS UNIX System Services, set the value of the *_BPXK_AUTOCVT* environment variable to ON to enable automatic codeset conversion between EBCDIC and ISO8859-1.

```
export _BPXK_AUTOCVT=ON
```

Procedure

1. In the `config` directory, open the `cli.config` file, review the values of all configuration properties, and update the values as appropriate.
2. Save and close the `cli.config` file.
3. Again, launch the setup utility, as described in the following information:
   - "Installing the Problem Insights server on IBM z/OS UNIX System Services" on page 25
   - "Installing the Problem Insights server on a Linux system" on page 27

    - “Installing the Problem Insights server on a Windows system” on page 29

4. Configure the IBM Z Operations Analytics runtime environment, as described in the following steps:
    a) From the main menu of the setup utility, select option 3, which is `Configure runtime environment`.
    b) In the secondary menu, specify that you want to configure a runtime environment for the IBM Z Operations Analytics Problem Insights server (suboption 2).
    c) Provide your target directory for the runtime environment.
    d) Respond to any prompts in the window.
5. Exit the setup utility.

Configuring the rules engine

In the Problem Insights server configuration file for the rules engine (`analysisroutine.config`), you can specify the KPIs for which you want alerts to be sent if anomalies occur.

About this task

For information about the KPIs that you can specify, see Insight Packs.

Procedure

To have alerts sent if anomalies occur, specify the configuration properties **`izoa.kpi.anomaly.ar.alert.config`** and **`izoa.kpi.anomaly.ar.enabled`** in the `analysisroutine.config` file.

izoa.kpi.anomaly.ar.alert.config

This value specifies the list of KPIs for which you want alerts to be sent if anomalies occur. These KPIs are described in “Insight Packs” on page 20. The KPI list must be specified in JavaScript Object Notation (JSON) format, as shown in the following example:

```
izoa.kpi.anomaly.ar.alert.config={"alertConfig":
[{
"kpiName":"ACTIVE_DBATS","subsystemIds":["ALL"],
"scoreThreshold":90,"snoozeMinutes":60,"subsystemType":"DB2","messageId":"DB2MLSSC"}
{"kpiName":"RESPONSE_SEC","subsystemIds":["CICS01","CICS02"],
"scoreThreshold":68,"snoozeMinutes":30,"subsystemType":"CICS","messageId":"CICSMLSSC"}
]}
```

This example defines two KPIs, Db2 KPI ACTIVE_DBATS and CICS API RESPONSE_SEC for sending alerts.

Table 9 on page 31 indicates the parameter values that are required for each KPI that you define in the configuration property **`izoa.kpi.anomaly.ar.alert.config`**.

*Table 9. Parameter values are required for each KPI that you define in the configuration property **izoa.kpi.anomaly.ar.alert.config***

Parameter	Description
`kpiName`	The name of the KPI for which alerts are to be sent. This name is persisted in the enterprise data warehouse. For information about the KPI names for each subsystem type, see “Insight Packs” on page 20.
`subsystemIds`	The Db2 subsystems or CICs regions for which alerts are to be sent for the specified KPI. To specify all subsystems or regions, use the value ALL.

Table 9. Parameter values are required for each KPI that you define in the configuration property ***izoa.kpi.anomaly.ar.alert.config*** *(continued)*

Parameter	Description
scoreThreshold	The anomaly score that is the threshold for triggering alerts for the specified KPI. If the KPI anomaly score exceeds this threshold, an alert is sent. For more information about scoring, see "Data scoring overview" on page 48.
snoozeMinutes	The number of minutes between each time that an alert is sent for the specified KPI. For example, you might want to have a KPI alert sent every 60 minutes.
subsystemType	The type of the subsystem for which alerts are to be sent for the specified KPI. The following values are valid: • CICS • DB2 These values correspond to the respective Insight Packs, as described in "Insight Packs" on page 20.
messageId	The message ID of the message that is associated with the specified KPI. The following values are valid: • CICSMLSSC (for the CICS subsystem type) • DB2MLSSC (for the DB2 subsystem type)

izoa.kpi.anomaly.ar.enabled

This value specifies a true or false indication of whether to enable (`true`) or disable (`false`) the anomaly analysis routine.

To have the alerts sent, this value must be `true`.

Starting the Problem Insights server and the message broker

After the configuration of the IBM Z Operations Analytics Problem Insights server is complete, you must start (or restart) the server and the message broker.

About this task

Any configuration update to the `amq.config` file takes effect only after you stop and restart the Apache ActiveMQ message broker. Any configuration update to the other Problem Insights server configuration files takes effect only after you stop and restart the Problem Insights server. For more information about the Problem Insights server configuration files, see "Problem Insights server: configuration file reference" on page 85.

For more information about the commands for operating the IBM Z Operations Analytics Problem Insights server, see "Problem Insights server: command reference" on page 83.

Procedure

After configuration, run the following command, based on your operating system.

1. To start the Problem Insights server, run the following command.

   Linux system

   ```
   bin/analysis.sh start
   ```

   Tip: If the Problem Insights server is active, run the following command first to stop it:

```
bin/analysis.sh stop
```

Windows system

```
bin\analysis.bat start
```

Tip: If the Problem Insights server is active, run the following command first to stop it:

```
bin\analysis.bat stop
```

z/OS system
System Display and Search Facility (SDSF) command:

```
/S GLAPISRV
```

Tip: If the Problem Insights server is active, run the following System Display and Search Facility (SDSF) command first to stop it:

```
/C GLAPISRV
```

2. To start the Apache ActiveMQ message broker, run the following command on the z/OS system.

```
cd IZOA_HOME/bin
./mqcontrol.sh start
```

Tip: If the Apache ActiveMQ message broker is active, run the following command first to stop it:

```
cd IZOA_HOME/bin
./mqcontrol.sh stop
```

Installing or updating Insight Packs

The IBM Z Operations Analytics Insight Packs for the Problem Insights server are available only for insights that are produced by the IBM Z Operations Analytics machine learning system (applicable only to IBM z/OS UNIX System Services). Install the appropriate Insight Packs into the IBM Z Operations Analytics Problem Insights server so that it can provide visualizations of the subsystem-specific insights from the machine learning system.

About this task

"Insight Packs" on page 20 includes more information about the Insight Packs that are provided with IBM Z Operations Analytics.

Procedure

To install or update Insight Packs, log in to the Problem Insights server by using the administrative user ID and password that is defined in the Problem Insights server `cli.config` file.

From the Problem Insights server GUI, you can manage the lifecycle of Insight Packs. Managing the life cycle includes the following tasks:

- Installing an Insight Pack
- Enabling an Insight Pack
- Disabling an Insight Pack
- Updating an Insight Pack
- Uninstalling an Insight Pack

Important: The administrator user ID cannot have multiple, concurrent sessions with the Problem Insights server. For example, the administrator user ID cannot be simultaneously logged in to the server from two different computers, or from two different browsers on the same computer.

Sending events from the Problem Insights server to other applications

You can configure the IBM Z Operations Analytics Problem Insights server to send events to IBM Tivoli Netcool/OMNIbus, to a cloud-based event management service, or to IBM Z ChatOps.

About this task

Forwarding events to these other applications involves the following activities:

Preparation of the receiving application
In this activity, you prepare the receiving application (for example, the event management system or service) to receive events from the IBM Z Operations Analytics Problem Insights server and to map them to the event format that the application uses.

Preparation of the event source
In this activity, you configure the communication between the IBM Z Operations Analytics Problem Insights server and the product or component that provides the event information to be forwarded.

Definition of the event forwarding criteria
In this activity, you define the criteria that is used to identify events that should be forwarded to the application.

Sending events to IBM Tivoli Netcool/OMNIbus

You can configure the IBM Z Operations Analytics Problem Insights server to send events to IBM Tivoli Netcool/OMNIbus.

Before you begin

Review "Configuration planning for the Problem Insights server" on page 18.

.

Procedure

1. To prepare IBM Tivoli Netcool/OMNIbus to receive events from the Problem Insights server, complete the following steps.
   a) Configure IBM Tivoli Netcool/OMNIbus Probe for Message Bus rules for IBM Z Operations Analytics events, as described in "Configuring IBM Tivoli Netcool/OMNIbus Probe for Message Bus rules for IBM Z Operations Analytics events" on page 36.
   b) Enable quick links from the IBM Tivoli Netcool/OMNIbus Event Viewer to the IBM Z Operations Analytics Problem Insights GUI, as described in "Enabling quick links from the IBM Tivoli Netcool/OMNIbus Web GUI Event Viewer to the Problem Insights GUI" on page 37.
   c) Define the connection parameters for your OMNIbus instance in the *IZOA_HOME*`/config/event.config` file.
2. In the IBM Z Operations Analytics `event.config` file, update the following configuration properties.

   eventapi.user
   The user name for the IBM Tivoli Netcool/OMNIbus event management system to which the Problem Insights server must connect.

   eventapi.password
   The password for the IBM Tivoli Netcool/OMNIbus user name.

eventapi.url
The URL for the IBM Tivoli Netcool/OMNIbus event management system.

3. To prepare the event source, complete the following steps that apply for your event source.

Event source	Steps
IBM Z Operations Analytics machine learning system (for metric-based machine learning events)	a. Enable continuous scoring by running the **`ml=scag_batchscore`** command, as described in "Analyzing and scoring new SMF data" on page 72. b. Enable the forwarding of continuous scoring results, which is also described in "Analyzing and scoring new SMF data" on page 72. c. Configure and enable the rules engine, as described in "Configuring the rules engine" on page 31. d. Configure and start the Apache ActiveMQ message broker, as described in "Starting the Problem Insights server and the message broker" on page 32.
Elastic Stack	a. Copy the `elk.config` file from *`IZOA_HOME`*`/samples` to *`IZOA_HOME`*`/config`. b. Edit *`IZOA_HOME`*`/config/elk.config` according to the comments in the file.
Splunk	a. Copy the `splunk.config` file from *`IZOA_HOME`*`/samples` to *`IZOA_HOME`*`/config`. b. Edit *`IZOA_HOME`*`/config/splunk.config` according to the comments in the file.

Configuring more than one event source: You can configure the machine learning system and either Elastic Stack or Splunk as event sources. However, you cannot configure both Elastic Stack and Splunk to be event sources at the same time.

4. Define the event forwarding criteria as indicated.

Event forwarding criteria for metric-based machine learning insights
Specify the configuration property **`forward_messages`** in the `event.config` file to include the same message IDs that you specify in the **`izoa.kpi.anomaly.ar.alert.config`** property in the `analysisroutine.config` file.

Event forwarding criteria for single message insights (Elastic Stack or Splunk)
Each insight from Elastic Stack or Splunk is based on a single message. To forward event notifications for insights from Elastic Stack or Splunk, specify the configuration property **`forward_messages`** in the `event.config` file to include one or more of the message IDs that are defined in the following Problem Insights server message libraries:

- `usr/lang/en/IZOACICSInsights.xml`
- `usr/lang/en/IZOADB2Insights.xml`
- `usr/lang/en/IZOAMQInsights.xml`
- `usr/lang/en/IZOANetworkInsights.xml`
- `usr/lang/en/IZOAzOSInsights.xml`

For more information about message libraries, see "Problem Insights server: message library reference" on page 86.

Tip: If you want to forward event notifications for insights from both machine learning and Elastic Stack or Splunk, specify the configuration property **`forward_messages`** in the `event.config` file to include the relevant message IDs for both types of insights.

Configuring IBM Tivoli Netcool/OMNIbus Probe for Message Bus rules for IBM Z Operations Analytics events

The IBM Z Operations Analytics Problem Insights server forwards events to IBM Tivoli Netcool/OMNIbus through the OMNIbus Probe for Message Bus. Other types of integration between IBM Z Operations Analytics and IBM Tivoli Netcool/OMNIbus are not supported. To enable event forwarding and the subsequent display of events in the NetCool/OMNIbus Web GUI Event Viewer, you must configure HTTP transport properties, event sources, and message bus rules for the OMNIbus Probe for Message Bus.

Before you begin

Install the IBM Tivoli Netcool/OMNIbus Probe for Message Bus, if it is not yet installed. For more information, see Probe for Message Bus in the IBM Tivoli Netcool/OMNIbus documentation.

Procedure

To configure the IBM Tivoli Netcool/OMNIbus Probe for Message Bus, complete the following steps:

1. Copy the file `$OMNIHOME/java/conf/httpTransport.properties` to the directory `$OMNIHOME/probes`.
2. In the file `$OMNIHOME/probes/httpTransport.properties`, uncomment the code `serverPort=http:80`.
3. In the file `$OMNIHOME/probes/`*`architecture_directory`*`/message_bus.props`, add the following code near the end of the file, immediately before the phrase `#PROTECTED REGION END#`:

   ```
   Server : 'AGG_P'
   TransportType : 'HTTP'
   TransportFile : '${OMNIHOME}/probes/httpTransport.properties'
   TransformerFile : '${OMNIHOME}/probes/architecture_directory/message_bus_parser_config.json'
   AuthUserName : 'root'
   AuthPassword : 'root_password'
   Port : 80
   ```

4. In the file `$OMNIHOME/probes/`*`architecture_directory`*`/message_bus_parser_config.json`, verify that one of the configurations for `eventSources` is set to the following value:

   ```
   {
       "endpoint" : "http:80",
       "name" : "NotificationAlarmParser",
       "config" : {
           "dataToRecord" : [
               "first-raise-time",
               "last-raise-time"
           ],
           "messagePayload" : "json",
           "messageHeader" : "json",
           "jsonNestedPayload" : "",
           "jsonNestedHeader" : "",
           "messageDepth" : 3
       }
   }
   ```

5. In the file `$OMNIHOME/probes/`*`architecture_directory`*`/message_bus.rules`, customize the field assignments so that IBM Tivoli Netcool/OMNIbus can properly process the JavaScript Object Notation (JSON) tokens that are included in IBM Z Operations Analytics events.

   The following fields must be added:

   - `@Evidence`
   - `@Subsystem`
   - `@Sys`
   - `@Sysplex`

- `@Url`

The following fields must be redefined:

- `@AlertGroup`
- `@AlertKey`
- `@Identifier`
- `@Node`
- `@Severity`
- `@Summary`

For a sample `message_bus.rules` file, see “Sample message_bus.rules file” on page 38.

6. Complete the following steps to add the newly added fields to the IBM Tivoli Netcool/OMNIbus database schema:
   a) Start the SQL interface to your IBM Tivoli Netcool/OMNIbus ObjectServer, as described in Starting the SQL interactive interface in the IBM Tivoli Netcool/OMNIbus documentation.

   b) Enter the following commands:

   ```
   alter table alerts.status add Subsystem varchar(1024);
   alter table alerts.status add Url varchar(1024);
   alter table alerts.status add Sys varchar(1024);
   alter table alerts.status add Evidence varchar(1024);
   alter table alerts.status add Sysplex varchar(1024);
   go
   exit
   ```
7. Start or restart the IBM Tivoli Netcool/OMNIbus Probe for Message Bus.

Enabling quick links from the IBM Tivoli Netcool/OMNIbus Web GUI Event Viewer to the Problem Insights GUI

The events that the IBM Z Operations Analytics machine learning system sends to IBM Tivoli Netcool/OMNIbus include a URL that is a direct link to a problem evidence page in the Problem Insights GUI. To simplify the use of this URL in the NetCool/OMNIbus Web GUI Event Viewer, you must create a script tool in IBM Tivoli Netcool/OMNIbus, and associate this tool with the **Alerts** menu in the Event Viewer.

Procedure

To create the menu item, complete the following steps:

1. Create a new tool as described in Defining tools in the IBM Tivoli Netcool/OMNIbus documentation.
   a) Assign a meaningful name, such as `ZOA_evidence`, to the new tool.

   b) Select **Script** as the tool type, and follow the appropriate instructions.

   c) Under **Script Commands**, enter the following code:

   ```
   window.open("{@Url}");
   ```

   d) Save the new tool.
2. Assign the new tool to the **Alerts** menu as described in Changing menus in the IBM Tivoli Netcool/OMNIbus documentation.
   a) Select **alerts** from the **Available menus** list, and click **Modify**.

   b) Under **Available items**, select the name that you used for the new tool (for example, `ZOA_evidence`), and add it under **Current items**.

   c) Save the updated menu.

Sample `message_bus.rules` file

A sample `message_bus.rules` file is provided for reference.

```
########################################################################
#
#       Licensed Materials - Property of IBM
#
#
#
#       (C) Copyright IBM Corp. 2015. All Rights Reserved
#
#       US Government Users Restricted Rights - Use, duplication
#       or disclosure restricted by GSA ADP Schedule Contract
#       with IBM Corp.
#
#
#######################################################################


if( match( @Manager, "ProbeWatch" ) )
{
    switch(@Summary)
    {
    case "Running ...":
                    @Severity = 1
                    @AlertGroup = "probestat"
                    @Type = 2
    case "Going Down ...":
                    @Severity = 5
                    @AlertGroup = "probestat"
                    @Type = 1
    case "Start resynchronization" | "Finish resynchronization":
                    @Severity = 2
                    @AlertGroup = "probestat"
                    @Type = 13
    case "Connection to source lost":
                    @Severity = 5
                    @AlertGroup = "probestat"
                    @Type = 1
    default:
                    @Severity = 1
    }
    @AlertKey = @Agent
    @Summary = @Agent + " probe on " + @Node + ": " + @Summary
}
else
{
    @Manager = %Manager + " probe running on " + hostname()
    @Node = $Node
    @NodeAlias = %Host + ":" + %Port
    @Class = 30505

    if (exists($TransformerName))
    {
        switch($TransformerName)
        {
            case "dummy case statement": ### This will prevent syntax errors in case no
includes are added below.


                        include "message_bus_netcool.rules"
                        include "message_bus_cbe.rules"
                        include "message_bus_wbe.rules"
                        include "message_bus_wef.rules"

            default:
                log(DEBUG, "<<<<< Rules are not supported for this format >>>>")

            @Summary = "Rules are not supported for this format - " + $TransformerName
        }
    }
    else
    {
        log(DEBUG, "<<<<< Entering... message_bus_netcool.rules >>>>>")


        @Manager = %Manager
        @Class = 89210
        @NodeAlias = $NodeAlias
        @Agent = $Agent
        @AlertKey = $sysplex + "." + $system + "." + $subsystem
        @Severity = int($severity)
        @Subsystem = $subsystem
        @Sysplex = $sysplex
        @Url     = $url
        @Evidence = $evidence
        @Sys = $system
        @Node = @Sysplex + "." + @Sys
        @Summary = $evidence + " " + $summary
        #@Summary = $summary
        @StateChange = $StateChange
        #@FirstOccurrence = $FirstOccurrence
        #@LastOccurrence = $LastOccurrence
        #@InternalLast = $InternalLast
        @Poll = $Poll
        @Type = $Type
        @Tally = $Tally
        @Class = $Class
        @Grade = $Grade
        @Location = $Location
        @OwnerUID = $OwnerUID
        @OwnerGID = $OwnerGID
        @Acknowledged = $Acknowledged
        @Flash = $Flash
        @EventId = $EventId
        @ExpireTime = $ExpireTime
        @ProcessReq = $ProcessReq
        @SuppressEscl = $SuppressEscl
        @Customer = $Customer
        @Service = $Service
        @PhysicalSlot = $PhysicalSlot
        @PhysicalPort = $PhysicalPort
        @PhysicalCard = $PhysicalCard
        @TaskList = $TaskList
        @NmosSerial = $NmosSerial
        @NmosObjInst = $NmosObjInst
        @NmosCauseType = $NmosCauseType
        @LocalNodeAlias = $LocalNodeAlias
        @LocalPriObj = $LocalPriObj
        @LocalSecObj = $LocalSecObj
        @LocalRootObj = $LocalRootObj
        @RemoteNodeAlias = $RemoteNodeAlias
        @RemotePriObj = $RemotePriObj
        @RemoteSecObj = $RemoteSecObj
        @RemoteRootObj = $RemoteRootObj
        @X733EventType = $X733EventType
        @X733ProbableCause = $X733ProbableCause
        @X733SpecificProb = $X733SpecificProb
        @X733CorrNotif = $X733CorrNotif
        @ExtendedAttr = $ExtendedAttr
        @ServerName = $ServerName
        @ServerSerial = $ServerSerial
        if (exists($evidence))
        {
            switch(substr($evidence,1,3))
            {
                case "DB2" | "CIC":
                        @AlertGroup = "MLInsight"
                        @AlertKey = $sysplex + "." + $system + "." + $subsystem + "." +
@LastOccurrence

                case "DFH" | "DSN" | "CSQ" | "EZZ" | "EZD" | "IVT" | "IST" | "IEA" | "IEF":
                        @AlertGroup = "SingleMsgInsight"
                default:
                    log(DEBUG, "<<<<< Non-ML evidence value found  >>>>")
            }
        }
        else
        {}


        @Identifier = @Node + " " + @AlertKey + " " + @AlertGroup + " " + @Type + " " + @Agent
+ " " + @Manager


        log(DEBUG, "<<<<< Leaving... message_bus_netcool.rules >>>>>")
    }
}
```

Sending events to cloud-based event management services

By using webhook integration, you can configure the IBM Z Operations Analytics Problem Insights server to send events to cloud-based event management services, such as IBM Cloud Event Management, IBM Netcool Operations Insight, and IBM Watson AIOps Event Manager.

Before you begin

Review “Configuration planning for the Problem Insights server” on page 18.

Tip: By default, the supported cloud-based event management systems aggregate into a single event the events that are sent by the Problem Insights server for the same KPI and the same subsystem. This aggregated event includes the following information:

- A time stamp that indicates the latest occurrence of the event
- A counter that indicates the number of individual event occurrences that are now aggregated
- The information that was submitted with the first occurrence of the event

The URLs that are included in the Problem Insights server events link to analytics evidence for a specified time window. Because of the event aggregation by the cloud-based event management system, the URL in the aggregated event links, by default, to the *first occurrence* of the event in the Problem Insights server.

About this task

Important: The configuration for sending events to IBM Cloud Event Management, IBM Netcool Operations Insight, or IBM Watson AIOps Event Manager is the same for each of these three services. In this procedure, the configuration is described by using the example of IBM Cloud Event Management.

You can insert event information into IBM Cloud Event Management from any event source that can send the information in JavaScript Object Notation (JSON) format. You use a webhook URL to set your event source to send event information to IBM Cloud Event Management. Then, by using an example incoming request in JSON format, you define the mapping between the event attributes from your event source and the event attributes in IBM Cloud Event Management.

Procedure

1. To prepare IBM Cloud Event Management to receive events from the Problem Insights server, complete the following steps.

   a) Create a webhook for defining the mapping between the event attributes from your event source and the event attributes in IBM Cloud Event Management.

   For more information, see Creating custom event sources with JSON in the IBM Cloud Event Management documentation.

   Important: As part of creating the webhook, you must create an API key, which provides you with a user name and password. This user name and password is populated in the Problem Insights server configuration file `event.config`.

   b) Define the mapping between the event attributes from the Problem Insights server event payload (the example incoming request) and the event attributes in IBM Cloud Event Management.

   Table 10 on page 41 shows the correlation of these event attributes.

   Example of Problem Insights server event payload

```
{
    "summary": "Abnormal behavior has been identified for a Db2 subsystem.
                  Investigate the metrics that have deviated from their normal
patterns.",
    "shortSummary": "Abnormal behavior has been identified for a Db2 subsystem.
                       Investigate the metrics that have deviated from their normal
patterns.",
    "time": "2020-08-28 01:00:00",
    "subsystem": "DC1V",
    "url": "https://localhost:9446/piFramework/protected/izoa/v1/pi/DB2MLSSC/4601",
    "system": "system",
    "eventId": "4601",
    "evidence": "DB2MLSSC",
    "kpiName": "ACTIVE_DBATS_MAX",
    "sysplex": "LPAR400J",
    "severity": "7"
}
```

Table 10. Correlation of the event attributes of IBM Cloud Event Management to the event attributes in the IBM Z Operations Analytics Problem Insights server event payload

IBM Cloud Event Management event attributes	**Event attributes in Problem Insights server event payload**
`resource name` In the IBM Cloud Event Management UI, events that have the same resource name are grouped together.	`sysplex.system` The values of `sysplex` and `system` in the event payload (in that order) must be concatenated to form the value of the event attribute `resource name`.
`resource type` In the IBM Cloud Event Management UI, events that have the same resource type are grouped together. The evidence URL of the first occurrence of the event is used for all events of that resource type.	`sysplex.system.subsystem.kpiName` The values of `sysplex`, `system`, `subsystem`, and `kpiName` in the event payload (in that order) must be concatenated to form the value of the event attribute `resource type`.
`event type`	`evidence` (the message ID)
`sender name`	"IZOA Webhook"
`url description`	"Evidence URL"
`severity`	• `1=Critical` • `2=Critical` • `3=Warning` • `4=Information` • `5=Indeterminate` • `6=Indeterminate` • `7=Indeterminate`

c) Define the connection parameters for your IBM Cloud Event Management instance in the *IZOA_HOME*`/config/event.config` file.

2. In the IBM Z Operations Analytics `event.config` file, update the following configuration properties.

   izoa.notification.cem.user
   The user name for the IBM Cloud Event Management, IBM Netcool Operations Insight, or IBM Watson AIOps Event Manager service to which the Problem Insights server must connect.

   izoa.notification.cem.password
   The password for the IBM Cloud Event Management, IBM Netcool Operations Insight, or IBM Watson AIOps Event Manager user name.

   izoa.notification.cem.url
   The URL for the IBM Cloud Event Management, IBM Netcool Operations Insight, or IBM Watson AIOps Event Manager webhook integration.

3. To prepare the event source, complete the following steps that apply for your event source.

Event source	Steps
IBM Z Operations Analytics machine learning system (for metric-based machine learning events)	a. Enable continuous scoring by running the **`ml=scag_batchscore`** command, as described in "Analyzing and scoring new SMF data" on page 72. b. Enable the forwarding of continuous scoring results, which is also described in "Analyzing and scoring new SMF data" on page 72. c. Configure and enable the rules engine, as described in "Configuring the rules engine" on page 31. d. Configure and start the Apache ActiveMQ message broker, as described in "Starting the Problem Insights server and the message broker" on page 32.
Elastic Stack	a. Copy the `elk.config` file from *`IZOA_HOME`*`/samples` to *`IZOA_HOME`*`/config`. b. Edit *`IZOA_HOME`*`/config/elk.config` according to the comments in the file.
Splunk	a. Copy the `splunk.config` file from *`IZOA_HOME`*`/samples` to *`IZOA_HOME`*`/config`. b. Edit *`IZOA_HOME`*`/config/splunk.config` according to the comments in the file.

Configuring more than one event source: You can configure the machine learning system and either Elastic Stack or Splunk as event sources. However, you cannot configure both Elastic Stack and Splunk to be event sources at the same time.

4. Define the event forwarding criteria as indicated.

Event forwarding criteria for metric-based machine learning insights

Specify the configuration property **`forward_messages`** in the `event.config` file to include the same message IDs that you specify in the **`izoa.kpi.anomaly.ar.alert.config`** property in the `analysisroutine.config` file.

Event forwarding criteria for single message insights (Elastic Stack or Splunk)

Each insight from Elastic Stack or Splunk is based on a single message. To forward event notifications for insights from Elastic Stack or Splunk, specify the configuration property **`forward_messages`** in the `event.config` file to include one or more of the message IDs that are defined in the following Problem Insights server message libraries:

- `usr/lang/en/IZOACICSInsights.xml`
- `usr/lang/en/IZOADB2Insights.xml`
- `usr/lang/en/IZOAMQInsights.xml`
- `usr/lang/en/IZOANetworkInsights.xml`
- `usr/lang/en/IZOAzOSInsights.xml`

For more information about message libraries, see "Problem Insights server: message library reference" on page 86.

Tip: If you want to forward event notifications for insights from both machine learning and Elastic Stack or Splunk, specify the configuration property **`forward_messages`** in the `event.config` file to include the relevant message IDs for both types of insights.

Sending events to IBM Z ChatOps

You can configure the IBM Z Operations Analytics Problem Insights server to send events to IBM Z ChatOps.

Before you begin

Review “Configuration planning for the Problem Insights server” on page 18.

Procedure

1. In your Slack or Mattermost chat application, create a new channel, and add "Bot User" as a member of the newly created channel, as described in the following IBM Z ChatOps documentation:
   - For Slack, see Add your bot user to your Slack channel.
   - For Mattermost, see Inviting the created bot to your Mattermost channel.
2. In the IBM Z Operations Analytics `event.config` file, update the following configuration properties.

   izoa.notification.chatops.user
   The user account name for the IBM Z ChatOps server to which the Problem Insights server must connect.

   izoa.notification.chatops.password
   The password for the IBM Z ChatOps user account name.

   izoa.notification.chatops.url
   The URL for the IBM Z ChatOps Incident API.

   izoa.notification.chatops.post.to.channel
   The name of the channel in your chat application that is connected to IBM Z ChatOps. This channel is the channel that you created in step “1” on page 43, and it is the channel where events from the Problem Insights server are to be posted.
3. To prepare the event source, complete the following steps that apply for your event source.

Event source	Steps
IBM Z Operations Analytics machine learning system (for metric-based machine learning events)	a. Enable continuous scoring by running the **`ml=scag_batchscore`** command, as described in “Analyzing and scoring new SMF data” on page 72. b. Enable the forwarding of continuous scoring results, which is also described in “Analyzing and scoring new SMF data” on page 72. c. Configure and enable the rules engine, as described in “Configuring the rules engine” on page 31. d. Configure and start the Apache ActiveMQ message broker, as described in “Starting the Problem Insights server and the message broker” on page 32.
Elastic Stack	a. Copy the `elk.config` file from *`IZOA_HOME`*`/samples` to *`IZOA_HOME`*`/config`. b. Edit *`IZOA_HOME`*`/config/elk.config` according to the comments in the file.
Splunk	a. Copy the `splunk.config` file from *`IZOA_HOME`*`/samples` to *`IZOA_HOME`*`/config`. b. Edit *`IZOA_HOME`*`/config/splunk.config` according to the comments in the file.

Configuring more than one event source: You can configure the machine learning system and either Elastic Stack or Splunk as event sources. However, you cannot configure both Elastic Stack and Splunk to be event sources at the same time.

4. Define the event forwarding criteria as indicated.

Event forwarding criteria for metric-based machine learning insights

Specify the configuration property **`forward_messages`** in the `event.config` file to include the same message IDs that you specify in the **`izoa.kpi.anomaly.ar.alert.config`** property in the `analysisroutine.config` file.

Event forwarding criteria for single message insights (Elastic Stack or Splunk)

Each insight from Elastic Stack or Splunk is based on a single message. To forward event notifications for insights from Elastic Stack or Splunk, specify the configuration property **`forward_messages`** in the `event.config` file to include one or more of the message IDs that are defined in the following Problem Insights server message libraries:

- `usr/lang/en/IZOACICSInsights.xml`
- `usr/lang/en/IZOADB2Insights.xml`
- `usr/lang/en/IZOAMQInsights.xml`
- `usr/lang/en/IZOANetworkInsights.xml`
- `usr/lang/en/IZOAzOSInsights.xml`

For more information about message libraries, see "Problem Insights server: message library reference" on page 86.

Tip: If you want to forward event notifications for insights from both machine learning and Elastic Stack or Splunk, specify the configuration property **`forward_messages`** in the `event.config` file to include the relevant message IDs for both types of insights.

Verifying the identity of the server for a cloud-based event management service or Z ChatOps

For communication to be secure for sending events, the IBM Z Operations Analytics Problem Insights server must verify the identity of the server for the cloud-based event management service or for IBM Z ChatOps. The communication between the Problem Insights server and any of these other servers occurs over the Hypertext Transfer Protocol Secure (HTTPS).

About this task

You must manually import a security certificate for the Problem Insights server into the `cacerts` keystore file for the Java Runtime Environment (JRE) that the Problem Insights server 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 the Problem Insights server 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`.

Procedure

To secure communication between the Problem Insights server and server for IBM Cloud Event Management, IBM Netcool Operations Insight, IBM Watson AIOps Event Manager, or IBM Z ChatOps, complete the following steps:

1. Obtain the certificate from the server keystore for IBM Cloud Event Management, IBM Netcool Operations Insight, IBM Watson AIOps Event Manager, or IBM Z ChatOps.

2. To import the server certificate from any of the previously listed servers, run the following command (all on one line):

   ```
   JRE_path/bin/keytool -import -alias alias -file
       certificate_path -keystore cacerts_path
       -storepass changeit
   ```

   Tip:

   - The default truststore password is `changeit`.
   - The `cacerts` file is typically in the path *`java.home`*`/lib/security/cacerts`. A best practice is to put the certificate in this directory also.

3. In response to the prompt `Trust this certificate`, type Yes.
4. Restart the Problem Insights server, as described in "Starting the Problem Insights server and the message broker" on page 32.

Chapter 5. Deploying the machine learning system

If you want to use machine learning to gain more extensive actionable insights (including trends and patterns of activity), install and configure the IBM Z Operations Analytics machine learning system.

About this task

This task includes the following steps:

1. Plan for deployment.

   See the following information:

   - "Machine learning system overview" on page 47
   - "Planning for deployment of the machine learning system" on page 49
2. Deploy the machine learning system on IBM z/OS UNIX System Services.

   See "Deploying the machine learning system on IBM z/OS UNIX System Services" on page 61.

Machine learning system overview

The IBM Z Operations Analytics machine learning system leverages IBM Watson Machine Learning for z/OS to provide early insight into potential IT operations problems in z/OS-based data centers. It can detect changes in subsystem usage patterns and identify emerging problems. The goal of the machine learning system is to guide operators and subject matter experts in identifying potential problems and taking the appropriate actions to resolve these problems.

Flow of data among metric-based machine learning components

Figure 2 on page 48 illustrates the flow of data among the primary components of the IBM Z Operations Analytics machine learning system for metric-based machine learning.

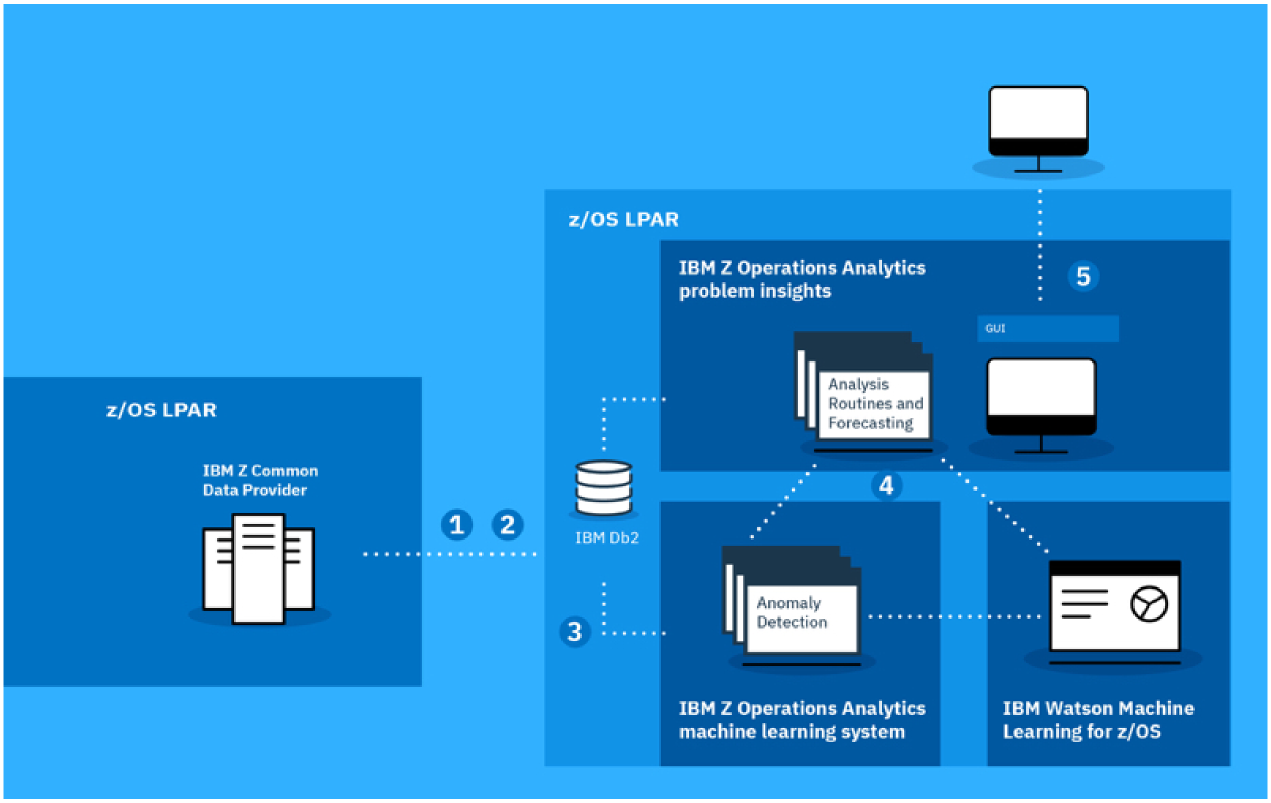
z/OS LPAR
IBM Z Common Data Provider
1
2
z/OS LPAR
IBM Z Operations Analytics problem insights
5
GUI
Analysis Routines and Forecasting
IBM Db2
4
3
Anomaly Detection
IBM Z Operations Analytics machine learning system
IBM Watson Machine Learning for z/OS

Figure 2. Flow of data among IBM Z Operations Analytics machine learning components

The following steps describe the data flow among components. The step numbers correspond to the numbers that are used in the illustration.

1. Historical data is loaded in batch mode.

   System Management Facilities (SMF) data that is stored in generation data groups (GDGs) is processed into the IBM Db2 for z/OS enterprise data warehouse through the IBM Z Common Data Provider, which is running in a single z/OS logical partition (LPAR).
2. For continuous anomaly scoring, IBM Z Common Data Provider, which is running in each LPAR from which you stream data, retrieves the data from the respective SMF source in near real-time and sends it to the IBM Db2 for z/OS enterprise data warehouse.
3. The machine learning system retrieves SMF data from the enterprise data warehouse, processes the data for model training and anomaly scoring, and stores the results in the enterprise data warehouse.
4. Anomaly scores are sent to subsystem-specific analysis routines for further analysis.
5. The IBM Z Operations Analytics Problem Insights GUI provides the capability to explore historical anomaly scores for all monitored subsystems.

Data scoring overview

By using its analysis model, the IBM Z Operations Analytics machine learning system computes anomaly scores for historical or near real-time operational data from the systems that it is monitoring.

Anomaly scores are calculated for each minute that data is available. You can view the analysis results in the Problem Insights GUI. You can also use this insight to determine the cause of past or current anomalies.

Scoring modes: batch or continuous

Scoring can be performed in the following two modes:

Batch scoring

In this mode, the machine learning system performs a one-time analysis of historical SMF data.

Continuous scoring

In this mode, the machine learning system launches a long-running scoring agent. The scoring agent analyzes newly collected SMF data at time intervals that are configurable. The default value for a time interval is 5 minutes.

Important: To make use of continuous scoring, you must configure IBM Z Common Data Provider to stream data from the category **IBM Z Operations Analytics – Machine Learning** into the enterprise data warehouse, as described in "Definition of your SMF data streams in the Configuration Tool" on page 10.

Presentation of anomaly scores in the Problem Insights GUI

In the Problem Insights GUI, the *subsystem scorecard* presents the anomaly scores for key performance indicators (KPIs) over a period of time, where each score is visually highlighted in a color shade that represents the level of anomaly. The *anomaly level* is an indication of the degree to which a KPI value is atypical, based on the analysis model. Anomaly levels are associated with anomaly scores, as shown in Table 11 on page 49.

Table 11. Correlation of each anomaly level to its associated range of anomaly scores

Anomaly level	Associated range of anomaly scores
Normal	0 - 39
Low	40 - 89
High	90 - 100

The following definitions describe the difference between an anomaly score and a deviation score:

anomaly score

A value of 0 –100 that is assigned by a predefined rule set from IBM Z Operations Analytics or by a rule set that is defined by a user. An anomaly score is based on the most extreme value (extremely low or extremely high) for a KPI at an indicated time.

Example

An anomaly score of 0 is based on a deviation score of less than 90 for each of the last 10 intervals.

deviation score

A value of 0 –100 that indicates the variance of a KPI value from the baseline for the analysis model. This score is assigned by the IBM Z Operations Analytics machine learning system.

Planning for deployment of the machine learning system

The IBM Z Operations Analytics machine learning system must be installed on IBM z/OS UNIX System Services. The standard deployment configuration is to install the machine learning system, together with the following components, on a single z/OS LPAR: the Problem Insights server, IBM Watson Machine Learning for z/OS, IBM Open Data Analytics for z/OS, and IBM Db2 for z/OS.

Before you begin

For system requirements and other planning information for the Problem Insights server, see "Planning for deployment of the Problem Insights server" on page 15.

About this task

The components of the machine learning system are a base set of machine learning functionality and a set of Insight Packs that provide z/OS subsystem-specific capabilities. These components are always installed together, but you can select which Insight Packs you want to use in your machine learning runtime environment.

When you install the Problem Insights server and the machine learning system, the following two directory structures are created:

runtime directory (*IZOA_HOME* environment variable)

The runtime directory contains the Problem Insights server and a reference copy of the machine learning system for system-wide use. When the machine learning system is in use, this directory structure is identified by the environment variable *IZOA_HOME*.

Use the following guidelines in deciding which directory to use as the runtime directory:

- The directory must not be on a read-only file system.
- The directory must have at least 0.5 GB of storage space available.
- To avoid possible conflicts, do not use the SMP/E installation directory.

Important: With the exception of the machine learning system configuration files, do not update, delete, or move the files in the runtime directory.

For information about the machine learning system configuration files, see "Machine learning system: configuration file reference" on page 103.

instance directory (*IZOA_INSTANCE* environment variable)

The instance directory contains a copy of the machine learning system that is configured for a specific purpose (such as for the monitoring and analysis of a specific set of Db2 subsystems). When the machine learning system is in use, this directory structure is identified by the environment variable *IZOA_INSTANCE*.

Although multiple instance directories can be created, only one instance (the instance that is identified by the *IZOA_INSTANCE* environment variable) can be used by one user ID at a time.

Use the following guidelines in deciding which directory to use as the instance directory:

- The directory must not be on a read-only file system.
- The directory must have at least 10 MB of storage space available.
- To avoid possible conflicts, do not use the SMP/E installation directory or the directory that is defined as a runtime directory.

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

System requirements for the machine learning system

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

- "Operating system requirements" on page 51
- "Prerequisite software" on page 51
- "Data storage, memory, and CPU requirements" on page 52

Operating system requirements

Table 12. Operating systems on which the machine learning system can run

Operating system	Version
IBM z/OS UNIX System Services **Restriction:** The machine learning system can be run only on IBM z/OS UNIX System Services. Also, the machine learning system and the Problem Insights server must be located on the same computer system. Therefore, if you plan to use the machine learning system, the Problem Insights server must be installed on IBM z/OS UNIX System Services.	IBM z/OS 2.2 or 2.3, with APAR OA56570, or IBM z/OS 2.4

Tip: A best practice is to install IBM Z Operations Analytics in the same z/OS system where you install the IBM Db2 for z/OS enterprise data warehouse and IBM Watson Machine Learning for z/OS.

Prerequisite software

Table 13 on page 51 indicates the software that must be operational for the machine learning system to run.

Table 13. Runtime requirements for the machine learning system

Runtime requirement	Version
IBM Z Common Data Provider For planning, installation, and configuration information, see Chapter 3, “Z Common Data Provider: Planning for installation and configuration,” on page 7.	2.1
IBM Db2 for z/OS, which serves as the enterprise data warehouse for operational data and the data that results from the analysis of operational data. For installation information, see Installing and migrating IBM Db2 for z/OS.	11 or later
IBM Watson Machine Learning for z/OS, which is an end-to-end enterprise machine learning platform that helps you create, train, and deploy machine learning models to extract value from your mission critical data on IBM Z, while keeping the data where it resides. For installation information, see the IBM Watson Machine Learning for z/OS documentation.	2.1.0.3, 2.2.0, or 2.2.1

Table 13. Runtime requirements for the machine learning system (continued)

Runtime requirement	Version
IBM Open Data Analytics for z/OS, which provides an extensive set of open source runtime environments and analytics capabilities (including Apache Spark) that IBM Watson Machine Learning for z/OS uses. For installation information, see Abstract for IBM Open Data Analytics for z/OS Installation and Customization Guide. Bash, Zip, Python are also required by the machine learning system, and are included with IBM Open Data Analytics for z/OS.	1.1, with the required program temporary fixes (PTFs) for your version of IBM Watson Machine Learning for z/OS. For more information about the required PTFs, see https://www.ibm.com/support/pages/node/608273.
IBM SDK for Node.js, which is an extended implementation of the well-known Node.js Javascript runtime. IBM SDK for Node.js is a separately orderable prerequisite for the IBM Watson Machine Learning for z/OS administrative web user interface.	• **For IBM Watson Machine Learning for z/OS 2.1.0.3:** Version 8, at maintenance level 8.16.2 or later • **For IBM Watson Machine Learning for z/OS 2.2.0:** Version 12, at maintenance level 12.15.0 or 12.16.1 • **For IBM Watson Machine Learning for z/OS 2.2.1:** Version 12, at maintenance level 12.15.0, 12.16.1, 12.8.0, or 12.8.4 **Important:** The IBM Z Operations Analytics machine learning system does not support other versions of IBM SDK for Node.js that are not listed here.
64-bit Java Runtime Environment (JRE)	IBM SDK for z/OS, Java Technology Edition 8, with Service Refresh (SR) 5 Fix Pack 20 or a later update
IBM z/OS Integrated Cryptographic Service Facility (ICSF), which is a software element of z/OS that works with the hardware cryptographic feature and the Security Server (RACF) to provide secure, high-speed cryptographic services. ICSF provides the application programming interfaces through which applications request the cryptographic services.	Not applicable (element of the z/OS system)

Data storage, memory, and CPU requirements

Data storage, memory and CPU requirements for the machine learning system are primarily based on the requirements for IBM Watson Machine Learning for z/OS. The machine learning system environment also requires a storage management subsystem (SMS) data class with extended addressability.

For planning information, see the IBM Watson Machine Learning for z/OS documentation. Especially, see the information about planning the system capacity for the IBM Watson Machine Learning for z/OS base in Planning system capacity for WML for z/OS base.

SMF record collection requirements

For the IBM Z Operations Analytics machine learning system, the SMF records that are used to produce analysis models and scores must be generated in 1-minute intervals.

Configuration planning for the machine learning system

During the installation of IBM Z Operations Analytics, the setup script `izoa-setup.sh` creates the configuration file `izoaml.config` and populates it with values for some configuration properties. Before you create a machine learning system instance, all properties in the `izoaml.config` file must have a specified value, unless otherwise stated in a comment in the file. For some properties, default values are specified, but you must verify that these values are appropriate for your environment, and change them if they are not.

After you verify the configuration property values in the `izoaml.config` file, you must run the `izoa-setup.sh` script to create an instance of the machine learning system and populate that instance with these values.

- "Machine learning system instances" on page 53
- "Machine learning system CLI" on page 53
- "Enterprise data warehouse configuration" on page 53
- "Preparation that must be done before a machine learning instance is created" on page 53

Machine learning system instances

You can have multiple instances of the machine learning system. For example, you might want to use multiple instances to separate development environments from production environments. Each instance uses its own IBM Db2 for z/OS database for its enterprise data warehouse.

Machine learning system CLI

To operate the machine learning system, you use a command-line interface (CLI). For information about the commands, see "Machine learning system: command reference" on page 91.

A CLI command can override its associated configuration property in the `izoaml.config` file.

Enterprise data warehouse configuration

The IBM Db2 for z/OS enterprise data warehouse is used for storing the operational data for the IBM Z Operations Analytics machine learning system. It is automatically created by the `izoa-setup.sh` script as part of deploying the machine learning system. However, before a machine learning system instance is created, you must complete some tasks to prepare for the automatic creation of the enterprise data warehouse. You must also verify that the values of associated configuration properties in the `izoaml.config` file are appropriate for your environment. Table 14 on page 54 describes the preparation tasks and lists associated configuration properties.

Tip: The system on which the enterprise data warehouse is installed is identified by the configuration property **`ZOS_HOST`** in the `izoaml.config` file.

Preparation that must be done before a machine learning instance is created

Table 14 on page 54 describes tasks to do or information to gather before a machine learning instance is created. It also indicates the configuration properties in the `izoaml.config` file that are associated with the task or information.

Tip: For descriptions of configuration properties, see "Machine learning system: command reference" on page 91, which includes descriptions of the command parameters that correspond to the configuration properties. You set these configuration properties as part of step "8.b" on page 64 of "Deploying the machine learning system on IBM z/OS UNIX System Services" on page 61.

Table 14. Tasks to do or information to gather before a machine learning instance is created

Task to do or information to gather	**Associated configuration parameter in `izoaml.config` file**
SMS data class with extended addressability: For the machine learning system, you must have a storage management subsystem (SMS) data class with extended addressability. For more information, see "Loading historical SMF data" on page 69.	**`IZOA_DATACLAS`**
Db2 subsystem that hosts the enterprise data warehouse database: Before the machine learning system instance is created, the Db2 subsystem that is to host the enterprise data warehouse database must exist or must be created by a Db2 for z/OS administrator.	**`IZOA_DBSID`**
Port on which the Db2 subsystem listens for connection requests: Before the machine learning system instance is created, determine the port on which the Db2 subsystem that hosts the enterprise data warehouse listens for connection requests.	**`IZOA_DBPORT`**
Name for the enterprise data warehouse database: Before the machine learning system instance is created, decide on the name that you want to use for the enterprise data warehouse database. The name must comply with all Db2 for z/OS rules for a database name, and no existing database can have this name. ⚠ **Attention:** The enterprise data warehouse database must be ***dedicated to the use of the IBM Z Operations Analytics machine learning system***. Do not use the same database that you use for IBM Watson Machine Learning for z/OS or for any other purpose. If you use the enterprise data warehouse database for other purposes also, data loss might occur for those other purposes during the regular maintenance operations for the IBM Z Operations Analytics machine learning system.	**`IZOA_DATABASE`**

Table 14. Tasks to do or information to gather before a machine learning instance is created (continued)

Task to do or information to gather	Associated configuration parameter in `izoaml.config` file
Names of each schema for the enterprise data warehouse database: Before the machine learning system instance is created, decide on the names that you want to use for the following schemas: • The schema in the enterprise data warehouse database that contains the tables and other objects that are related to loaded and streamed data. • The schema in the enterprise data warehouse database that contains the analysis models and scoring results for the machine learning system. You can use the same schema for both purposes, or you can use a unique schema for each purpose. **Tip:** The default schema name for both schemas is `IZOAML1`. These schema names must comply with all Db2 for z/OS rules for a schema name.	• **`IZOA_SCHEMA`** for the tables and other objects that are related to loaded and streamed data • **`IZOA_ML_SCHEMA`** for the analysis models and scoring results
Buffer pools for the enterprise data warehouse database: Before the machine learning system instance is created, the following buffer pools must be assigned for use exclusively by the enterprise data warehouse database: • The buffer pool to be used for caching all tables in the enterprise data warehouse database as they are read from disk. The name of this buffer pool is the value of the **`BUFFERPOOL`** parameter in the Db2 **`CREATE DATABASE`** command. • The buffer pool to be used for caching all indexes in the enterprise data warehouse database as they are read from disk. The name of this buffer pool is the value of the **`INDEXBP`** parameter in the Db2 **`CREATE DATABASE`** command. Use a buffer pool with a page size of 32K, and use separate buffer pools for tables and indexes. **Tip:** The default buffer pool name for both table and index buffer pools is BP32K. **Remember:** To enable buffer pool isolation at database creation time, the values of the **`IZOA_BP`** and **`IZOA_INDEXBP`** parameters in the `izoaml.config` file must designate buffer pools that are not used by other databases.	• **`IZOA_BP`** for the buffer pool that is used to cache all tables in the enterprise data warehouse database as they are read from disk • **`IZOA_INDEXBP`** for the buffer pool that is used to cache all indexes in the enterprise data warehouse database as they are read from disk

Table 14. Tasks to do or information to gather before a machine learning instance is created (continued)

Task to do or information to gather	Associated configuration parameter in `izoaml.config` file
SMS storage group for the enterprise data warehouse database: Before the machine learning system instance is created, a storage management subsystem (SMS) storage group for storing all table spaces in the enterprise data warehouse database must exist or must be created by a Db2 for z/OS administrator. The name of this storage group is the value of the **STOGROUP** parameter in the Db2 **CREATE DATABASE** command. **Tip:** The default storage group name is SYSDEFLT, which is the default storage group that is defined when IBM Db2 for z/OS is installed.	**IZOA_STOGROUP**
Encoding scheme for data in the enterprise data warehouse database: Before the machine learning system instance is created, determine the encoding scheme for data that is stored in the enterprise data warehouse. The encoding scheme applies to the table spaces that contain loaded and streamed data. The name of this scheme is the value of the **CCSID** parameter in the Db2 **CREATE DATABASE** command. The table spaces that contain machine learning system analysis models and scoring results use the Unicode encoding scheme. **Tip:** The default encoding scheme is EBCDIC.	**IZOA_CCSID**
Prefix of the SDSNLOAD and SDSNEXIT libraries in Db2 for z/OS: Before the machine learning system instance is created, determine the prefix of the SDSNLOAD and SDSNEXIT libraries in Db2 for z/OS. Db2 jobs that are run by the machine learning system must be able to load these libraries.	**IZOA_DB2HLQ**

Table 14. Tasks to do or information to gather before a machine learning instance is created (continued)

Task to do or information to gather	**Associated configuration parameter in `izoaml.config` file**
IBM Watson Machine Learning for z/OS base host name: Before the machine learning system instance is created, determine the IBM Watson Machine Learning for z/OS base host name, which is the fully qualified domain name or IP address for the system on which IBM Watson Machine Learning for z/OS is installed. For more information, see Configuring WML for z/OS base in the IBM Watson Machine Learning for z/OS documentation. This base host name is also specified by the **`MLZ_HOST_IP`** property in the `$`*`IML_HOME`*`/iml-portal/mmd_on_z.cfg` file on the system where IBM Watson Machine Learning for z/OS is installed. To access the contents of this file, you might need to contact the administrator who installed and configured IBM Watson Machine Learning for z/OS.	**`MLZ_BASE_HOST_NAME`**
IBM Watson Machine Learning for z/OS base core services port: Before the machine learning system instance is created, determine the IBM Watson Machine Learning for z/OS base core services port. This port is defined during the UI and Core Services configuration of IBM Watson Machine Learning for z/OS, as described in Configuring WML for z/OS base in the IBM Watson Machine Learning for z/OS documentation. **Tip:** The default value for this port is 11442. For more information, see Configuring ports for WML for z/OS base in the IBM Watson Machine Learning for z/OS documentation. This base core services port is also specified by the **`MLZ_SER_PORT_HTTPS`** property in the `$`*`IML_HOME`*`/iml-portal/mmd_on_z.cfg` file on the system where IBM Watson Machine Learning for z/OS is installed. To access the contents of this file, you might need to contact the administrator who installed and configured IBM Watson Machine Learning for z/OS.	**`MLZ_BASE_CORE_SERVICES_PORT`**

Table 14. Tasks to do or information to gather before a machine learning instance is created (continued)

Task to do or information to gather	Associated configuration parameter in `izoaml.config` file
IBM Watson Machine Learning for z/OS base UI service port: Before the machine learning system instance is created, determine the IBM Watson Machine Learning for z/OS base UI service port. This port is defined during the UI and Core Services configuration of IBM Watson Machine Learning for z/OS as described in Configuring WML for z/OS base in the IBM Watson Machine Learning for z/OS documentation. **Tip:** The default value for this port is 9888. For more information, see Configuring ports for WML for z/OS base in the IBM Watson Machine Learning for z/OS documentation. This base UI service port is also specified by the **`MLZ_WEB_PORT_HTTPS`** property in the `$IML_HOME/iml-portal/mmd_on_z.cfg` file on the system where IBM Watson Machine Learning for z/OS is installed. To access the contents of this file, you might need to contact the administrator who installed and configured IBM Watson Machine Learning for z/OS.	**`MLZ_BASE_UI_PORT`**

Roles and responsibilities for managing the machine learning system

The tasks of installation, configuration, and use of the IBM Z Operations Analytics machine learning system span different roles within the same IT environment. Depending on the organizational structure of your company, these roles might be assigned to one or multiple employees. To help you understand the tasks that must be done to manage the workflow of the machine learning system, each role is listed and described.

The machine learning system CLI is used in managing the workflow of the machine learning system.

Role descriptions

system administrator

- Ensures that two dedicated TSO user IDs are created.

  After the installation of the machine learning system, these user IDs must be specified as the values for the following configuration parameters in the `izoaml.config` file:

  - **`ZOS_ADM_USER`**
  - **`ZOS_USER_ID`**
- Completes the installation and configuration of the machine learning system by using the `izoa-setup.sh` script.
- By using the IBM Watson Machine Learning for z/OS administration GUI, assigns the appropriate user ID to the Model developer role within IBM Watson Machine Learning for z/OS.

  After the installation of the machine learning system, this user ID must be specified as the value for the configuration parameter **`ZOS_USER_ID`** in the `izoaml.config` file.
- Configures and maintains the started task for the Problem Insights server.

database administrator

- Configures the Problem Insights server to interact with IBM Db2 for z/OS as the enterprise data warehouse.

IBM Z Operations Analytics administrator

- Completes the following tasks by using the IBM Z Operations Analytics machine learning system CLI, as described in Chapter 6, “Operating the machine learning system,” on page 67:
  - Loads data into the enterprise data warehouse.
  - Trains the analysis model.
  - Scores data that is loaded in batch mode.
  - Defines subsystems to model and score.
  - Deploys continuous scoring agents.
- To help users better understand the presentation of the analysis results in the Problem Insights GUI, the administrator informs these users of the parameters that are used for training and scoring.

IBM Z Common Data Provider administrator

- Manages the various aspects of IBM Z Common Data Provider configuration, including defining a policy and installing data stream definitions, protocol definitions, and update definitions.

Security administrator

- Creates and manages the administrator and user IDs that are described in “Required user IDs for managing the machine learning system” on page 59.

Required user IDs for managing the machine learning system

You must define two user IDs for managing the IBM Z Operations Analytics machine learning system. To minimize effort and reduce complexity, define these user IDs so that they can be shared across IBM Z Operations Analytics, IBM Watson Machine Learning for z/OS, and IBM Open Data Analytics for z/OS.

Administrator user ID

After the installation of the machine learning system, this user ID must be specified as the value for the configuration parameter **ZOS_ADM_USER** in the `izoaml.config` file.

Tasks performed under this user ID

- Install and run IBM Watson Machine Learning for z/OS and IBM Open Data Analytics for z/OS.
- Install IBM Z Operations Analytics, and perform related administrative tasks.
- Assign the appropriate privileges to the user ID that is specified as the value for the configuration parameter **ZOS_USER_ID** in the `izoaml.config` file.

Required privileges for this user ID

- TSO and OMVS access
- Write access to IBM Watson Machine Learning for z/OS, IBM Open Data Analytics for z/OS, and IBM Z Operations Analytics installation directories and associated subdirectories
- Write access to IBM Z Operations Analytics machine learning system instance directory and associated subdirectories
- IBM Db2 for z/OS database administration authority (DBADM)

Recommended privileges for this user ID

- Root authority so that the user with this ID can create and mount file systems and update file ownership and permissions.

Non-administrative user ID

After the installation of the machine learning system, this user ID must be specified as the value for the configuration parameter **ZOS_USER_ID** in the `izoaml.config` file.

Tasks performed under this user ID

- Create and run IBM Z Operations Analytics machine learning system instances.

Required privileges for this user ID

- TSO and OMVS access
- Read access to IBM Watson Machine Learning for z/OS, IBM Open Data Analytics for z/OS, and IBM Z Operations Analytics installation directories and associated subdirectories
- Write access to IBM Z Operations Analytics machine learning system instance directory and associated subdirectories
- Read access to the high-level qualifier where the System Management Facilities (SMF) data sets that are to be used during collection and loading of historical SMF data are located
- Read and alter access to the high-level qualifier to be used for the dynamic allocation of z/OS data sets that the IBM Z Operations Analytics machine learning system requires for its tasks.

  After the installation of the machine learning system, this high-level qualifier must be specified as the value for the configuration parameter **IZOA_HLQ** in the `izoaml.config` file.
- IBM Db2 for z/OS database user authority

Information about the privileges for IBM Watson Machine Learning for z/OS user IDs

For information about the privileges for each IBM Watson Machine Learning for z/OS user ID, see Allocating other required user IDs for WML for z/OS base in the IBM Watson Machine Learning for z/OS documentation.

Insight Packs

IBM Z Operations Analytics includes Insight Packs for IBM CICS Transaction Server for z/OS and IBM Db2 for z/OS. Depending on the insights that you want, you must enable the appropriate Insight Pack in 1) the Problem Insights server so that it can provide visualizations of the subsystem-specific insights, and 2) the machine learning system so that it can provide subsystem-specific machine learning capabilities.

Key performance indicators (KPIs)

Each Insight Pack analyzes key performance indicators (KPIs) for the subsystem to which it applies. For information about these KPIs, see the following topics:

- "KPIs that are analyzed by the IBM CICS Transaction Server for z/OS Insight Pack" on page 20
- "KPIs that are analyzed by the IBM Db2 for z/OS Insight Pack" on page 22

Installation of Insights Packs

For information about installing or updating Insight Packs in the Problem Insights server, see "Installing or updating Insight Packs" on page 33.

Deploying the machine learning system on IBM z/OS UNIX System Services

After you complete the installation steps in the *Program Directory for IBM Z Operations Analytics*, you can deploy the IBM Z Operations Analytics Problem Insights server and machine learning system on IBM z/OS UNIX System Services.

Before you begin

Review "Planning for deployment of the machine learning system" on page 49.

Procedure

1. Start a z/OS UNIX System Services session by using one of the following methods:
   - From the Interactive System Productivity Facility (ISPF), by using the **omvs** command
   - By using Telnet
   - If a Secure Shell (SSH) daemon is available on the z/OS system, by using an SSH client. Typically, for z/OS UNIX System Services operations, the use of an SSH client provides the best user experience and a high level of security.
2. Go to the file system location where hierarchical file system (HFS) artifacts were installed during SMP/E processing.

   The default location is `/usr/lpp/IBM/zoa/V4R1M0`, but your system administrator might use a different location.
3. Set and export the following environment variables.

   WLPHOME

   Set the value of this variable to the installation directory for IBM WebSphere Application Server for z/OS Liberty.

   Tip: Although IBM WebSphere Application Server for z/OS Liberty is provided as part of IBM Z Common Data Provider FMID HHBO21L, any other copy of IBM WebSphere Application Server for z/OS Liberty Version 19, or later, can be used.

   JAVA_HOME

   Set the value of this variable to the installation directory for the Java Runtime Environment (JRE) for the machine learning system. For information about the runtime requirements, see "System requirements for the machine learning system" on page 50.

   PYTHON_HOME

   Set the value of this variable to the root directory of Anaconda, which is a Python distribution. Anaconda is included with IBM Open Data Analytics for z/OS.

   PATH

   Add the following information to the value of this variable:
   - `$JAVA_HOME/bin`
   - `$PYTHON_HOME/bin`

   LIBPATH

   Add the following information to the value of this variable:
   - `$JAVA_HOME/lib/s390x`
   - `$JAVA_HOME/lib/s390x/classic`
   - `$PYTHON_HOME/lib`

   Tips for z/OS UNIX System Services environment variables:

   - To set and export an environment variable, use the following command:

```
export name="value"
```

Example

```
export JAVA_HOME="/Java/J8-0_64"
```

- To append information to a path environment variable, use the following command:

```
export name="$name:new_directory"
```

Example

```
export PATH="$PATH:$JAVA_HOME/bin:$PYTHON_HOME/bin
```

4. Verify that the following programs are available in the system path ($PATH).

   These programs are provided by either the operating system or one of the prerequisite software packages.

   - `gunzip`
   - `jar`
   - `keytool`
   - `python`

   Tip for verifying that the programs are in the path: To determine the path name or command that the shell uses to call a program, use the following command:

```
command -v command-name
```

   Example

```
command -v gunzip
```

   If no path name or command is returned, the program is not available in the system path.

   If `gunzip` or `python` is not available, the probable cause is one of the following issues:

   - The path $*PYTHON_HOME*/`bin` was not added to the *PATH* environment variable, according to the instructions in step "3" on page 61.
   - The *PYTHON_HOME* environment variable is not set correctly, according to the instructions in step "3" on page 61.
   - Anaconda, which is included with IBM Open Data Analytics for z/OS, is not installed correctly.

   If `jar` or `keytool` is not available, the probable cause is one of the following issues:

   - The path `$JAVA_HOME/bin` was not added to the *PATH* environment variable, according to the instructions in step "3" on page 61.
   - The *JAVA_HOME* environment variable is not set correctly, according to the instructions in step "3" on page 61.
   - Java is not installed correctly.

5. To start the IBM Z Operations Analytics setup, run the setup command (or one of its variations), as indicated in the following description.

   The setup command presents a set of menus to help you navigate the setup process. The menus prompt you for input. After you provide the requested input, press **Enter** to proceed to the next step.

   To launch the setup utility, run one of the following commands from the installation directory:

Setup command

```
./bin/izoa-setup.sh
```

Setup command with verbose option

To obtain more verbose output for some of the setup processes, run the setup command with the `--verbose` option, as indicated:

```
./bin/izoa-setup.sh --verbose
```

Setup command with trace option

To control the trace level for some of the setup processes, run the setup command with the `--tracelevel` option, as indicated:

```
./bin/izoa-setup.sh --tracelevel [ 0|1|2|3|4 ]
```

The default value for the `--tracelevel` option is 0.

The main menu of the setup utility is shown with the following options:

```

************************************************************
*           Z Operations Analytics Setup Menu              *
************************************************************
Select one of the following processing options:
1       Perform prerequisite check
2       Create runtime environment
3       Configure runtime environment and machine learning system instance
4       Remove machine learning system instance
5       Remove runtime environment
6       Print configuration settings
7       Deploy machine learning support into Problem Insights server
8       Remove machine learning support from Problem Insights server
9       Manage rules engine
10      Change passwords
11      Collect logs
12      Quit
```

Tip: You can process only one setup menu option each time that you run the setup command. Therefore, to process multiple options, you must run the setup command multiple times.

6. Optional: To verify that your specified runtime directory meets requirements, run a prerequisite check, as described in the following steps:

   a) From the main menu of the setup utility, select option 1, which is `Perform prerequisite check`.

   b) In the secondary menu, specify whether you want to run a check for only the IBM Z Operations Analytics machine learning system, only the IBM Z Operations Analytics Problem Insights server, or both.

   c) Provide your runtime (*IZOA_HOME*) directory.

7. Create the IBM Z Operations Analytics runtime environment, as described in the following steps:

   a) From the main menu of the setup utility, select option 2, which is `Create runtime environment`.

   b) In the secondary menu, specify whether you want to set up a runtime environment for only the IBM Z Operations Analytics machine learning system, only the IBM Z Operations Analytics Problem Insights server, or both.

   c) Provide your runtime (*IZOA_HOME*) directory.

8. Manually update the IBM Z Operations Analytics configuration files, as described in the following steps.

   Although the IBM Z Operations Analytics setup process helps you to populate the configuration files for the IBM Z Operations Analytics Problem Insights server and machine learning system, some manual updates are required after the runtime environment is created and before the runtime environment is configured.

a) Verify that the following configuration files are copied from the `$IZOA_HOME/samples` directory to the `$IZOA_HOME/config` directory:

   - For the machine learning system: `izoaml.config`

     This file is automatically copied when the runtime environment is created, if you chose to include the machine learning system in the runtime environment.

   - For the Problem Insights server: `cli.config`

     This file is automatically copied when the runtime environment is created, if you chose to include the Problem Insights server in the runtime environment.

   If these files were not automatically copied, copy them by running the following command:

   ```
   cp $IZOA_HOME/*.config $IZOA_HOME/config
   ```

b) If you are setting up a runtime environment for the machine learning system, update the file `$IZOA_HOME/config/izoaml.config` with the appropriate values for each configuration property.

   The `izoaml.config` file includes a description for each configuration property, and each property is set to a value. Change an existing value only if it is not appropriate for your installation.

   The configuration properties in `$IZOA_HOME/config/izoaml.config` correspond to the configuration parameters that are described in "Machine learning system: command reference" on page 91.

   Important: Do not change the value of any configuration property that is designated as "automatically set by the `izoa-setup.sh` script."

c) If you are setting up a runtime environment for the Problem Insights server, update the file `$IZOA_HOME/config/cli.config` with the appropriate values for each configuration property. However, do not change the default passwords in the configuration file at this time. You are later prompted to provide the required passwords, and these passwords are then automatically encrypted before they are written to the configuration file.

d) For each configuration file, save your updates.

Important: The configuration files are encoded in ISO8859-1 and are tagged in the HFS accordingly. To view the tagging of the files, run the following command from the runtime (*IZOA_HOME*) directory:

```
ls -T *.config
```

To update these files, you must use a file editor that is capable of reading and writing ISO8859-1 data. If you plan to use the vi editor, or a similar command line editor on z/OS UNIX System Services, set the value of the *_BPXK_AUTOCVT* environment variable to ON to enable automatic codeset conversion between EBCDIC and ISO8859-1.

```
export _BPXK_AUTOCVT=ON
```

9. Configure the IBM Z Operations Analytics runtime environment, and create an instance of the machine learning system. From the main menu of the setup utility, select option 3, which is `Configure runtime environment and Machine Learning system instance`.

   This option results in the following actions:

   a. You are prompted for the required passwords, which are the passwords for the user IDs **ZOS_ADM_USER** and **ZOS_USER_ID** that are described in "Roles and responsibilities for managing the machine learning system" on page 58.

   b. The passwords are automatically encrypted before they are written to the configuration file.

   c. Configuration information from the two core configuration files (`izoaml.config` and `cli.config`) is propagated to other configuration files as needed.

d. If you chose to include the machine learning system in the runtime environment, a single instance of the machine learning system is created, based on the values of the configuration properties in the `izoaml.config` file. In this step, IBM Z Operations Analytics must communicate with IBM Db2 for z/OS and IBM Watson Machine Learning for z/OS by using the communication-related properties that are defined in the `izoaml.config` file.

Important: Depending on your selections in earlier steps of this procedure, the results of option 3 vary in the following ways:

- If you chose to create a runtime environment for both the machine learning system and the Problem Insights server, the IBM Db2 for z/OS communication properties that are required for the Problem Insights server are automatically derived from the `izoaml.config` file.
- If you chose to create a runtime environment for only the Problem Insights server, you are prompted to enter the IBM Db2 for z/OS communication properties that are required for the Problem Insights server.

Important: If you chose to configure a runtime environment for the machine learning system, you are prompted at the end of the configuration process to add the following environment variables to the `$HOME/.profile` file of any user ID under which the machine learning system is to be used:

IZOA_HOME
This is the directory into which you installed the runtime environment.

IZOA_INSTANCE
This is the directory in which you created the machine learning system instance.

To ensure that the Python scripts that are provided by the machine learning system instance function properly, you are also prompted to update the *PYTHONPATH* and *PATH* environment variables for the user IDs.

10. Run the Db2 RUNSTATS utility against the tables in the schema to be queried.

    For more information, see "Maintaining the enterprise data warehouse" on page 75.

11. Start the Problem Insights server.

    You can use either of the following methods to start and stop the Problem Insights server:

    Start and stop as a started task
    The best practice is to run the Problem Insights server as a long-running task. IBM Z Operations Analytics includes the sample procedure GLAPISRV to help you configure this started task. Use the instructions in the sample to update the started task for your environment. Then, you can start and stop the Problem Insights server in the same way that you start and stop any other started task.

    Start

    ```
    /S GLAPISRV
    ```

    Stop

    ```
    /C GLAPISRV
    ```

    Start and stop from z/OS UNIX System Services
    If you prefer to start and stop the Problem Insights server from z/OS UNIX System Services, you can use the `analysis.sh` script that is provided in the runtime (*IZOA_HOME*) directory.

    Start

    ```
    ./analysis.sh start
    ```

    Stop

    ```
    ./analysis.sh stop
    ```

    If you plan to use this script to control the Problem Insights server, you must comment out the following configuration properties in the `cli.config` file:

- Comment out the property that defines that the use of an IBM WebSphere Application Server for z/OS Liberty angel process is required: `com.ibm.ws.zos.core.angelRequired`
- Comment out the property that defines the name of the IBM WebSphere Application Server for z/OS Liberty angel process to use: `com.ibm.ws.zos.core.angelName`

Tip: If you run the Problem Insights server as a started task, these properties are required.

Chapter 6. Operating the machine learning system

To use the IBM Z Operations Analytics machine learning system to gain insight into your IT operations environment, you must first load historical System Management Facilities (SMF) data into the IBM Db2 for z/OS enterprise data warehouse, and train the analysis model. Then, new data can be analyzed and scored in comparison to the analysis model.

About this task

The following topics contain important reference information:

- "Machine learning system: command reference" on page 91
- "Machine learning system: configuration file reference" on page 103

To collect and load historical SMF data, train the analysis model, and analyze and score new SMF data, you must use the IBM Z Operations Analytics machine learning system CLI. "Machine learning system: command reference" on page 91 contains information about the commands and associated parameters for each task.

The CLI commands depend on the execution of z/OS jobs and SQL scripts. The resulting logs are saved in the following directories for reference by the IBM Z Operations Analytics administrator:

z/OS job logs
`$IZOA_INSTANCE/log/jcl`

SQL script logs
`$IZOA_INSTANCE/log/sql`

Other related logs
`$IZOA_INSTANCE/log`

Procedure

1. Load historical SMF data, as described in "Loading historical SMF data" on page 69.
2. Train the analysis model, as described in "Training the analysis model" on page 70.
3. Analyze and score new SMF data, as described in "Analyzing and scoring new SMF data" on page 72.

Requirements and guidelines for the data that is loaded

Review these requirements and guidelines for getting the appropriate IBM CICS Transaction Server for z/OS and IBM Db2 for z/OS insights, for loading the historical data, and for continuous scoring of streaming data.

Requirements for getting the appropriate CICS and Db2 insights

To obtain insights into IBM CICS Transaction Server for z/OS subsystems, you must collect SMF record type 110 data. Also, if you want to limit the amount of CICS Transaction Server for z/OS data that is processed by the machine learning system, see the instructions in "Filtering of CICS transactions for machine learning" on page 68.

To obtain insights into IBM Db2 for z/OS subsystems, you must collect SMF record type 100 data.

For information about the key performance indicators (KPIs) that are analyzed for these subsystems, see "Insight Packs" on page 20.

Requirements and guidelines for the historical data

The machine learning system expects historical SMF data to be in generation data groups (GDGs).

For best results, the machine learning models should be trained with 30 days of historical data. Although training can be performed with a smaller set of historical data, the quality of the model might be reduced.

The data must be contiguous, which means that it must cover all days in the selected time period and all hours within each day. Gaps in the historical data result in less reliable models because the model does not adequately reflect normal variations in system behavior on different days, or at different times of a specific day of the week.

Important: The collection and loading of one or more days of production SMF data might require very large data set allocations and might consume significant processing time.

Requirements and guidelines for continuous scoring of streaming data

If you plan to use continuous scoring, you must install IBM Z Common Data Provider on each managed z/OS system. Also, each instance of IBM Z Common Data Provider must be configured with a policy for streaming data from the Configuration Tool category **IBM Z Operations Analytics - Machine Learning** to the IBM Z Operations Analytics enterprise data warehouse, as described in "Definition of your SMF data streams in the Configuration Tool" on page 10. For information about installing and configuring IBM Z Common Data Provider, see the IBM Z Common Data Provider documentation.

Filtering of CICS transactions for machine learning

To limit the amount of CICS Transaction Server for z/OS data that is processed by the IBM Z Operations Analytics machine learning system, you can restrict the transactions that are being monitored to those that are most important to your environment. To do this, configure filters for both initial data collection and real-time data streaming.

Transaction filtering for initial data collection

To filter transactions in the initial data collection, you must update the file `cicshealth_GLAAKCUN.jcl`, as described in the following steps:

1. In the `features/cicshealth/jcl` directory of the directory where IBM Z Operations Analytics is installed, edit the file `cicshealth_GLAAKCUN.jcl`.
2. Update the `TRAN_FILTER` statement as appropriate. For example, to track only the transactions ESDA or `CWXN`, use the following code:

```
SET TRAN_FILTER = 'AND (TRAN = ''EDSA'' OR TRAN = ''CWXN'')';
```

Tip: If you add transaction filters, the value of the `TRAN_FILTER` statement must start with the `'AND (` qualifier.

Transaction filtering for real-time data streaming

To filter transactions in the real-time data streaming, you must update members of the IBM Z Operations Analytics data set that is referenced in the `concats.json` file that is in the working directory for the IBM Z Common Data Provider Configuration Tool.

Tip: `"IZOA" : "ZOA.V4R1M0.SGLADEFS"` is an example of this data set definition in the `concats.json` file. For more information about the `concats.json` file, see "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

The following steps describe how to update the data set members:

1. Edit the members `GLA110MT`, `GLA110PT`, and `GLA110TT`.
2. In each of these members, update the `DEFINE UPDATE` section to include a filter that is based on transaction IDs. For example, in the `GLA110MT` member, to include a filter to send data only for the transactions ESDA or `CWXN`, use the following code to include these transaction IDs in the `WHERE` logic:

```
DEFINE UPDATE A_KC_MON_TRAN_I
```

```
   VERSION 'IZOA.V410'
    FROM SMF_CICS_T
    WHERE SMFMNRVN >= '0650'
    AND (TRAN = 'EDSA' OR TRAN = 'CWXN')
    TO &IBM_UPDATE_TARGET
    AS &IBM_FILE_FORMAT
```

3. To have the member updates take effect, either create a new policy, or open and save your existing policy in the Configuration Tool.

Loading historical SMF data

Collect your historical System Management Facilities (SMF) data, and load it into the IBM Db2 for z/OS enterprise data warehouse for analysis by the IBM Z Operations Analytics machine learning system.

Before you begin

Review "Requirements and guidelines for the data that is loaded" on page 67.

Restriction: If multiple end users must perform data collection and loading activities simultaneously, each user must define and use a unique user ID. The use of shared user IDs can cause Job Entry Subsystem (JES) queue conflicts, which can result in in failed job submissions.

About this task

The **database=collect** and **database=load** commands direct IBM Z Common Data Provider to run and load the historical SMF data for each subsystem. This process populates the database schema that is defined in the izoaml.config file.

For more information about the commands and associated parameters that are mentioned in this procedure, see "database command" on page 92.

Important:

If the configuration property **IZOA_DATACLAS** in the izoaml.config file does not have an associated default **STORCLAS** parameter in the SMS automatic class selection (ACS) routines, the COLLECT and LOAD jobs cannot allocate data sets.

Procedure

1. Collect the data by running the **database=collect** command, as shown in the following example:

   ```
   $IZOA_HOME/bin/izoaml.py --feature=db2health --database=collect
   --run --IZOA_SMFGDG=AN.SMF.GDG.DATASET
   ```

   In the **database=collect** command, the parameter **IZOA_SMFGDG** is required, and *AN.SMF.GDG.DATASET* represents the name of the relevant SMF generation data group (GDG) data set.

   Important: The **database=collect** command must be run for each SMF GDG data set. Therefore, if more than one GDG data set must be processed, **database=collect** must be run repeatedly until the data from all data sets is collected.
2. If data collection is successful, load the data by running the **database=load** command, as shown in the following example:

   ```
   $IZOA_HOME/bin/izoaml.py --feature=db2health --database=load --run
   ```
3. After all data sets are loaded, run the following commands to prepare the enterprise data warehouse for continuous data loading and analysis.

a. Create a backup copy of the database tables in the enterprise data warehouse by running the **database=copy** command, as shown in the following example:

```
$IZOA_HOME/bin/izoaml.py --feature=db2health --database=copy --run
```

b. Prepare the enterprise data warehouse for concurrent read and write operations by running the **database=concurrent** command, as shown in the following example:

```
$IZOA_HOME/bin/izoaml.py --feature=db2health --database=concurrent --run
```

Training the analysis model

Based on the historical System Management Facilities (SMF) data that you loaded, the IBM Z Operations Analytics machine learning system defines an analysis model of normal system behavior and uses this model as a baseline for detecting anomalies in your Z environment. To train the analysis model, you must define the subsystems that you want to analyze to the machine learning system.

Before you begin

Load historical SMF data, as described in "Loading historical SMF data" on page 69.

About this task

For more information about the commands and associated parameters that are mentioned in this procedure, see the following topics:

- "define command" on page 96
- "ml command" on page 100
- "status command" on page 102

Procedure

1. To provide the list of subsystem IDs for the subsystems that you want to analyze, run the **define** command with the appropriate options and parameters, as shown in the following examples.

   Tips:

   - You can define the subsystems to the machine learning system in either of the following ways:
     - Use the **define=create** command to define a subsystem explicitly.
     - Use the **define=find** command to have the system search the loaded data for subsystems and create associated subsystem definitions for you.
   - At a minimum, you must run the following commands:
     - `--define=find` (to search the historical data for subsystems to analyze)
     - `--define=update --IZOA_SYSPLEX_NAME=`*`sysplex_name`* (to provide the enterprise data warehouse with the sysplex name for the subsystems that are to be analyzed)
     - If the time zone of the IBM z/OS UNIX System Services environment in which the machine learning system CLI runs does not match the time zone of the LPARs with the data that is to be analyzed, you must run the following command:

       ```
       --define=update --IZOA_TIMEZONE=time_zone
       ```

   Define a subsystem: define=create command

   ```
   $IZOA_HOME/bin/izoaml.py --feature=db2health --define=create --run
   --IZOA_SYS_ID=MVS system ID for system being monitored and analyzed
   --IZOA_SUBSYS_ID=Subsystem ID for system being monitored and analyzed
   ```

```
--IZOA_TIMEZONE=Time zone of system specified in IZOA_SYS_ID
--IZOA_SYSPLEX_NAME=Sysplex name for system being monitored and analyzed
--IZOA_SYSTEM=System name for system being monitored and analyzed
```

The following parameters are required:

- **IZOA_SYS_ID**
- **IZOA_SUBSYS_ID**
- **IZOA_TIMEZONE**

Search for subsystems: `define=find` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health --define=find --run
--IZOA_SYS_ID=MVS system ID for system being monitored and analyzed
--IZOA_SUBSYS_ID=Subsystem ID for system being monitored and analyzed
--IZOA_START_DATE=Generic start time stamp for monitored system discovery
--IZOA_END_DATE=Generic end time stamp for monitored system discovery
--IZOA_TIMEZONE=Time zone of system specified in IZOA_SYS_ID
--IZOA_SYSPLEX_NAME=Sysplex name for system being monitored and analyzed
--IZOA_SYSTEM=System name for system being monitored and analyzed
```

This command directs the machine learning system to search for subsystems and related SMF data that match the criteria that is specified by the following parameters:

- **IZOA_SYS_ID**
- **IZOA_SUBSYS_ID**
- **IZOA_START_DATE**
- **IZOA_END_DATE**

Update a subsystem definition: `define=update` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health --define=update --run
--IZOA_SYS_ID=MVS system ID for system being monitored and analyzed
--IZOA_SUBSYS_ID=Subsystem ID for system being monitored and analyzed
--IZOA_SUBSYS_STATUS=Subsystem status for system being monitored and analyzed
--IZOA_USER_LABEL=User label for saving model in IBM Watson Machine Learning for z/OS
--IZOA_TIMEZONE=Time zone of system specified in IZOA_SYS_ID
--IZOA_SYSPLEX_NAME=Sysplex name for system being monitored and analyzed
--IZOA_SYSTEM=System name for system being monitored and analyzed
```

Delete a subsystem definition: `define=delete` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health --define=delete --run
--IZOA_SYS_ID=MVS system ID for system being monitored and analyzed
--IZOA_SUBSYS_ID=Subsystem ID for system being monitored and analyzed
--IZOA_USER_LABEL=User label for saving model in IBM Watson Machine Learning for z/OS
```

2. Run the **`ml=train`** command, as shown in the following example.

```
$IZOA_HOME/bin/izoaml.py
--feature=db2health --ml=train --run
--IZOA_DATA_TIME_RANGE=Specified time range of data to be used for training
--IZOA_SYS_ID=MVS system ID for system being monitored and analyzed
--IZOA_SUBSYS_ID=Subsystem ID for system being monitored and analyzed
```

The following parameter is required:

- **IZOA_DATA_TIME_RANGE** in format "*YYYYMMDD-HHMMSS;YYYYMMDD-HHMMSS*"

  The quotation marks are required also.

Although the following parameters are optional, specify one or both of them if you want to train the analysis model for a specific system or subsystem:

- **IZOA_SYS_ID**
- **IZOA_SUBSYS_ID**

3. Optional: To monitor the status of scoring and training jobs, use the **status** command, as shown in the following example.

   View the status of a specific job

   ```
   $IZOA_HOME/bin/izoaml.py
   --status=submission ID of the job to be viewed
   ```

   View the status of all jobs

   ```
   $IZOA_HOME/bin/izoaml.py
   --status=all
   --JOB_MAX=maximum number of jobs to show
   --JOB_TYPE=show only jobs of this type
   --JOB_STATUS=show only jobs with this status
   ```

   Although the following parameters are optional, specify one or more of them if you want to limit the number of jobs that are shown:

   - **JOB_MAX**

     If you specify the maximum number of jobs to show, the most recently submitted job is shown first.

   - **JOB_TYPE**

     The following values are valid:

     - score
     - train

   - **JOB_STATUS**

     The following values are valid:

     - CANCELED
     - COMPLETED
     - FAILED
     - QUEUED
     - RUNNING
     - SUBMITTED

Analyzing and scoring new SMF data

By using its analysis model, the IBM Z Operations Analytics machine learning system computes anomaly scores for historical or near real-time operational data from the systems that it is monitoring. To start the scoring, you must run the appropriate commands, depending on whether you want to score the data only once or score it continuously.

Before you begin

Train the analysis model, as described in "Training the analysis model" on page 70.

Install IBM Z Common Data Provider on each managed z/OS system. Also, configure each instance of IBM Z Common Data Provider with a policy for streaming data from the Configuration Tool category **IBM Z Operations Analytics - Machine Learning** to the IBM Z Operations Analytics enterprise data warehouse, as described in "Definition of your SMF data streams in the Configuration Tool" on page 10.

If, in the loading of historical SMF data for initial data collection, you created filters to limit the amount of CICS Transaction Server for z/OS data that is processed by the machine learning system, you must also apply the same filters to new SMF data. For more information, see the instructions in "Filtering of CICS transactions for machine learning" on page 68.

About this task

For more information about scoring, see "Data scoring overview" on page 48.

For more information about the commands and associated parameters that are mentioned in this procedure, see the following topics:

- "ml command" on page 100
- "status command" on page 102

Procedure

1. Depending on whether you want to score the data only once or score it continuously, run either the **ml=batchscore** or **ml=scag_batchscore** command, as shown in the following examples.

   Tip: If you need to customize the configuration properties that are used for scoring, copy the IZOA_*subsystem*HEALTH_ML.properties file for the relevant subsystem type (for example, the `IZOA_DB2HEALTH_ML.properties` file) from the `samples` directory to the `config` directory, and update it. When you run the scoring command, you must reference the path to this file.

   Score the incoming SMF data once: ml=batchscore command

   ```
   $IZOA_HOME/bin/izoaml.py
   --feature=db2health --ml=batchscore --run
   --IZOA_DATA_TIME_RANGE=Specified time range of data to be scored
   --IZOA_SYS_ID=MVS system ID for system being monitored and analyzed
   --IZOA_SUBSYS_ID=Subsystem ID for system being monitored and analyzed
   --IZOA_ML_PROP_FILE=Path to IZOA_DB2HEALTH_ML.properties or IZOA_CICSHEALTH_ML.properties
   ```

   The following parameter is required:

   - **IZOA_DATA_TIME_RANGE** in format *YYYYMMDD-HHMMSS;YYYYMMDD-HHMMSS*

   Although the following parameters are optional, specify one or both of them if you want to score data for a specific system or subsystem:

   - **IZOA_SYS_ID**
   - **IZOA_SUBSYS_ID**

   Score the incoming SMF data continuously: ml=scag_batchscore command

   ```
   $IZOA_HOME/bin/izoaml.py
   --feature=db2health --ml=scag_batchscore --run
   --IZOA_SYS_ID=MVS system ID for system being monitored and analyzed
   --IZOA_SUBSYS_ID=Subsystem ID for system being monitored and analyzed
   --IZOA_SCORE_FREQ=Scoring interval in seconds
   --IZOA_ML_PROP_FILE=Path to IZOA_DB2HEALTH_ML.properties or IZOA_CICSHEALTH_ML.properties
   ```

   With continuous scoring, you specify the scoring interval (in seconds), which determines how often the data scoring process is run. However, regardless of the scoring interval, distinct scoring results are produced for each minute. For example, if your scoring interval is 300 seconds (5 minutes), one scoring run occurs every 5 minutes, and each scoring run produces score values for each minute within that 5-minute interval.

   To enable the forwarding of continuous scoring results to the rules engine of the Problem Insights server, and to enable integration with an event management system, the following conditions must be met:

   - The **ml=scag_batchscore** command must be run with the **IZOA_ML_PROP_FILE** parameter. The value of the **IZOA_ML_PROP_FILE** parameter is the fully qualified path to the `IZOA_DB2HEALTH_ML.properties` or `IZOA_CICSHEALTH_ML.properties` file (depending on the subsystem type that is supported by the scoring agent).

     Important: The **IZOA_USER_LABEL** parameter must ***not*** be included on the **ml=scag_batchscore** command.

- The following lines must be uncommented in the respective files:

  IZOA_DB2HEALTH_ML.properties file

  ```
  STAGE.isNotificationEnabled.Notification.system = true
  STAGE.isNotificationEnabled.Notification.latch = true
  STAGE.isNotificationEnabled.Notification.storage = true
  ```

  IZOA_CICSHEALTH_ML.properties

  ```
  STAGE.isNotificationEnabled.Notification.a_kc_mon_tran = true
  STAGE.isNotificationEnabled.Notification.a_kc_rmi_perf_t = true
  STAGE.isNotificationEnabled.Notification.a_kc_s_stor_dsa_t = true
  STAGE.isNotificationEnabled.Notification.a_kc_t_tran = true
  STAGE.isNotificationEnabled.Notification.a_kc_t_tran_t_sum = true
  ```

Although the following parameters are optional, specify one or both of them if you want to score data for a specific system or subsystem:

- **IZOA_SYS_ID**
- **IZOA_SUBSYS_ID**

Stop the scoring agent: ml=scag_stop command

```
$IZOA_HOME/bin/izoaml.py
--feature=db2health --ml=scag_stop --run
```

2. Optional: To monitor the status of scoring and training jobs, use the **status** command, as shown in the following example.

View the status of a specific job

```
$IZOA_HOME/bin/izoaml.py
--status=submission ID of the job to be viewed
```

View the status of all jobs

```
$IZOA_HOME/bin/izoaml.py
--status=all
--JOB_MAX=maximum number of jobs to show
--JOB_TYPE=show only jobs of this type
--JOB_STATUS=show only jobs with this status
```

Although the following parameters are optional, specify one or more of them if you want to limit the number of jobs that are shown:

- **JOB_MAX**

  If you specify the maximum number of jobs to show, the most recently submitted job is shown first.

- **JOB_TYPE**

  The following values are valid:

  - score
  - train

- **JOB_STATUS**

  The following values are valid:

  - CANCELED
  - COMPLETED
  - FAILED

    - QUEUED
    - RUNNING
    - SUBMITTED

Maintaining the enterprise data warehouse

To ensure optimal performance, you must maintain the IBM Db2 for z/OS enterprise data warehouse for the IBM Z Operations Analytics machine learning system.

About this task

For more information about the commands and associated parameters that are mentioned in this procedure, see "database command" on page 92.

Procedure

- **Run the Db2 RUNSTATS utility when necessary**, as described in the following information.
    - Under the following conditions, run RUNSTATS commands against the *historical data tables* (`A_%`) in the enterprise data warehouse database:
        - After the initial data load is complete.
        - If operational data is continuously streamed into historical data tables, run the commands at least nightly.

      Tip: To run RUNSTATS against only the historical data tables, use the **`--database=stat_history`** command.
    - Under the following conditions, run RUNSTATS commands against the *machine learning data tables* (`CICSHEALTH_%, DB2HEALTH_%`) in the enterprise data warehouse database:
        - After first-time training of the system and before first-time scoring of data.
        - After each subsequent training of the system
        - After first-time scoring of data and before first-time access to the scoring results from the Problem Insights server UI.
        - After each subsequent scoring of data
        - If operational data is continuously scored in machine learning data tables, run the commands at least nightly.

      Tip: To run RUNSTATS against all the machine learning data tables, use the **`--database=stat_ml`** command.
    - To run RUNSTATS against both historical and machine learning data tables in a single process, use the **`--database=stat_all`** command.
    - More considerations for running RUNSTATS against machine learning data tables:
        - After initial training and after initial scoring, RUNSTATS must be run against all machine learning tables.
        - If, for subsequent training or scoring, you have a reason for not running RUNSTATS against all machine learning data tables, you can use the machine learning system CLI to target only specific tables. Use the following guidelines:
            - To run RUNSTATS against a single table, use the following basic syntax:

              ```
              $IZOA_HOME/bin/submit.py -z ZOS_HOST -u ZOS_USER_ID -p ZOS_USER_PWD -j
              $IZOA_INSTANCE/jcl/subsystem_TABLENAME_table_stat.job
              ```
            - After training, run RUNSTATS against the *SUBSYSTEM*`_MODEL_DETAILS` table at a minimum.
            - After scoring, run RUNSTATS against the following tables at a minimum:

    - *SUBSYSTEM*_SCORE_AGGREGATION
    - *SUBSYSTEM*_SCORE_RESULT
    - *SUBSYSTEM*_FORECAST_RESULT
  - A *SUBSYSTEM*_SCORE_RESULT table can become very large, and under extreme conditions, its size can cause RUNSTATS failures.

Use the following **database** commands to run RUNSTATS:

Run the Db2 RUNSTATS utility against the historical database tables: database=stat_history command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=stat_history —run
```

Run the Db2 RUNSTATS utility against the machine learning database tables: database=stat_ml command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=stat_ml —run
```

Run the Db2 RUNSTATS utility against all database tables: database=stat_all command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=stat_all —run
```

- **Run the Db2 REORG TABLESPACE utility when necessary**, as described in the following information.
  - For detailed guidelines, see Determining whether an object requires reorganization.
  - Run REORG commands under the following conditions:
    - If any enterprise data warehouse tables are placed in advisory REORG-pending state (AREO* or AREOR)
    - If the indexes are substantially fragmented, a reorganization of the indexes should be done.

Tips:

- The REORG jobs are provided in files with the naming convention *subsystem_table-name*_table_reorg.job, as shown in the following example:

```
db2health_DB2HEALTH_MODEL_DETAILS_table_reorg.job
```

- Because reorganization is a very expensive operation for large database tables, run these commands only if necessary.
- By default, the REORG jobs that are provided by IBM Z Operations Analytics reorganize the tablespace and indexes with the REORG SHRLEVEL NONE option.
- If there are pending changes to the enterprise data warehouse database schema, reorganization must be done with the REORG SHRLEVEL CHANGE option, and an image copy must be taken. The options for this operation are included in the REORG jobs that are provided by IBM Z Operations Analytics, but are commented out due to the significant resource requirements for this operation. If you need these options, you must uncomment them.
- Due to the size of the *SUBSYSTEM*_SCORE_RESULT tables, you might need to add more DATAWRKxx DD statements to provide more temporary space if the REORG command fails for these tables. Because of the potential for this command to fail, reorganization of these tables is set to run only if the table space is in REORG-pending state.
- Because fragmentation occurs under the following conditions, limit these conditions as much as possible:
  - Concurrent inserts into the same table from multiple IBM Z Common Data Provider data streams

    - Concurrent runs to train the system
    - Most significantly, concurrent scoring of data from multiple subsystems

  Use the following **`database`** commands to run REORG:

  Run the Db2 REORG TABLESPACE utility against the historical database tables: `database=reorg_history` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=reorg_history --run
```

  Run the Db2 REORG TABLESPACE utility against the machine learning database tables: `database=reorg_ml` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=reorg_ml –-run
```

  Run the Db2 REORG TABLESPACE utility against all database tables: `database=reorg_all` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=reorg_all –-run
```

- **Run the following miscellaneous commands when necessary** to maintain the enterprise data warehouse:

  Alter the buffer pool for the historical database tables: `database=bufferpool_history` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=bufferpool_history –-run
```

  Alter the buffer pool for the machine learning database tables: `database=bufferpool_ml` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=bufferpool_ml –-run
```

  Alter the buffer pool for all database tables: `database=bufferpool_all` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=bufferpool_all –-run
```

  Check the buffer pool in use for tables that are associated with a specific subsystem type: `database=sql` command

```
$IZOA_HOME/bin/izoaml.py --feature=db2health
--database=sql --IZOA_FN=db2health_tables_list.sql
–-runcat $IZOA_INSTANCE/sql/db2health_tables_list.sql.log
```

  Run a job to display database information

```
$IZOA_HOME/bin/submit.py -z ZOS_HOST
-u ZOS_USER_ID -p ZOS_USER_PWD
-j $IZOA_INSTANCE/jcl/mlbase_database_display.job
```

  Run a job to display buffer pool information

```
$IZOA_HOME/bin/submit.py -z ZOS_HOST
-u ZOS_USER_ID -p ZOS_USER_PWD
-j $IZOA_INSTANCE/jcl/mlbase_bufferpool_display.job
```

Remove unneeded analysis result data from the enterprise data warehouse: `database=cleandata` command

```
$IZOA_HOME/lib/izoaml.py
--feature=db2health
--database=cleandata
--run
--IZOA_SYS_ID=MVS system ID for system being monitored and analyzed
--IZOA_SUBSYS_ID=Subsystem ID for system being monitored and analyzed
--IZOA_START_ENTRYTIME=Beginning of time range for data to be removed
--IZOA_END_ENTRYTIME=End of time range for data to be removed
```

Attention: The following parameters, which have the indicated default values, are optional. Ensure that these parameters are specified with the values that you want. For example, if you do not change the default value of % for the **`IZOA_SYS_ID`** or **`IZOA_SUBSYS_ID`** parameter, the analysis result data for all systems or subsystems is removed.

- **`IZOA_SYS_ID`** with default value %, which includes all systems
- **`IZOA_SUBSYS_ID`** with default value %, which includes all subsystems
- **`IZOA_START_ENTRYTIME`** with default value `2018-01-01 00:00:00`
- **`IZOA_END_ENTRYTIME`** with default value `2025-12-31 59:59:59`

Chapter 7. Tuning the machine learning system

Best practices are provided for tuning IBM Z Operations Analytics with machine learning. They are based on observations that IBM testers made after performing measurements for various functions of IBM Z Operations Analytics machine learning, including training of the analysis model, batch scoring of operational data, and continuous scoring of operational data.

About this task

Disclaimer: This information will be periodically updated as more data becomes available from continued lab testing. All performance data published herein is for reference only as it is not representative of any non-testing environment.

Reference information: The following information was used as reference by the testers who compiled these best practices:

- IBM Z Operations Analytics V4.1.0 documentation
- IBM Watson Machine Learning for z/OS V2.1.0.1 documentation
- IBM Db2 12 for z/OS Managing Performance
- IBM Z Common Data Provider V2.1.0 documentation
- IBM Z Common Data Provider and IBM Z Operations Analytics Best Practices Guide

Best practices

Best practices for tuning the IBM Z Operations Analytics machine learning system apply to 1) the IBM Db2 for z/OS enterprise data warehouse, 2) the Apache Spark component of IBM Watson Machine Learning for z/OS, and 3) training the analysis model and scoring the data.

- "IBM Db2 for z/OS" on page 79
- "Apache Spark" on page 79
- "Training the analysis model and scoring data" on page 80

IBM Db2 for z/OS

The IBM Db2 for z/OS enterprise data warehouse database is essential to the operation of the IBM Z Operations Analytics machine learning system. Input data from monitored subsystems and scoring results are stored in this Db2 database.

To keep IBM Db2 for z/OS databases running in an optimal environment, apply typical Db2 tuning techniques.

Apache Spark

IBM Z Operations Analytics uses IBM Watson Machine Learning for z/OS for its machine learning functions, and IBM Watson Machine Learning for z/OS uses Apache Spark for its runtime environment. In this context, Apache Spark uses a driver and an executor to perform machine learning functions.

You can use the following general equations for estimating resource requirements for Apache Spark:

Table 15. Estimating the resource requirements for Apache Spark

Cores =	number of subsystems × (number of driver cores per subsystem + number of executor cores per subsystem)
Memory =	number of subsystems × (amount of driver memory per subsystem + amount of executor memory per subsystem)

In continuous scoring, a separate job is spawned for each subsystem that is being analyzed by the IBM Z Operations Analytics machine learning system. A production environment with 4 subsystems requires 8 cores and 32 GB of memory.

Although it is a best practice to use 2 cores and 8 GB of memory for each subsystem (each continuous scoring job), you might need to use other values in certain situations. For example, an installation might collect data less frequently. Each scoring job then has fewer records to process, and the CPU requirement is less.

For continuous scoring of one subsystem (one continuous scoring job), the following example shows the values of the configuration properties in the `izoaml.config` file that are related to Apache Spark:

- `MLZ_SPARK_DRIVER_CORES`: 1
- `MLZ_SPARK_DRIVER_MEMORY`: 4 GB
- `MLZ_SPARK_EXECUTOR_CORES`: 1
- `MLZ_SPARK_EXECUTOR_MEMORY`: 4 GB

In this example, resources are not shared among multiple jobs.

Running concurrent scoring jobs requires a proportional increase in resources. For example, running 5 concurrent jobs requires 10 cores and 40 GB of memory. If enough resources are not available, some of the jobs cannot be started until resources are free, whereas training and batch scoring jobs for multiple systems run in sequential mode, and the resource requirement is constant.

Laboratory testing shows 2 cores per subsystem can process 5 minutes of data (1 collected record per minute) in 2 minutes on an IBM z14® processor. Different processor models result in different run times.

Training the analysis model and scoring data

These best practices apply to training the analysis model and to scoring data (in batch and continuous modes).

Training the analysis model and batch scoring

If you want to use multiple subsystems to train the analysis model, or if you want to run batch scoring against multiple subsystems, you gain system performance benefits if you run one Apache Spark job to train or score all subsystems rather than running a separate Apache Spark job for each subsystem. Running one job for all provides the following benefits:

- It eliminates the extra JVM startup processing that is required for a separate Spark job for each subsystem.
- It reduces the elapsed time for job completion.
- Because any training job (assuming fixed resource allocation) usually consumes the same amount of CPU in a specific time period, the reduced elapsed time results in reduced time of CPU consumption.

Continuous scoring

If you use a long scoring interval for continuous scoring, the run time might impact production.

Tip: Continuous scoring jobs run on the interval that is specified by the configuration property **`IZOA_SCORE_FREQ`** in the `izoaml.config` file. The **`IZOA_SCORE_FREQ`** value is specified in seconds and must be greater than or equal to 60. The default value for the continuous scoring interval is 300 seconds.

If you notice long run times for continuous scoring jobs, you might need to change the value of the **`IZOA_SCORE_FREQ`** configuration property as appropriate for your environment.

To determine the optimal scoring interval, record the run times of jobs that are scoring varying amounts of data. To obtain the run times, log onto the Apache Spark Master UI, and look for the *duration* of each of your applications. Based on these results, add a time buffer to the run times to derive the appropriate scoring interval.

The result of using a longer interval is a possible delay in scoring data and in detecting a system anomaly. The result of using a shorter interval is a possible increase in elapsed time for scoring the system data.

Test environment that was used for determining the best practices

Information is provided about the systems, subsystems, and applications that were used in the IBM Z Operations Analytics test environment for determining the tuning best practices.

- "Applications and subsystems" on page 81
- "Systems" on page 81

Applications and subsystems

- IBM Z Operations Analytics V4.1.0
- IBM Watson Machine Learning for z/OS V2.1.0.1
- IBM Z Common Data Provider V2.1
- IBM Open Data Analytics for z/OS V1.1
- IBM Db2 for z/OS V12

Systems

- 1 LPAR hosting z/OS-only installation of IBM Z Operations Analytics and IBM Watson Machine Learning for z/OS:
  - Z14 model M02, running z/OS V2.3
  - 8 total 3906.7 processors active:
    - 4 central processors
    - 4 z Systems® Integrated Information Processors (zIIPs)
  - 205 GB of online storage
- 1 LPAR hosting IBM Z Common Data Provider (for streaming test data to IBM Z Operations Analytics):
  - EC12 model HA1, running z/OS V2.3
  - 2 total 2827.7 processors active:
    - 2 central processors
    - No z Systems Integrated Information Processors (zIIPs)
  - 16 GB of online storage

Chapter 8. Reference for Problem Insights server and machine learning system

For the IBM Z Operations Analytics Problem Insights server, reference information is provided for associated commands, configuration files, and message libraries. For the IBM Z Operations Analytics machine learning system, reference information is provided for associated commands and configuration files.

Problem Insights server: command reference

This reference lists and describes the commands for operating the IBM Z Operations Analytics Problem Insights server.

Table 16. Commands for operating the Problem Insights server

Action	Command

Start the server.	**Linux system**
	`bin/analysis.sh start`
	Windows system
	`bin\analysis.bat start`
	z/OS system
	System Display and Search Facility (SDSF) command:
	`/S GLAPISRV`
	These commands load the message libraries.

Stop the server.	**Linux system**
	`bin/analysis.sh stop`
	Windows system
	`bin\analysis.bat stop`
	z/OS system
	System Display and Search Facility (SDSF) command:
	`/C GLAPISRV`

List the message libraries that were imported into the embedded database.	**z/OS UNIX System Services or Linux system**
	`bin/analysis.sh getlibrarylist`
	Windows system
	`bin\analysis.bat getlibrarylist`
	These commands give you the message library ID and locale, which you must use in the commands for exporting or deleting a message library.

Table 16. Commands for operating the Problem Insights server (continued)

Action	Command

Export a message library.	Use one of the following commands, depending on whether you want to export the message library to the console or to a file:
	Export to the console
	z/OS UNIX System Services or Linux system
	`bin/analysis.sh exportlibrary library_id locale`
	Windows system
	`bin\analysis.bat exportlibrary library_id locale`
	Export to a file
	z/OS UNIX System Services or Linux system
	`bin/analysis.sh exportlibrary library_id locale path_to_file`
	Windows system
	`bin\analysis.bat exportlibrary library_id locale path_to_file`
	Tip: To get the message library ID and locale, run the **`getlibrarylist`** command.

Delete a message library.	Use one of the following commands, depending on whether you want to have a prompt for user input:
	Delete *with* prompting for user input
	z/OS UNIX System Services or Linux system
	`bin/analysis.sh deletelibrary library_id locale`
	Windows system
	`bin\analysis.bat deletelibrary library_id locale`
	Delete *without* prompting for user input
	These commands include the argument **`-f`**, which forces the deletion without prompting for user input.
	z/OS UNIX System Services or Linux system
	`bin/analysis.sh deletelibrary -f library_id locale`
	Windows system
	`bin\analysis.bat deletelibrary -f library_id locale`
	Tip: To get the message library ID and locale, run the **`getlibrarylist`** command.

Table 16. Commands for operating the Problem Insights server (continued)

Action	Command

Import a message library.	**z/OS UNIX System Services or Linux system**
	`bin/analysis.sh importlibrary path_to_file`
	Windows system
	`bin\analysis.bat importlibrary path_to_file`

Delete data match results that are over 30 days old.	Use one of the following commands, depending on whether you want to have a prompt for user input:
	Delete *with* prompting for user input
	z/OS UNIX System Services or Linux system
	`bin/analysis.sh purgeResults`
	Windows system
	`bin\analysis.bat purgeResults`
	Delete *without* prompting for user input
	These commands include the argument **`-f`**, which forces the deletion without prompting for user input.
	z/OS UNIX System Services or Linux system
	`bin/analysis.sh purgeResults -f`
	Windows system
	`bin\analysis.bat purgeResults -f`

Generate a new
`bootstrap.properties` file
from the current configuration setup.
	z/OS UNIX System Services or Linux system
	`bin/analysis.sh createBootstrap`
	Windows system
	`bin\analysis.bat createBootstrap`

Encrypt text, such as a password.	**z/OS UNIX System Services or Linux system**
	`bin/analysis.sh encrypt text_to_encrypt`
	Windows system
	`bin\analysis.bat encrypt text_to_encrypt`

Problem Insights server: configuration file reference

This reference lists and describes the configuration files for the IBM Z Operations Analytics Problem Insights server.

`samples/amq.config`

This file contains information for configuring the Apache ActiveMQ message broker for use with the IBM Z Operations Analytics machine learning system and rules engine.

`samples/analysisroutine.config`
This file contains information for configuring the rules engine and analysis routines that are provided by IBM Z Operations Analytics for analyzing operational data.

`samples/cli.config`
This file contains information for defining the Problem Insights server, including, for example, the host, port, SSL configuration, and IBM IBM WebSphere Application Server for z/OS Liberty angel process.

`samples/db2zos.config`
This file is available only on IBM z/OS UNIX System Services. It contains information for connecting the Problem Insights server to the IBM Db2 for z/OS enterprise data warehouse for the IBM Z Operations Analytics machine learning system.

`samples/elk.config`
This file contains information for connecting the Problem Insights server to the Elastic Stack platform deployment.

`samples/event.config`
This file contains information for sending events from the Problem Insights server to other software (such as an event management system or service, or a chat application). For example, it includes the list of events to be sent.

`samples/microservices.config`
This file is available only on IBM z/OS UNIX System Services. It contains information for connecting subsystem-specific Problem Insights extensions to the IBM Db2 for z/OS enterprise data warehouse for the IBM Z Operations Analytics machine learning system.

`samples/splunk.config`
This file contains information for connecting the Problem Insights server to the Splunk platform deployment.

Problem Insights server: message library reference

In the IBM Z Operations Analytics Problem Insights GUI, the problem insights are provided by message libraries in the Problem Insights server. This reference includes, and describes, an example of a message library XML file and provides examples of command usage for operating message libraries.

Message libraries

The message libraries in the Problem Insights server are defined by the XML Schema Definition (XSD) file `messageLibrary.xsd`.

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 server 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.

Commands for operating message libraries

"Problem Insights server: command reference" on page 83 includes commands for operating (such as importing, exporting, or deleting) message libraries. "Examples that illustrate command usage for operating message libraries" on page 87 includes examples of how to use these commands.

Key parts of a message library XML file

"Example of a message library" on page 87 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 GUI:

- 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 GUI, 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 Z Common Data Provider 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 GUI 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="4.1.0"
 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>
```

Examples that illustrate command usage for operating message libraries

This reference includes examples that illustrate how to use various commands for operating the message libraries.

- “Get command usage information” on page 87
- “Initial import of the IBM-provided message libraries” on page 88
- “List all imported message libraries” on page 89
- “Import a message library” on page 89
- “Export a message library” on page 90
- “Delete a message library” on page 90
- “Update a message library” on page 91

Get command usage information

Usage information for the command-line interface (CLI) for the IBM Z Operations Analytics Problem Insights server is shown in the following situations:

- When you run the `analysis.sh` command without specifying any arguments

- When the command syntax is incorrect

Linux or z/OS UNIX System Services system

```
bin/analysis.sh
```

Windows system

```
bin\analysis.bat
```

Command results

```
GLAX059I Command usage: ./analysis.sh command arguments [optargs]
  start                    Start the server.
  stop                     Stop the server.
  getLibraryList           List the imported message library IDs with their locales and
descriptions.
  exportLibrary library_id locale [path_to_file]  Export a message library to the console.
                           To export a message library to a file,
                           specify the optional path_to_file argument.
  deleteLibrary library_id locale [-f]  Delete a message library.
                           To force the deletion without input, specify the optional -f
argument.
  importLibrary path_to_file  Import a message library.
  purgeResults [-f]        Delete data match results that are over 30 days old.
                           To force the deletion without input, specify the optional -f
argument.
  createBootstrap          Generate a new bootstrap.properties file from the current
configuration setup.
  encrypt text_to_encrypt  Encrypt text, such as a password.
```

Initial import of the IBM-provided message libraries

IBM-provided message libraries are shipped in the Problem Insights server directory `usr/lang/en/*.xml` and are automatically imported into the Problem Insights server embedded database during the initial start of the Problem Insights server. The import status for each IBM-provided message library is shown in message GLAX067I of the command results.

Tip: If a Problem Insights server is installed on an IBM z/OS UNIX System Services system, the best practice is to run it as a long-running task. IBM Z Operations Analytics includes the sample procedure GLAPISRV to help you configure this started task.

Linux or z/OS UNIX System Services system

```
bin/analysis.sh start
```

Windows system

```
bin\analysis.bat start
```

Command results

```
GLAX051I Starting the Problem Insights server. This may take a moment.
.........................
Enter the Problem Insights server administrative password:

GLAX084I The Problem Insights server has started.

GLAX066I Message Library: IZOAMQInsights.xml
GLAX067I Import status: GLAX038I message library: IZOAMQInsights.xml added successfully

GLAX066I Message Library: IZOACICSInsights.xml
GLAX067I Import status: GLAX038I message library: IZOACICSInsights.xml added successfully

GLAX066I Message Library: IZOADB2Insights.xml
GLAX067I Import status: GLAX038I message library: IZOADB2Insights.xml added successfully

GLAX066I Message Library: IZOANetworkInsights.xml
GLAX067I Import status: GLAX038I message library: IZOANetworkInsights.xml added successfully
```

```
GLAX066I Message Library: IZOAzOSInsights.xml
GLAX067I Import status: GLAX038I message library: IZOAzOSInsights.xml added successfully
```

List all imported message libraries

The message library ID, locale, and description are listed for each message library that was imported into the Problem Insights server embedded database.

Tip: The case-sensitive values for the message library ID and locale are required command input for exporting or deleting a message library.

Linux or z/OS UNIX System Services system

```
bin/analysis.sh getLibraryList
```

Windows system

```
bin\analysis.bat getLibraryList
```

Command results

```

Enter the Problem Insights server administrative password:

Library ID: NetworkforzOS
Locale: en_US
Description: Network message library description
#######################################################
Library ID: MQforzOS
Locale: en_US
Description: MQ message library description
#######################################################
Library ID: CICSforzOS
Locale: en_US
Description: CICS message library description
#######################################################
Library ID: DB2forzOS
Locale: en_US
Description: DB2 message library description
#######################################################
Library ID: zOS
Locale: en_US
Description: zOS message library description
#######################################################
```

Import a message library

In this example, a new user-provided message library is imported into the Problem Insights server embedded database. The user-provided message library is located in the home directory for the user, which is not part of the Problem Insights server directory structure. The file path to the message library is required command input.

Tip: Only a new message library can be imported. To verify that your message library was not previously imported into the Problem Insights server embedded database, issue the **`getLibraryList`** command (described in "Problem Insights server: command reference" on page 83), and verify that the associated message library ID and locale are not in the resulting list.

Linux or z/OS UNIX System Services system

```
bin/analysis.sh importlibrary ~/userDefinedInsights.xml
```

Windows system

```
bin\analysis.bat importlibrary %homepath%\userDefinedInsights.xml
```

Command results

```

Enter the Problem Insights server administrative password:
```

```
GLAX038I message library: userDefinedInsights.xml added successfully
```

Export a message library

This example shows how to export the user-provided message library that was imported in "Import a message library" on page 89.

To get the associated message library ID and locale, issue the **getLibraryList** command (described in "Problem Insights server: command reference" on page 83).

Sample result of issuing getLibraryList command

```
Library ID: userDefined
Locale: en_US
Description: user defined message library description
```

Tip: The case-sensitive values for the message library ID and locale are required command input for exporting a message library.

Optionally, you can provide a target file path to the exported message library. If you do not provide a value for the target file path, the exported message library is written to the console.

Linux or z/OS UNIX System Services system

```
bin/analysis.sh exportlibrary userDefined en_US ~/exportUserInsights.xml
```

Windows system

```
bin\analysis.bat exportlibrary userDefined en_US %homepath%\exportUserInsights.xml
```

Command results

```
Enter the Problem Insights server administrative password:

GLAX036I userDefined message library exported successfully.
```

Delete a message library

This example shows how to delete the user-provided message library that was imported in "Import a message library" on page 89.

To get the associated message library ID and locale, issue the **getLibraryList** command (described in "Problem Insights server: command reference" on page 83).

Sample result of issuing getLibraryList command

```
Library ID: userDefined
Locale: en_US
Description: user defined message library description
```

Tips:

- The case-sensitive values for the message library ID and locale are required command input for deleting a message library.
- Deleting a message library from the Problem Insights server embedded database does not delete the file from your server.

Linux or z/OS UNIX System Services system

```
bin/analysis.sh deleteLibrary userDefined en_US
```

Windows system

```
bin\analysis.bat deleteLibrary userDefined en_US
```

Command results

```
GLAX090I Are you sure you wish to delete this file? [y/n].
y

Enter the Problem Insights server administrative password:

GLAX036I userDefined message library deleted.
```

Update a message library

This example shows how to update an IBM-provided message library. You must update a message library as a whole. You cannot import, export, or delete a single message ID or section of a message library.

To update a message library that was previously imported to the Problem Insights server embedded database, you must complete the following steps:

1. Export the message library to a directory (such as the home directory for the user) that is not in the Problem Insights server directory structure. Also, make a backup copy of the exported message library so that you can return to a working state if you have any problems with your updated message library.
2. Update the exported message library.
3. Delete the message library from the Problem Insights server embedded database.
4. Import the updated copy of the message library into the Problem Insights server embedded database.

Linux or z/OS UNIX System Services system

```
bin/analysis.sh exportlibrary NetworkforzOS en_US ~/update_IZOANetworkforzOS.xml
*** make a backup copy of the exported message library ***
*** update the exported message library ***
bin/analysis.sh deletelibrary NetworkforzOS en_US
bin/analysis.sh importlibrary ~/update_IZOANetworkInsights.xml
```

Windows system

```
bin\analysis.bat exportlibrary NetworkforzOS en_US %homepath%\update_IZOANetworkforzOS.xml
*** make a backup copy of the exported message library ***
*** update the exported message library ***
bin\analysis.bat deletelibrary NetworkforzOS en_US
bin\analysis.bat importlibrary %homepath%\update_IZOANetworkInsights.xml
```

Machine learning system: command reference

This reference describes the commands, and their associated options and parameters, for operating the IBM Z Operations Analytics machine learning system.

Python scripts for running commands

You use the following Python scripts to run commands that are related to the machine learning system:

`$IZOA_HOME/bin/izoaml.py`
The main script for running the machine learning system commands.

`$IZOA_HOME/bin/submit.py`
The script for submitting JCL jobs.

In this command reference, the command examples indicate how to use these scripts.

feature and **run** parameters for **izoaml.py** script

The parameters **feature** and **run** are used in most of the izoaml.py commands, as shown in example commands in this command reference.

The value of the **feature** parameter corresponds to one or more of the Insight Packs that you are using. For example, it can have either or both of the values that are shown in Table 17 on page 92.

For information about the Insight Packs, see the following topics:

- "Insight Packs" on page 20
- "Installing or updating Insight Packs" on page 33

*Table 17. Correlation of **feature** parameter value to the Insight Packs that you use*

Value for feature parameter	Associated Insight Pack
cicshealth	CICS Transaction Server for z/OS
db2health	IBM Db2 for z/OS

For example, each of the following three **database=load** commands uses a different value for the **feature** parameter:

Load CICS Transaction Server for z/OS SMF data

```
$IZOA_HOME/bin/izoaml.py --feature=cicshealth --database=load --run
```

Load Db2 for z/OS SMF data

```
$IZOA_HOME/bin/izoaml.py --feature=db2health --database=load --run
```

Load both CICS Transaction Server for z/OS and Db2 for z/OS SMF data

```
$IZOA_HOME/bin/izoaml.py --feature=cicshealth,db2health --database=load --run
```

database command

Use the **database** command to manage the IBM Db2 for z/OS enterprise data warehouse database for the IBM Z Operations Analytics machine learning system.

Command options

- "database=bufferpool_all" on page 93
- "database=bufferpool_history" on page 93
- "database=bufferpool_ml" on page 93
- "database=cleandata" on page 93
- "database=collect" on page 94
- "database=concurrent" on page 94
- "database=copy" on page 94
- "database=load" on page 94
- "database=reorg_all" on page 95
- "database=reorg_history" on page 95
- "database=reorg_ml" on page 95
- "database=sql" on page 95
- "database=stat_all" on page 95

- “database=stat_history” on page 95
- “database=stat_ml” on page 96

database=bufferpool_all

Command option description
Alter the buffer pool for all database tables.

Task where this command option is used
“Maintaining the enterprise data warehouse” on page 75

Valid parameters for this option
Not applicable.

database=bufferpool_history

Command option description
Alter the buffer pool for the historical database tables.

Task where this command option is used
“Maintaining the enterprise data warehouse” on page 75

Valid parameters for this option
Not applicable.

database=bufferpool_ml

Command option description
Alter the buffer pool for the machine learning database tables.

Task where this command option is used
“Maintaining the enterprise data warehouse” on page 75

Valid parameters for this option
Not applicable.

database=cleandata

Command option description
Remove unneeded analysis result data from the enterprise data warehouse.

Task where this command option is used
“Maintaining the enterprise data warehouse” on page 75

Valid parameters for this option

IZOA_END_ENTRYTIME
Specifies the end time of the time range for data to be removed.

Default value
`2025-12-31 23:59:59`

IZOA_START_ENTRYTIME
Specifies the start time of the time range for data to be removed.

Default value
`2018-01-01 00:00:00`

IZOA_SUBSYS_ID
Specifies the ID of the subsystem that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **`define=delete`** command, this value is optional.

- On the **define=delete**, **define=find**, and **define=update** commands, this value can contain the wildcard character % (to represent the IDs of all subsystems that are monitored and analyzed).
- On the **define=create** command, this value cannot contain the wildcard character.

IZOA_SYS_ID
Specifies the ID of the MVS™ system that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **define=delete** command, this value is optional.
- On the **define=delete**, **define=find**, and **define=update** commands, this value can contain the wildcard character % (to represent the IDs of all MVS systems that are monitored and analyzed).
- On the **define=create** command, this value cannot contain the wildcard character.

database=collect

Command option description
Collect SMF data.

Task where this command option is used
"Loading historical SMF data" on page 69

Valid parameters for this option

Important: The parameter **IZOA_SMFGDG** is required.

IZOA_SMFGDG
Specifies the high-level qualifier for the SMF generation data group (GDG) data set that is used to run IBM Z Common Data Provider collect jobs.

database=concurrent

Command option description
Prepare the enterprise data warehouse for concurrent read and write operations.

Task where this command option is used
"Loading historical SMF data" on page 69

Valid parameters for this option
Not applicable.

database=copy

Command option description
Create a backup copy of the database tables in the enterprise data warehouse.

Task where this command option is used
"Loading historical SMF data" on page 69

Valid parameters for this option
Not applicable.

database=load

Command option description
Load SMF data.

Task where this command option is used
"Loading historical SMF data" on page 69

Valid parameters for this option
Not applicable.

`database=reorg_all`

Command option description
Run the Db2 REORG TABLESPACE utility against all database tables.

Task where this command option is used
"Maintaining the enterprise data warehouse" on page 75

Valid parameters for this option
Not applicable.

`database=reorg_history`

Command option description
Run the Db2 REORG TABLESPACE utility against the historical database tables.

Task where this command option is used
"Maintaining the enterprise data warehouse" on page 75

Valid parameters for this option
Not applicable.

`database=reorg_ml`

Command option description
Run the Db2 REORG TABLESPACE utility against the machine learning database tables.

Task where this command option is used
"Maintaining the enterprise data warehouse" on page 75

Valid parameters for this option
Not applicable.

`database=sql`

Command option description
Check the buffer pool in use for tables that are associated with a specific subsystem type.

Task where this command option is used
"Maintaining the enterprise data warehouse" on page 75

Valid parameters for this option

IZOA_FN
Specifies the path for SQL scripts that are to be processed by the **`database=sql`** command.

`database=stat_all`

Command option description
Run the Db2 RUNSTATS utility against all database tables.

Task where this command option is used
"Maintaining the enterprise data warehouse" on page 75

Valid parameters for this option
Not applicable.

`database=stat_history`

Command option description
Run the Db2 RUNSTATS utility against the historical database tables.

Task where this command option is used
"Maintaining the enterprise data warehouse" on page 75

Valid parameters for this option
Not applicable.

database=stat_ml

Command option description
Run the Db2 RUNSTATS utility against the machine learning database tables.

Task where this command option is used
"Maintaining the enterprise data warehouse" on page 75

Valid parameters for this option
Not applicable.

define command

Use the **define** command to define the subsystems that you want to analyze to the IBM Z Operations Analytics machine learning system.

Command options

- "define=create" on page 96
- "define=delete" on page 97
- "define=find" on page 98
- "define=update" on page 99

define=create

Command option description
Define a subsystem.

Task where this command option is used
"Training the analysis model" on page 70

Valid parameters for this option

Important: The following parameters are required:

- **IZOA_SYS_ID**
- **IZOA_SUBSYS_ID**
- **IZOA_TIMEZONE**

IZOA_SUBSYS_ID
Specifies the ID of the subsystem that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **define=delete** command, this value is optional.
- On the **define=delete**, **define=find**, and **define=update** commands, this value can contain the wildcard character % (to represent the IDs of all subsystems that are monitored and analyzed).
- On the **define=create** command, this value cannot contain the wildcard character.

IZOA_SYS_ID
Specifies the ID of the MVS system that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **define=delete** command, this value is optional.

- On the **define=delete**, **define=find**, and **define=update** commands, this value can contain the wildcard character % (to represent the IDs of all MVS systems that are monitored and analyzed).
- On the **define=create** command, this value cannot contain the wildcard character.

IZOA_TIMEZONE

Specifies the time zone of the system that generates the data to be processed. The value of this parameter must be the name of a time zone as assigned by the Internet Assigned Numbers Authority (IANA). The value also must be enclosed in quotation marks, as shown in the following examples:

- `"America/New_York"`
- `"Europe/London"`

Tip:

- On the **define=create**, **define=find**, and **define=update** commands, the **IZOA_TIMEZONE** parameter specifies the time zone of the MVS system that is specified by the **IZOA_SYS_ID** parameter.
- On the **define=find** and **define=update** commands, if the **IZOA_SYS_ID** is not specified, the time zone applies to all discovered systems.
- On the **define=find** and **define=update** commands, if the **IZOA_SYS_ID** contains one or more wildcard characters, the time zone applies to all discovered systems that match the wildcard.
- **IZOA_SYS_ID** is optional for the **define=find** and **define=update** commands. If it is not specified, **IZOA_TIMEZONE** applies to all discovered systems.

IZOA_USER_LABEL

Specifies a label that identifies the analysis model in the machine learning system.

Tip: On the **define=delete** command, this value is required.

define=delete

Command option description

Delete one or more subsystem definitions.

Task where this command option is used

"Training the analysis model" on page 70

Valid parameters for this option

Tip:

- On the **define=delete** and **define=update** commands, if the **IZOA_SUBSYS_ID** parameter is not specified, these operations apply to all defined subsystems for the specified system. If neither the **IZOA_SYS_ID** parameter nor the **IZOA_SUBSYS_ID** parameter is specified, these operations apply to all defined systems and subsystems (and the value of the **IZOA_TIMEZONE** parameter is applied to all of them).
- On the **define=find** command, if the **IZOA_SUBSYS_ID** parameter is not specified, this operation discovers all subsystems in the loaded data for the specified system. If neither the **IZOA_SYS_ID** parameter nor the **IZOA_SUBSYS_ID** parameter is specified, this operation discovers all systems and subsystems in the loaded data (and the value of the **IZOA_TIMEZONE** parameter is applied to all of them).

IZOA_SUBSYS_ID

Specifies the ID of the subsystem that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **define=delete** command, this value is optional.

- On the **define=delete**, **define=find**, and **define=update** commands, this value can contain the wildcard character % (to represent the IDs of all subsystems that are monitored and analyzed).
- On the **define=create** command, this value cannot contain the wildcard character.

IZOA_SYS_ID

Specifies the ID of the MVS system that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **define=delete** command, this value is optional.
- On the **define=delete**, **define=find**, and **define=update** commands, this value can contain the wildcard character % (to represent the IDs of all MVS systems that are monitored and analyzed).
- On the **define=create** command, this value cannot contain the wildcard character.

IZOA_USER_LABEL

Specifies a label that identifies the analysis model in the machine learning system.

Tip: On the **define=delete** command, this value is required.

define=find

Command option description

Direct the machine learning system to search the loaded data for subsystems and to create associated subsystem definitions.

Task where this command option is used

"Training the analysis model" on page 70

Valid parameters for this option

IZOA_END_DATE

Specifies the end time in the time stamp for monitored system discovery.

IZOA_START_DATE

Specifies the start time in the time stamp for monitored system discovery.

IZOA_SUBSYS_ID

Specifies the ID of the subsystem that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **define=delete** command, this value is optional.
- On the **define=delete**, **define=find**, and **define=update** commands, this value can contain the wildcard character % (to represent the IDs of all subsystems that are monitored and analyzed).
- On the **define=create** command, this value cannot contain the wildcard character.

IZOA_SYS_ID

Specifies the ID of the MVS system that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **define=delete** command, this value is optional.
- On the **define=delete**, **define=find**, and **define=update** commands, this value can contain the wildcard character % (to represent the IDs of all MVS systems that are monitored and analyzed).
- On the **define=create** command, this value cannot contain the wildcard character.

IZOA_SYSPLEX_NAME

Specifies the name of the sysplex that is monitored and analyzed by IBM Z Operations Analytics.

IZOA_SYSTEM

Specifies the system name for the system that is monitored and analyzed IBM Z Operations Analytics.

IZOA_TIMEZONE

Specifies the time zone of the system that generates the data to be processed. The value of this parameter must be the name of a time zone as assigned by the Internet Assigned Numbers Authority (IANA). The value also must be enclosed in quotation marks, as shown in the following examples:

- `"America/New_York"`
- `"Europe/London"`

Tip:

- On the **`define=create`**, **`define=find`**, and **`define=update`** commands, the **`IZOA_TIMEZONE`** parameter specifies the time zone of the MVS system that is specified by the **`IZOA_SYS_ID`** parameter.
- On the **`define=find`** and **`define=update`** commands, if the **`IZOA_SYS_ID`** is not specified, the time zone applies to all discovered systems.
- On the **`define=find`** and **`define=update`** commands, if the **`IZOA_SYS_ID`** contains one or more wildcard characters, the time zone applies to all discovered systems that match the wildcard.
- **`IZOA_SYS_ID`** is optional for the **`define=find`** and **`define=update`** commands. If it is not specified, **`IZOA_TIMEZONE`** applies to all discovered systems.

IZOA_USER_LABEL

Specifies a label that identifies the analysis model in the machine learning system.

Tip: On the **`define=delete`** command, this value is required.

define=update

Command option description

Update one or more subsystem definitions.

Task where this command option is used

"Training the analysis model" on page 70

Valid parameters for this option

IZOA_SUBSYS_ID

Specifies the ID of the subsystem that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **`define=delete`** command, this value is optional.
- On the **`define=delete`**, **`define=find`**, and **`define=update`** commands, this value can contain the wildcard character % (to represent the IDs of all subsystems that are monitored and analyzed).
- On the **`define=create`** command, this value cannot contain the wildcard character.

IZOA_SUBSYS_STATUS

Specifies the status of the subsystem that is monitored and analyzed by IBM Z Operations Analytics.

IZOA_SYS_ID

Specifies the ID of the MVS system that is monitored and analyzed by IBM Z Operations Analytics.

Tip:

- On the **`define=delete`** command, this value is optional.

- On the **define=delete**, **define=find**, and **define=update** commands, this value can contain the wildcard character % (to represent the IDs of all MVS systems that are monitored and analyzed).
- On the **define=create** command, this value cannot contain the wildcard character.

IZOA_SYSPLEX_NAME
Specifies the name of the sysplex that is monitored and analyzed by IBM Z Operations Analytics.

IZOA_SYSTEM
Specifies the system name for the system that is monitored and analyzed IBM Z Operations Analytics.

IZOA_TIMEZONE
Specifies the time zone of the system that generates the data to be processed. The value of this parameter must be the name of a time zone as assigned by the Internet Assigned Numbers Authority (IANA). The value also must be enclosed in quotation marks, as shown in the following examples:

- `"America/New_York"`
- `"Europe/London"`

Tip:

- On the **define=create**, **define=find**, and **define=update** commands, the **IZOA_TIMEZONE** parameter specifies the time zone of the MVS system that is specified by the **IZOA_SYS_ID** parameter.
- On the **define=find** and **define=update** commands, if the **IZOA_SYS_ID** is not specified, the time zone applies to all discovered systems.
- On the **define=find** and **define=update** commands, if the **IZOA_SYS_ID** contains one or more wildcard characters, the time zone applies to all discovered systems that match the wildcard.
- **IZOA_SYS_ID** is optional for the **define=find** and **define=update** commands. If it is not specified, **IZOA_TIMEZONE** applies to all discovered systems.

IZOA_USER_LABEL
Specifies a label that identifies the analysis model in the machine learning system.

Tip: On the **define=delete** command, this value is required.

ml command

Use the **ml** command for interacting with the IBM Z Operations Analytics machine learning system. For example, you use the **ml** command 1) for training the analysis model that the machine learning system uses for detecting anomalies in your Z environment and 2) for scoring the incoming data.

Command options

- “ml=batchscore” on page 100
- “ml=scag_batchscore” on page 101
- “ml=scag_stop” on page 101
- “ml=train” on page 101

ml=batchscore

Command option description
Score the incoming SMF data once.

Task where this command option is used
“Analyzing and scoring new SMF data” on page 72

Valid parameters for this option

Important: The parameter **`IZOA_DATA_TIME_RANGE`** is required.

IZOA_DATA_TIME_RANGE
Specifies the time range for collection of the data that is to be used for training the analysis model and scoring the data. The value of this parameter must be in the following format and must be enclosed in quotation marks, as indicated:

```
"YYYYMMDD-HHMMSS;YYYYMMDD-HHMMSS"
```

IZOA_ML_PROP_FILE
Specifies the name of the property file that contains the machine learning parameters.

IZOA_USER_LABEL
Specifies a label that identifies the analysis model in the machine learning system.

Tip: On the **`define=delete`** command, this value is required.

ml=scag_batchscore

Command option description
Score the incoming SMF data continuously.

Task where this command option is used
"Analyzing and scoring new SMF data" on page 72

Valid parameters for this option

IZOA_ML_PROP_FILE
Specifies the name of the property file that contains the machine learning parameters.

ml=scag_stop

Command option description
Stop the scoring agent, which is used in continuous scoring.

Task where this command option is used
"Analyzing and scoring new SMF data" on page 72

Valid parameters for this option
Not applicable

ml=train

Command option description
Train the analysis model.

Task where this command option is used
"Training the analysis model" on page 70

Valid parameters for this option

Important: The parameter **`IZOA_DATA_TIME_RANGE`** is required.

IZOA_DATA_TIME_RANGE
Specifies the time range for collection of the data that is to be used for training the analysis model and scoring the data. The value of this parameter must be in the following format and must be enclosed in quotation marks, as indicated:

```
"YYYYMMDD-HHMMSS;YYYYMMDD-HHMMSS"
```

IZOA_USER_LABEL
Specifies a label that identifies the analysis model in the machine learning system.

Tip: On the **`define=delete`** command, this value is required.

status command

Use the **status** command to view the status of the training and scoring jobs that are started by using the **ml** command.

Command options and explanation of command output

- **Options:**
  - "status=all" on page 102
  - "status=submission ID of the job to be viewed" on page 103
- **Explanation of command output:**
  - "Status command output" on page 103

status=all

Command option description

View the status of all jobs.

Task where this command option is used

- "Training the analysis model" on page 70
- "Analyzing and scoring new SMF data" on page 72

Valid parameters for this option

JOB_MAX

Specifies the maximum number of jobs to show in the status view for training and scoring jobs of the IBM Z Operations Analytics machine learning system.

When you specify the maximum number of jobs to show, the most recently submitted job is shown first.

JOB_STATUS

Specifies that you want to filter jobs by status in the status view for training and scoring jobs of the IBM Z Operations Analytics machine learning system.

The following values are valid:

SUBMITTED

The job was submitted but is not yet processed and in queue to run.

QUEUED

The job is in the queue to run but is not yet started.

RUNNING

The job is running.

COMPLETED

The job completed successfully. The resulting enterprise data warehouse tables should be updated.

FAILED

The job completed unsuccessfully.

CANCELED

The job was canceled by a user before it completed.

JOB_TYPE

Specifies that you want to filter jobs by type in the status view for training and scoring jobs of the IBM Z Operations Analytics machine learning system.

Either or both of the following values are valid:

score
To view only the submitted scoring jobs.

train
To view only the submitted training jobs.

status=*submission ID of the job to be viewed*

Command option description
View the status of a specific job.

Task where this command option is used

- "Training the analysis model" on page 70
- "Analyzing and scoring new SMF data" on page 72

Valid parameters for this option
Not applicable

Status command output

The output for the **status** command is a table format with the following columns:

Submitted
The date and time at which the job was submitted.

Submission_ID
The submission ID of the remote job request. This ID is used to track the status of the job.

Type
The type of job that was submitted. You can specify the **JOB_TYPE** parameter on the **status=all** command to filter jobs by type in the output.

Status
The current status of the submitted job. You can specify the **JOB_STATUS** parameter on the **status=all** command to filter jobs by status in the output.

Tip: For details about the jobs, including STDOUT and STDERR logs, see the Apache Spark Web UI.

Machine learning system: configuration file reference

This reference lists and describes the configuration files for the IBM Z Operations Analytics machine learning system.

Important: Unlike the configuration file `samples/izoaml.config`, the working copies of the configuration files `samples/IZOA_CICSHEALTH_ML.properties` and `samples/IZOA_DB2HEALTH_ML.properties` can be in the location of your choice. However, rather than editing the original files in the `samples` directory, create a copy of the files `samples/IZOA_CICSHEALTH_ML.properties` and `samples/IZOA_DB2HEALTH_ML.properties` in a different location, and customize each copy for your environment.

samples/izoaml.config
This file contains information for configuring the interaction between the machine learning system and the following software:

- IBM Z Common Data Provider
- IBM Db2 for z/OS
- IBM Watson Machine Learning for z/OS
- Other related software

It also contains the configuration properties that define the basic behavior of the machine learning system.

samples/IZOA_CICSHEALTH_ML.properties

This file contains properties for customizing the parameters for the CICS Health analysis model for machine learning.

samples/IZOA_DB2HEALTH_ML.properties

This file contains properties for customizing the parameters for the Db2 Health analysis model for machine learning.

Chapter 9. Deploying the Z Log and Data Analytics dashboards and searches on your analytics platform

IBM Z Log and Data Analytics dashboards and predefined searches give you the capability to search, visualize, and analyze large amounts of structured and unstructured operational data across Z systems IT environments. To use these dashboards and searches, deploy the IBM Z Log and Data Analytics application on the analytics platform of your choice (Z Log Analysis, Elastic Stack, or Splunk). The user interface and available functions can vary depending on the platform.

Z Log and Data Analytics on the Z Log Analysis platform

On the Z Log Analysis platform, IBM Z Operations Analytics provides dashboards and searches for Z operational insights.

Z Log and Data 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" on page 105
- "Flow of source data on the Elastic Stack platform" on page 105
- "Flow of source data on the Splunk platform" on page 106
- "Summary of system requirements" on page 107

Component overview

For each platform, IBM Z Operations Analytics includes the following components:

1. IBM Z Common Data Provider

   See Chapter 3, "Z Common Data Provider: Planning for installation and configuration," on page 7.
2. IBM Z Operations Analytics application

   See "Deploying the Z Operations Analytics application on the Elastic Stack platform" on page 108 or "Deploying the Z Operations Analytics application on the Splunk platform" on page 112.
3. Problem Insights server

   See Chapter 4, "Deploying the Problem Insights server," on page 15.
4. Machine learning system

   See Chapter 5, "Deploying the machine learning system," on page 47.

Flow of source data on the Elastic Stack platform

Figure 3 on page 106 illustrates the flow of source data among the primary components of IBM Z Operations Analytics on the Elastic Stack platform.

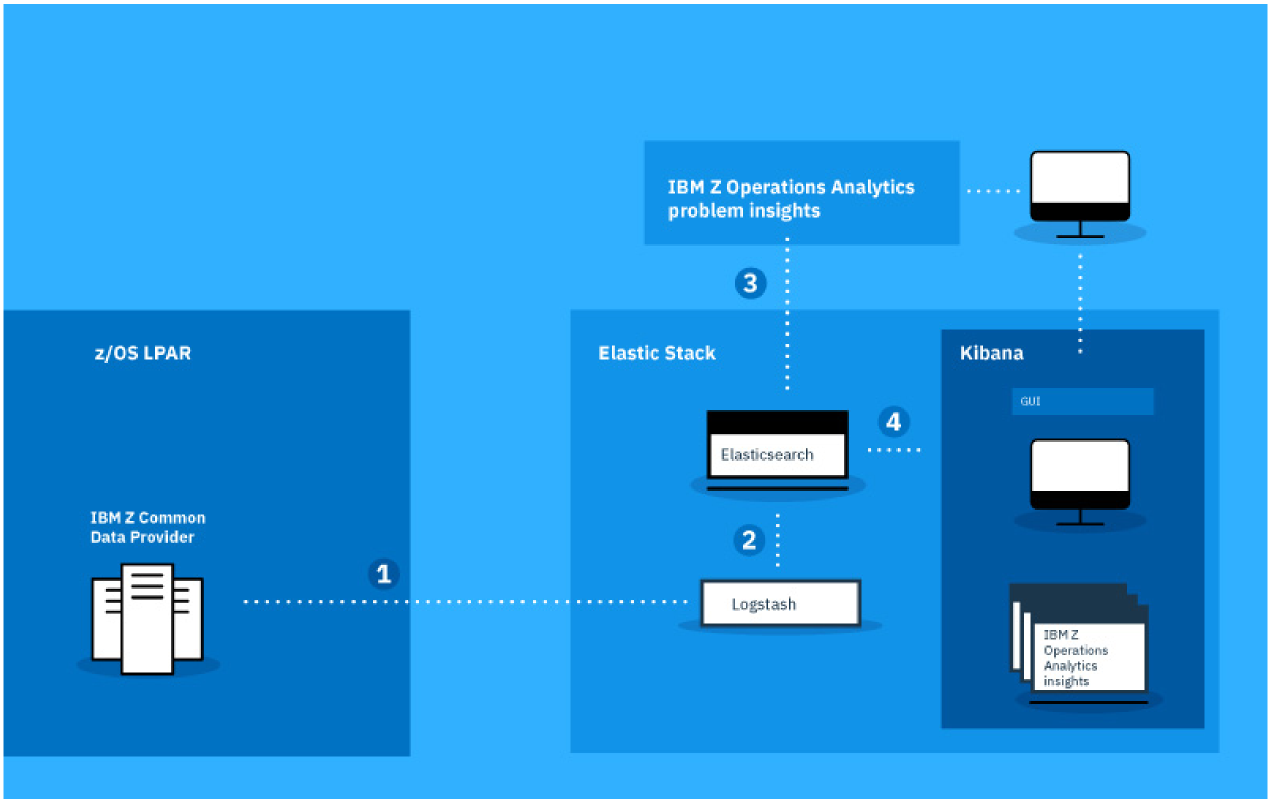
IBM Z Operations Analytics problem insights
3
z/OS LPAR
Elastic Stack
Kibana
GUI
4
Elasticsearch
2
IBM Z Common Data Provider
1
Logstash
IBM Z Operations Analytics insights

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

The following steps describe the data flow among components. The step numbers correspond to the numbers that are used in the illustration.

1. In each z/OS logical partition (LPAR), the IBM Z Common Data Provider 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 on the Splunk platform

Figure 4 on page 107 illustrates the flow of source data among the primary components of IBM Z Operations Analytics on the Splunk platform.

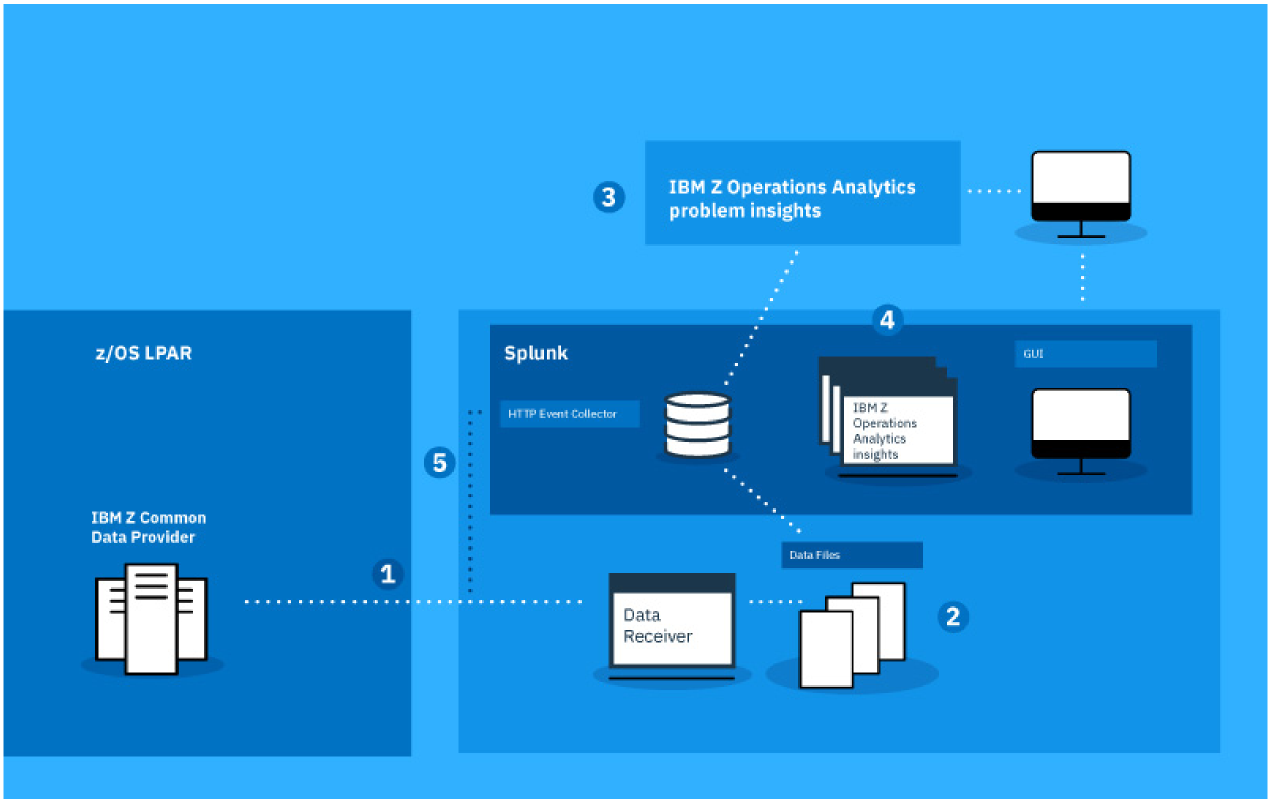
3
IBM Z Operations Analytics problem insights
z/OS LPAR
Splunk
4
GUI
HTTP Event Collector
IBM Z Operations Analytics insights
5
IBM Z Common Data Provider
1
Data Files
Data Receiver
2

Figure 4. 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 Z Common Data Provider retrieves the data from the respective source and sends it to the Splunk Enterprise server.
2. The source data is received by the IBM Z Common Data Provider Data Receiver and is written to local data files. Splunk reads and processes the local data files based on rules that are provided by IBM Z Operations Analytics and the IBM Z Common Data Provider Buffered Splunk Ingestion App.
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 the Splunk GUI. 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. As an alternative to the IBM Z Common Data Provider Data Receiver, customers can ingest data directly into Splunk by using the Splunk HTTP Event Collector (HEC).

Summary of system requirements

IBM Z Common Data Provider

Planning to use IBM Z Common Data Provider in the IBM Z Common Data Provider documentation

z/OS systems and subsystems: version requirements

“Software version requirements for the z/OS systems and subsystems from which operational data is gathered” on page 108

IBM Z Operations Analytics application on the Elastic Stack platform

“System requirements for deploying the IBM Z Operations Analytics application on the Elastic Stack platform” on page 112

IBM Z Operations Analytics application on the Splunk platform

“System requirements for deploying the IBM Z Operations Analytics application on the Splunk platform” on page 113

IBM Z Operations Analytics Problem Insights server and machine learning system

- “System requirements for the Problem Insights server” on page 16
- “System requirements for the machine learning system” on page 50

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, 2.3, or 2.4
- IBM CICS Transaction Server for z/OS 5.2, 5.3, 5.4, 5.5, or 5.6
- IBM Db2 for z/OS 11.1 or 12.1
- IBM IMS for z/OS 14.1, 15.1, or 15.2
- IBM MQ for z/OS 9.0 or 9.1
- IBM Z NetView 6.2 or 6.3
- IBM WebSphere Application Server for z/OS 8.5.5 or later, or 9.0
- Access Monitor component of IBM Security zSecure Admin 2.2.1 with APAR OA52273, 2.3, or 2.4

Deploying the Z Operations Analytics application on the Elastic Stack platform

To deploy the IBM Z Operations Analytics application on the Elastic Stack platform, complete this task.

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

If you are using the Elasticsearch ingestion kit for IBM Z Common Data Provider:

Do not use the Logstash pipeline configuration files that are provided by the following kits together:

- IBM Z Operations Analytics Elastic Stack ingestion kit
- IBM Z Common Data Provider Elasticsearch ingestion kit

Having these files in the same pipeline configuration directory is not supported and can lead to unpredictable results.

If you need to use both ingestion kits within the same Elastic Stack environment, implement one of the following options:

- Configure a unique Logstash instance for each ingestion kit.
- If you prefer to use a single Logstash instance, configure a unique ingestion pipeline for each ingestion kit.

  For information about how to configure multiple pipelines in a single Logstash instance, see the product documentation for your version of Logstash. The latest version of this documentation is available at https://www.elastic.co/guide/en/logstash/current/multiple-pipelines.html.

Regardless of which option you choose, ensure that the contents of each ingestion kit are stored in two unique directories.

Important: Logstash processes all files that are included in the configuration file directory *config_file_dir*. Processing is done in lexicographical order. Either copying more files into the configuration file directory, or making backups of existing files in the configuration file directory, can change Logstash startup behavior and, in some cases, might cause Logstash to fail.

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 `IZOA-IngestionKit-4.1.0.zip`.

   Table 18 on page 109 indicates the prefixes that are used in the file names for the Logstash configuration files in the IBM Z Operations Analytics Elastic Stack ingestion kit. The file name prefix is an indication of the configuration file content.

   Table 18. 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
E_	Preparation stage
H_	Field name annotation stage
Q_	Output stage

   The following descriptions further explain these Logstash configuration files:

   `B_cdpz.conf` file
   This file contains the input stage that specifies the TCP/IP port on which Logstash listens for data from the IBM Z Common Data Provider Data Streamer. Update the port number in this file as appropriate for your environment.

   `E_transform.conf` file
   This file prepares the data that is received from the IBM Z Common Data Provider for further processing.

Files with `H_` prefix in file name

Each of these files contains a unique field name annotation stage that maps to a unique data stream that IBM Z Common Data Provider can send to Logstash.

`Q_elasticsearch.conf` file

This file contains an output stage that sends all records to a single Elasticsearch server. Update the **`hosts`** parameter in this file as appropriate for your environment. The value of the **`index`** parameter is assigned during ingestion so that the data for each source type is sent to a different index.

2. 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-*`.

3. 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`

`.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`

4. Verify that Elasticsearch, Logstash, and Kibana are running.

   To verify the status of Elasticsearch, Logstash, and Kibana, run the following commands on the respective servers. The commands list each process, if that process is running.

   Elasticsearch

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

Logstash

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

Kibana

```
# ps -ef | grep node
```

5. Log in to the Kibana server.
6. Extract the IBM Z Operations Analytics package `IZOA-Elastic-V4.1.0.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 server is bound.

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

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

   Tip: When the dashboard files are imported into Kibana, some exceptions might occur, but they do not represent a functional problem and can be ignored.

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 Connect Enterprise Edition API Dashboard
- z/OS Connect Enterprise Edition Request URI Dashboard
- z/OS Connect Enterprise Edition Service Dashboard
- z/OS Job Dashboard
- z/OS Security Server RACF Dashboard
- zSecure Access Monitor Dashboard

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:

- One of the following Elastic Stack releases:
  - Any release of Elastic Stack Version 6, except Elastic Stack 6.0, that is still supported by Elasticsearch B.V.
  - Any release of Elastic Stack Version 7 that is still supported by Elasticsearch B.V.
- 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 IBM 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" on page 113, 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 124.

Clustered Splunk environment
See "Deploying Z Operations Analytics in a clustered Splunk environment" on page 125.

As part of your deployment, you can also install the Splunk IT Service Intelligence (ITSI) module for IBM Z Operations Analytics. For more information, see "Splunk ITSI module for IBM Z Operations Analytics" on page 129.

Important: To send data to Splunk, you can use either the IBM Z Common Data Provider Data Receiver or the Splunk HTTP Event Collector (HEC) as the subscriber, as indicated in "Subscribers for each type of source data" on page 10.

The Splunk HEC is an HTTP API endpoint that enables you to send data directly to Splunk over HTTP or HTTPS. If the Splunk HEC feature is enabled in Splunk, IBM Z Common Data Provider can send data directly to Splunk through the HEC rather than sending data through the Data Receiver.

- If you want to use the IBM Z Common Data Provider Data Receiver, you must also install the IBM Z Common Data Provider Buffered Splunk Ingestion App, as described in both "Deploying Z Operations Analytics on a single Splunk Enterprise system" on page 124 and "Installing the Splunk ITSI module for IBM Z Operations Analytics" on page 138.
- If you want to use the Splunk HEC, you must complete the steps in "Sending data directly to Splunk by using Splunk HEC as the subscriber" on page 128 **after** you deploy the IBM Z Operations Analytics application either on a single Splunk Enterprise system or in a clustered Splunk environment.

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:

- Any release of Splunk Enterprise Version 7 or 8 that is still supported by Splunk Inc.
- For the Splunk ITSI module for IBM Z Operations Analytics, the required version of Splunk ITSI is dependent on the version of Splunk Enterprise that you use, as shown in the following table:

Splunk Enterprise version	Required version of Splunk ITSI
7.2	4.0.1 or later
7.3	4.2.0 or later
8.0	4.4.0 or later
8.1	4.7.0 or 4.7.1 **Known problem:** Although the required Splunk ITSI version for Splunk Enterprise 8.1 is either 4.7.0 or 4.7.1, the Splunk ITSI module for IBM Z Operations Analytics does not work with Splunk ITSI 4.7.

For more information about compatibility among versions of Splunk Enterprise and other Splunk products, see the Splunk Products Version Compatibility Matrix.

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

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.

For more information about these macros, see "Splunk macros in IBM Z Operations Analytics application on the Splunk platform" on page 115.

The default values for indexes are different depending on how the data is streamed to Splunk, as described in the following sections:

- "Indexes when using the Data Receiver as the subscriber" on page 114
- "Indexes when using the Splunk HEC as the subscriber" on page 115

Indexes when using the Data Receiver as the subscriber

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. However, the following information describes two exceptions to this naming convention:

Data source type for z/OS SYSLOG Console data

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 Z Common Data Provider as part of the Buffered Splunk Ingestion App. The index name for z/OS SYSLOG Console data is `zosdex`.

All log data source types for Version 4.1.02 and later versions of the IBM Z Operations Analytics application on the Splunk platform

For Version 4.1.02 and later versions of the IBM Z Operations Analytics application on the Splunk platform, all log data source types are an exception to this naming convention. Input and index configurations for log data source types, which were previously defined in the IBM Z Operations Analytics application on the Splunk platform, are now defined in the Buffered Splunk Ingestion App in the IBM Z Common Data Provider Version 2.1 Continuous Delivery PTF (UJ03273). The following indexes for the respective source types were removed from the IBM Z Operations Analytics V4.1.02 application on the Splunk platform and are included in the macros for migration purposes only:

- `zOS-CICS-EYULOG`
- `zOS-CICS-EYULOGDMY`
- `zOS-CICS-EYULOGYMD`
- `zOS-CICS-MSGUSR`
- `zOS-CICS-MSGUSRDMY`
- `zOS-CICS-MSGUSRYMD`
- `zOS-NetView`
- `zOS-syslogd`
- `zOS-WAS-SYSOUT`
- `zOS-WAS-SYSPRINT`
- `zOS-zSecure`

The index name for all of the log 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.

Also, to better align with Splunk best practices, the names for IBM-provided indexes in Version 4.1.04, and later versions of the IBM Z Operations Analytics application on the Splunk platform, now use lowercase characters (rather than mixed case characters).

Indexes when using the Splunk HEC as the subscriber

You can send data directly to Splunk by using the Splunk HTTP Event Collector (HEC). Use of the Splunk HEC enables the quick setup of your environment without the need for the IBM Z Common Data Provider Data Receiver and Buffered Splunk Ingestion App.

A new index `zosdex_kv` is provided with the IBM Z Operations Analytics application on the Splunk platform. This `zosdex_kv` index is used to show data that is streamed to a Splunk HEC subscriber for all IBM-provided dashboards and predefined searches. When you create a HEC token for communicating with Splunk, you must set the value for the **Default index** field to `zosdex_kv` so that the dashboards and predefined searches continue to show the expected results.

Splunk macros in IBM Z Operations Analytics application on the Splunk platform

Previously, the Splunk macros that were provided with the IBM Z Operations Analytics application defined the indexes that are used by the predefined dashboards and searches. However, with the addition of Splunk HTTP Event Collector (HEC) support in IBM Z Operations Analytics Version 4.1.0.4, these macros were updated so that they define **not only the indexes but also the data source types**. The macros include default values, but you can update these values to match local naming schemas.

For information about editing a Splunk macro that is provided by Z Operations Analytics, see "Editing a Splunk macro that is provided by Z Operations Analytics" on page 116.

As described in "Index schema in IBM Z Operations Analytics application on the Splunk platform" on page 114, the default values for indexes are different depending on how the data is streamed to Splunk. The same is true for the default values for source types. The source type of all data that is streamed to a Splunk HEC subscriber is the value of source type, appended with the suffix _KV. For example, the following table illustrates the value of the **sourcetype** field for each type of Splunk subscriber.

Data stream subscriber	Value of sourcetype field in the subscriber definition
IBM Z Common Data Provider Data Receiver	`zOS-SYSLOG-Console`
Splunk HEC	`zOS-SYSLOG-Console_KV`

Important: When you define a data stream in a policy in the IBM Z Common Data Provider Configuration Tool, you can customize the Splunk HEC data source type. However, the customized field is applicable only for a subscriber that is defined with one of the following values for the subscriber protocol:

- `CDP Splunk via HEC via HTTP`
- `CDP Splunk via HEC via HTTPS`

A subscriber that is defined with one of the following values for the subscriber protocol ignores the customized field, which ensures that HEC data is shown for all IBM-provided dashboards and predefined searches:

- `IZOA on Splunk via HEC via HTTP`
- `IZOA on Splunk via HEC via HTTPS`

To customize the data source type for HEC, use the guidelines in the IBM Z Common Data Provider documentation.

By default, the macros search for data that is ingested to Splunk by either the IBM Z Common Data Provider Data Receiver or the Splunk HEC. A search for all data can be useful for a first time install, or when you are migrating from one data ingestion type to another, but the search performance might be degraded. To improve search performance, edit the macro definitions so that they correspond with the protocol (either the Data Receiver or HEC) of your Splunk subscriber.

For more information about the subscribers for each type of source data, see "Subscribers for each type of source data" on page 10.

For information about the available macros for log data and SMF data, including some example macro definitions, see the following topics:

- "Macros for log data" on page 116
- "Macros for SMF data" on page 120

Editing a Splunk macro that is provided by Z Operations Analytics

You can update the Splunk macros that are provided with the IBM Z Operations Analytics application.

Procedure

1. Log in to Splunk.
2. Click **Settings** > **Advanced search** > **Search macros**.
3. Set the App filter to `IBM Z Operations Analytics (ibm_zoa_insights)`, and select **Created in the App**.
4. Click the name of a macro to edit.

   When you save the macro, the `macros.conf` file is stored as a local editable file. For information about editing a configuration file, see How to edit a configuration file in the Splunk documentation.

Macros for log data

This reference includes some example macro definitions for log data and lists the available macros for each log data source type that is used in the IBM Z Operations Analytics dashboards and predefined searches.

For each macro, the following information is also listed:

- The expected values for the index and the log data source type for both an IBM Z Common Data Provider Data Receiver subscriber and a Splunk HEC subscriber
- For data to be streamed to the Splunk subscriber, the name of the associated IBM Z Operations Analytics data stream that must be defined for the respective log data source type in the IBM Z Common Data Provider Configuration Tool policy

- "Example macro definitions for syslogd data" on page 117
- "Available macros for each log data source type" on page 117

Remember: The following indexes for the respective source types were removed from the IBM Z Operations Analytics V4.1.02 application on the Splunk platform and are included in the macros for migration purposes only:

- `zOS-CICS-EYULOG`
- `zOS-CICS-EYULOGDMY`
- `zOS-CICS-EYULOGYMD`
- `zOS-CICS-MSGUSR`
- `zOS-CICS-MSGUSRDMY`
- `zOS-CICS-MSGUSRYMD`
- `zOS-NetView`
- `zOS-syslogd`
- `zOS-WAS-SYSOUT`
- `zOS-WAS-SYSPRINT`
- `zOS-zSecure`

Example macro definitions for syslogd data

The following examples are macro definitions for UNIX System Services system log (syslogd) data:

- By default, the macro is defined to show results in the dashboards and predefined searches for a Data Receiver subscriber (with migration for the pre-V4.1.02 log data) or a HEC subscriber.

  ```
  ((index=zOS-syslogd OR index=zosdex OR index=zosdex_kv) AND
  (sourcetype=zOS-syslogd OR sourcetype=zOS-syslogd_KV))
  ```

- Edit the macro to show results for only a **Data Receiver subscriber**, **with** migration for pre-V4.1.02 log data.

  ```
  ((index=zOS-syslogd OR index=zosdex) AND (sourcetype=zOS-syslogd))
  ```

- Edit the macro to show results for only a **Data Receiver subscriber**, **without** migration for pre-V4.1.02 log data.

  ```
  (index=zosdex AND sourcetype=zOS-syslogd)
  ```

- Edit the macro to show results for only a **HEC subscriber**.

  ```
  (index=zosdex_kv AND sourcetype=zOS-syslogd_KV)
  ```

Available macros for each log data source type

Table 19 on page 117 lists the available macros for each log data source type that is used in the IBM Z Operations Analytics dashboards and predefined searches.

For each macro, the following information is also shown:

- The expected values for the index and the log data source type for both an IBM Z Common Data Provider Data Receiver subscriber and a Splunk HEC subscriber
- For data to be streamed to the Splunk subscriber, the name of the associated IBM Z Operations Analytics data stream that must be defined for the respective log data source type in the IBM Z Common Data Provider Configuration Tool policy

Table 19. Available macros for each log data source type

Macro	Index	Log data source type	Data stream in the policy
zOS-CICS-EYULOG	• Data receiver: – zosdex – zOS-CICS-EYULOG (pre -V4.1.02) • HEC: – zosdex_kv	• Data receiver: – zOS-CICS-EYULOG • HEC: – zOS-CICS-EYULOG_KV	CICS EYULOG

Table 19. Available macros for each log data source type (continued)

Macro	Index	Log data source type	Data stream in the policy
zOS-CICS-EYULOGDMY	- Data receiver: – zosdex – zOS-CICS-EYULOGDMY (pre -V4.1.02) - HEC: – zosdex_kv	- Data receiver: – zOS-CICS-EYULOGDMY - HEC: – zOS-CICS-EYULOGDMY_KV	CICS EYULOG DMY
zOS-CICS-EYULOGYMD	- Data receiver: – zosdex – zOS-CICS-EYULOGYMD (pre -V4.1.02) - HEC: – zosdex_kv	- Data receiver: – zOS-CICS-EYULOGYMD - HEC: – zOS-CICS-EYULOGYMD_KV	CICS EYULOG YMD
zOS-CICS-MSGUSR	- Data receiver: – zosdex – zOS-CICS-MSGUSR (pre -V4.1.02) - HEC: – zosdex_kv	- Data receiver: – zOS-CICS-MSGUSR - HEC: – zOS-CICS-MSGUSR_KV	CICS User Messages
zOS-CICS-MSGUSRDMY	- Data receiver: – zosdex – zOS-CICS-MSGUSRDMY (pre -V4.1.02) - HEC: – zosdex_kv	- Data receiver: – zOS-CICS-MSGUSRDMY - HEC: – zOS-CICS-MSGUSRDMY_KV	CICS User Messages DMY
zOS-CICS-MSGUSRYMD	- Data receiver: – zosdex – zOS-CICS-MSGUSRYMD (pre -V4.1.02) - HEC: – zosdex_kv	- Data receiver: – zOS-CICS-MSGUSRYMD - HEC: – zOS-CICS-MSGUSRYMD_KV	CICS User Messages YMD

Table 19. Available macros for each log data source type (continued)

Macro	Index	Log data source type	Data stream in the policy
zOS-NetView	• Data receiver: – zosdex – zOS-NetView (pre -V4.1.02) • HEC: – zosdex_kv	• Data receiver: – zOS-NetView • HEC: – zOS-NetView_KV	NetView Netlog
zOS-SYSLOG	• Data receiver: – zosdex • HEC: – zosdex_kv	• Data receiver: – zOS-SYSLOG-Console • HEC: – zOS-SYSLOG-Console_KV	z/OS SYSLOG
zOS-syslogd	• Data receiver: – zosdex – zOS-syslogd (pre -V4.1.02) • HEC: – zosdex_kv	• Data receiver: – zOS-syslogd • HEC: – zOS-syslogd_KV	USS Syslogd
zOS-WAS-SYSOUT	• Data receiver: – zosdex – zOS-WAS-SYSOUT (pre -V4.1.02) • HEC: – zosdex_kv	• Data receiver: – zOS-WAS-SYSOUT • HEC: – zOS-WAS-SYSOUT_KV	WebSphere SYSOUT
zOS-WAS-SYSPRINT	• Data receiver: – zosdex – zOS-WAS-SYSPRINT (pre -V4.1.02) • HEC: – zosdex_kv	• Data receiver: – zOS-WAS-SYSPRINT • HEC: – zOS-WAS-SYSPRINT_KV	WebSphere SYSPRINT
zOS-zSecure	• Data receiver: – zosdex – zOS-zSecure (pre -V4.1.02) • HEC: – zosdex_kv	• Data receiver: – zOS-zSecure • HEC: – zOS-zSecure_KV	zSecure Access Monitor

Macros for SMF data

This reference includes some example macro definitions for SMF data and lists the available macros for each SMF data source type that is used in the IBM Z Operations Analytics dashboards and predefined searches.

For each macro, the following information is also listed:

- The expected values for the index and the SMF data source type for both an IBM Z Common Data Provider Data Receiver subscriber and a Splunk HEC subscriber
- For data to be streamed to the Splunk subscriber, the name of the associated IBM Z Operations Analytics data stream that must be defined for the respective SMF data source type in the IBM Z Common Data Provider Configuration Tool policy
- "Example macro definitions for SMF100_1 data" on page 120
- "Available macros for each SMF data source type" on page 120

Example macro definitions for SMF100_1 data

The following examples are macro definitions for SMF100_1 data:

- By default, the macro is defined to show results in the dashboards and predefined searches for a Data Receiver subscriber or a HEC subscriber.

```
((index=zos-smf100_1 OR index=zosdex_kv) AND
(sourcetype=zOS-SMF100_1 OR sourcetype=zOS-SMF100_1_KV))
```

- Edit the macro to show results for only a **Data Receiver subscriber**.

```
(index=zos-smf100_1 AND sourcetype=zOS-SMF100_1)
```

- Edit the macro to show results for only a **HEC subscriber**.

```
(index=zosdex_kv AND sourcetype=zOS-SMF100_1_KV)
```

Available macros for each SMF data source type

Table 20 on page 120 lists the available macros for each SMF data source type that is used in the IBM Z Operations Analytics dashboards and predefined searches.

For each macro, the following information is also shown:

- The expected values for the index and the SMF data source type for both an IBM Z Common Data Provider Data Receiver subscriber and a Splunk HEC subscriber
- For data to be streamed to the Splunk subscriber, the name of the associated IBM Z Operations Analytics data stream that must be defined for the respective SMF data source type in the IBM Z Common Data Provider Configuration Tool policy

Table 20. Available macros for each SMF data source type

Macro	Index	SMF data source type	Data stream in the policy
zOS-SMF100_1	• Data receiver: – zos-smf100_1 • HEC: – zosdex_kv	• Data receiver: – zOS-SMF100_1 • HEC: – zOS-SMF100_1_KV	SMF100_1

Table 20. Available macros for each SMF data source type (continued)

Macro	Index	SMF data source type	Data stream in the policy
zOS-SMF101_SUMMARY	- Data receiver: - zos-smf101_summary - HEC: - zosdex_kv	- Data receiver: - zOS-SMF101_SUMMARY - HEC: - zOS-SMF101_SUMMARY_KV	SMF101_SUMMARY
zOS-SMF110_E	- Data receiver: - zos-smf110_e - HEC: - zosdex_kv	- Data receiver: - zOS-SMF110_E - HEC: - zOS-SMF110_E_KV	SMF110_E
zOS-SMF110_S_10	- Data receiver: - zos-smf110_s_10 - HEC: - zosdex_kv	- Data receiver: - zOS-SMF110_S_10 - HEC: - zOS-SMF110_S_10_KV	SMF110_S_10
zOS-SMF110_1_SUMMARY	- Data receiver: - zos-smf110_1_summary - HEC: - zosdex_kv	- Data receiver: - zOS-SMF110_1_SUMMARY - HEC: - zOS-SMF110_1_SUMMARY_KV	SMF110_1_SUMMARY
zOS-SMF120_REQAPPL	- Data receiver: - zos-smf120_reqappl - HEC: - zosdex_kv	- Data receiver: - zOS-SMF120_REQAPPL - HEC: - zOS-SMF120_REQAPPL_KV	SMF120_REQAPPL
zOS-SMF120_REQCONT	- Data receiver: - zos-smf120_reqcont - HEC: - zosdex_kv	- Data receiver: - zOS-SMF120_REQCONT - HEC: - zOS-SMF120_REQCONT_KV	SMF120_REQCONT

Table 20. Available macros for each SMF data source type (continued)

Macro	Index	SMF data source type	Data stream in the policy
zOS-SMF123_01_V2	• Data receiver: – zos-smf123_01_v2 • HEC: – zosdex_kv	• Data receiver: – zOS-SMF123_01_V2 • HEC: – zOS-SMF123_01_V2_KV	SMF123_01_V2
zOS-SMF123_01_V2_SUMRY	• Data receiver: – zos-smf123_01_v2_sumry • HEC: – zosdex_kv	• Data receiver: – zOS-SMF123_01_V2_SUMRY • HEC: – zOS-SMF123_01_V2_SUMRY_KV	SMF123_01_V2_SUMRY
zOS-SMF123_01_V2_OR_V2_SUMRY	• Data receiver: – Either zos-smf123_01_v2 or zos-smf123_01_v2_sumry, but **not both** • HEC: – zosdex_kv	• Data receiver: – Either zOS-SMF123_01_V2 or zOS-SMF123_01_V2_SUMRY, but **not both** • HEC: – Either zOS-SMF123_01_V2_KV or zOS-SMF123_01_V2_SUMRY_KV, but **not both**	• Either SMF123_01_V2 or SMF123_01_V2_SUMRY, but **not both**
zOS-SMF30	• Data receiver: – zos-smf30 • HEC: – zosdex_kv	• Data receiver: – zOS-SMF30 • HEC: – zOS-SMF30_KV	SMF30
zOS-SMF80_COMMAND	• Data receiver: – zos-smf80_command • HEC: – zosdex_kv	• Data receiver: – zOS-SMF80_COMMAND • HEC: – zOS-SMF80_COMMAND_KV	SMF80_COMMAND

Table 20. Available macros for each SMF data source type (continued)

Macro	Index	SMF data source type	Data stream in the policy
zOS-SMF80_LOGON	- Data receiver: - zos-smf80_logon - HEC: - zosdex_kv	- Data receiver: - zOS-SMF80_LOGON - HEC: - zOS-SMF80_LOGON_KV	SMF80_LOGON
zOS-SMF80_OPERATION	- Data receiver: - zos-smf80_operation - HEC: - zosdex_kv	- Data receiver: - zOS-SMF80_OPERATION - HEC: - zOS-SMF80_OPERATION_KV	SMF80_OPERATION
zOS-SMF80_OMVS_RES_1	- Data receiver: - zos-smf80_omvs_res_1 - HEC: - zosdex_kv	- Data receiver: - zOS-SMF80_OMVS_RES_1 - HEC: - zOS-SMF80_OMVS_RES_1_KV	SMF80_OMVS_RES_1
zOS-SMF80_OMVS_RES_2	- Data receiver: - zos-smf80_omvs_res_2 - HEC: - zosdex_kv	- Data receiver: - zOS-SMF80_OMVS_RES_2 - HEC: - zOS-SMF80_OMVS_RES_2_KV	SMF80_ OMVS_RES_2
zOS-SMF80_OMVS_SEC_1	- Data receiver: - zos-smf80_omvs_sec_1 - HEC: - zosdex_kv	- Data receiver: - zOS-SMF80_OMVS_SEC_1 - HEC: - zOS-SMF80_OMVS_SEC_1_KV	SMF80_OMVS_SEC_1

Table 20. Available macros for each SMF data source type (continued)

Macro	Index	SMF data source type	Data stream in the policy
zOS-SMF80_OMVS_SEC_2	• Data receiver: – zos-smf80_omvs_sec_2 • HEC: – zosdex_kv	• Data receiver: – zOS-SMF80_OMVS_SEC_2 • HEC: – zOS-SMF80_OMVS_SEC_2_KV	SMF80_OMVS_SEC_2
zOS-SMF80_RESOURCE	• Data receiver: – zos-smf80_resource • HEC: – zosdex_kv	• Data receiver: – zOS-SMF80_RESOURCE • HEC: – zOS-SMF80_RESOURCE_KV	SMF80_RESOURCE

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.

Remember: In this procedure, you must install the IBM Z Common Data Provider Buffered Splunk Ingestion App **only if** you are using the IBM Z Common Data Provider Data Receiver as a subscriber.

For more information, see Preparing to send data to Splunk in the IBM Z Common Data Provider documentation.

Procedure

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

1. **If you are using the IBM Z Common Data Provider Data Receiver as a subscriber**, install and configure the Data Receiver and the IBM Z Common Data Provider Buffered Splunk Ingestion App.
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 138.

What to do next

If you are using the Splunk HTTP Event Collector (HEC) as the subscriber (as indicated in "Subscribers for each type of source data" on page 10), also complete the steps in "Sending data directly to Splunk by using Splunk HEC as the subscriber" on page 128.

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.

About this task

Remember: In this procedure, you must install the IBM Z Common Data Provider Buffered Splunk Ingestion App **only if** you are using the IBM Z Common Data Provider Data Receiver as a subscriber.

For more information, see Preparing to send data to Splunk in the IBM Z Common Data Provider documentation.

Procedure

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

1. **On the Splunk indexer cluster master node**, complete the following steps.

   a) **If you are using the IBM Z Common Data Provider Data Receiver as a subscriber**, install the IBM Z Common Data Provider Buffered Splunk Ingestion App, as indicated in the following instructions:

      i) From the IBM Z Common Data Provider `/usr/lpp/IBM/cdpz/v2r1m0/DS/LIB` directory, download the IBM Z Common Data Provider 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`

      ii) 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 (one of the following files) that you downloaded in an earlier step, select that file, and click **Upload**:

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

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

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

ii) Log in to Splunk.

iii) From the Splunk Web Home page, click the gear icon that is next to the word "Apps."

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

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

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

vii) 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
```

e) Copy the app files from `$SPLUNK_HOME/etc/apps/ibm_zoa_insights` to `$SPLUNK_HOME/etc/master-apps/ibm_zoa_insights`.

f) In the `$SPLUNK_HOME/etc/master-apps/ibm_zoa_insights/default` directory, edit the `indexes.conf` file, and add the line `repFactor=auto` to each of the indexes.

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

h) 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 a standalone search head**, complete the following steps.

a) **If you are using the IBM Z Common Data Provider Data Receiver as a subscriber**, install the IBM Z Common Data Provider Buffered Splunk Ingestion App, as indicated in the following instructions:

i) From the IBM Z Common Data Provider `/usr/lpp/IBM/cdpz/v2r1m0/DS/LIB` directory, download the IBM Z Common Data Provider 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`

ii) 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 (one of the following files) that you downloaded in an earlier step, select that file, and click **Upload**:

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

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

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

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

ii) Log in to Splunk.

iii) From the Splunk Web Home page, click the gear icon that is next to the word "Apps."

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

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

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

vii) 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. **In a search head cluster**, complete the following steps on the Splunk Enterprise deployer instance.

   a) **If you are using the IBM Z Common Data Provider Data Receiver as a subscriber**, install the IBM Z Common Data Provider Buffered Splunk Ingestion App, as indicated in the following instructions:

      i) From the IBM Z Common Data Provider `/usr/lpp/IBM/cdpz/v2r1m0/DS/LIB` directory, download the IBM Z Common Data Provider 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`

      ii) 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 (one of the following files) that you downloaded in an earlier step, select that file, and click **Upload**:

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

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

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

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

      ii) Log in to Splunk.

      iii) From the Splunk Web Home page, click the gear icon that is next to the word "Apps."

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

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

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

      vii) 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
      ```

   c) Create a `.tar` file of both the `/opt/splunk/etc/apps/ibm_cdpz_buffer` directory and the `/opt/splunk/etc/apps/ibm_zoa_insights` directory, and move these `.tar` files to the `opt/splunk/etc/shcluster/apps` directory.

   Use the following commands:

   ```
   cd /opt/splunk/etc/apps
   tar -cvf ibm_cdpz_buffer.tar ibm_cdpz_buffer
   mv ibm_cdpz_buffer.tar /opt/splunk/etc/shcluster/apps/
   tar -cvf ibm_zoa_insights.tar ibm_zoa_insights
   mv ibm_zoa_insights.tar /opt/splunk/etc/shcluster/apps/
   ```

   d) Extract the contents of the `.tar` files that you created in the previous step into the `opt/splunk/etc/shcluster/apps` directory.

   Use the following commands:

   ```
   cd /opt/splunk/etc/shcluster
   ```

```
tar -xvf ibm_cdpz_buffer.tar
tar -xvf ibm_zoa_insights.tar
```

   e) Run the **`splunk apply shcluster-bundle`** command.

4. **On the heavy forwarder**, complete the following steps.
   a) **If you are using the IBM Z Common Data Provider Data Receiver as a subscriber**, install and configure the IBM Z Common Data Provider Data Receiver and Buffered Splunk Ingestion App.
   b) Install the IBM Z Operations Analytics application, as indicated in the following instructions:
      i) Mount the IBM Z Operations Analytics ISO installation image, or extract the IBM Z Operations Analytics `.tar` file.
      ii) Log in to Splunk.
      iii) From the Splunk Web Home page, click the gear icon that is next to the word "Apps."
      iv) Select **Install app from file**.
      v) Navigate to the ISO image, select the `ibm_zoa_insights.spl` file, and click **Upload**.
      vi) If you are prompted to restart Splunk Enterprise, restart it.
      vii) 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
      ```

5. 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 138.

What to do next

If you are using the Splunk HTTP Event Collector (HEC) as the subscriber (as indicated in "Subscribers for each type of source data" on page 10), also complete the steps in "Sending data directly to Splunk by using Splunk HEC as the subscriber" on page 128.

Sending data directly to Splunk by using Splunk HEC as the subscriber

You can send data directly to Splunk by using the Splunk HTTP Event Collector (HEC). Use of the Splunk HEC enables the quick setup of your environment without the need for the IBM Z Common Data Provider Data Receiver and Buffered Splunk Ingestion App. However, this method can increase the data ingestion size, the cost, and the CPU usage on the z/OS system.

About this task

To stream data to Splunk directly by using the HEC, you must enable the HEC in Splunk, and create a token that allows an application to communicate with Splunk without using user credentials.

For more information about the Splunk HEC, see Set up and use HTTP Event Collector in Splunk Web in the Splunk documentation. Especially see the following subtopics:

- Enable HTTP Event Collector
- Create an Event Collector token

Procedure

1. Enable the HEC in Splunk.
2. If possible, install the IBM Z Operations Analytics application (the `ibm_zoa_insights.spl` file) before configuring a HEC token so that the IBM-provided index for HEC data is available.
3. Configure at least one token to allow the IBM Z Operations Analytics application to communicate with Splunk. Use the following values for the following fields:

Default index field

Set this value to `zosdex_kv`. This index is provided with the IBM Z Operations Analytics application and is used to display HEC data for all predefined dashboards and searches in IBM Z Operations Analytics.

App Context field

Set this value to the IBM Z Operations Analytics application, for example, `IBM Z Operations Analytics (ibm_zoa_insights)`.

Tip: If you delete the IBM Z Operations Analytics application, the Splunk HEC token is automatically deleted.

4. Enabling HEC on the Splunk Enterprise system generates a token value for communicating with Splunk. Take note of this token value because you must provide it as part of the subscriber configuration when you create a policy in the IBM Z Common Data Provider Configuration Tool.

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, see Overview of Service Insights in ITSI 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 Overview of entity integrations in ITSI in the Splunk documentation.

Table 21 on page 130 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 22 on page 136 describes the KPI base searches that are listed in column 3 of Table 21 on page 130, and includes the associated metrics that are defined for each search.

Table 21. 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

Table 21. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	Description	KPI base search	Search and calculate
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
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

Table 21. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	Description	KPI base search	Search and calculate
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
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

Table 21. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	Description	KPI base search	Search and calculate
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 21. 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

Table 21. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	**Description**	**KPI base search**	**Search and calculate**
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
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

Table 21. KPIs in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI	Description	KPI base search	Search and calculate
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
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 22. 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

Table 22. KPI base searches in the Splunk ITSI module for IBM Z Operations Analytics (continued)

KPI base search	Description	Defined metrics
`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
`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_4.1.03.spl` as a Splunk application.

Before you begin

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

- IBM Z Common Data Provider V2.1.0 must be running on your z/OS system.
- On the Splunk Enterprise server or Splunk heavy forwarder, the following IBM Z Common Data Provider 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_4.1.03.spl` as a Splunk application.

   You can install from Splunk Web, from the Splunk command line, or by extracting the `DA-ITSI-IBMZOA_4.1.03.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_4.1.03.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 GUI, use the backup and restore function to create a restore job for reading the `backup_json_data_for_DA-ITSI-IBMZOA_4.1.03.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_4.1.03.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.	Unsuccessful restoration of the IBM-provided backup JSON data for the Splunk ITSI module for IBM Z Operations Analytics

Symptom	Problem
• 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.	
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_4.1.03.spl`)
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`)

Chapter 10. 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, APIs, databases, messaging, networks, security, transactions, and web servers. For each of these domains, this reference lists the data sources from which the operational data is retrieved, the dashboards that represent that operational data, and the predefined searches for searching that operational data.

Operational insights are presented in different ways, depending on your platform, as described in Table 23 on page 141.

Table 23. Presentation of operational insights depending on your platform and on whether you are using the machine learning system

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: – extensions for z/OS Problem Insights and client-side Expert Advice – Installing the , extensions, and data gatherer – Configuring the Problem Insights extension – Getting started with Problem Insights for z/OS • 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 183 – “Searches that are predefined for searching the operational data” on page 185
Elastic Stack and Splunk	• Problem Insights server, 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 Chapter 4, “Deploying the Problem Insights server,” on page 15. • 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 183 – “Searches that are predefined for searching the operational data” on page 185

Table 23. Presentation of operational insights depending on your platform and on whether you are using the machine learning system (continued)

Platform	Presentation
Machine learning system component of IBM Z Operations Analytics	• Problem Insights server, 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 Chapter 4, “Deploying the Problem Insights server,” on page 15.

Domains of interest with associated data sources, dashboards, and searches

Each IT domain of interest (such as the z/OS system or IBM Db2 for z/OS subsystem) is listed with links to more information about its associated IBM Z Operations Analytics data sources, dashboards, and predefined searches.

For each domain of interest, Table 24 on page 142 indicates the following information:

- The data sources that contribute to the operational insights for the domain
- The dashboards that represent the operational data for the domain
- The searches that are predefined for searching the operational data for the domain

Tip: For more information about a data source (including the configuration that you must do to send the data from that data source to IBM Z Operations Analytics), dashboard, or predefined search, click the relevant links in Table 24 on page 142.

You can also find the information about all data sources, dashboards, and searches in the following topics:

- “Data sources that contribute to the operational insights” on page 144
- “Dashboards that represent the operational data” on page 183
- “Searches that are predefined for searching the operational data” on page 185

Table 24. Domains of interest with associated data sources, dashboards, and searches

Domain of interest	Data sources that contribute to the insights	Associated dashboards	Associated predefined searches
z/OS system	• z/OS SYSLOG • “SMF 30 data” on page 147 • ***Log Analysis platform:*** zAware	• Log Analysis dashboards • “Elastic Stack dashboards” on page 184 • “Splunk dashboards” on page 185	“z/OS system searches” on page 192

Table 24. *Domains of interest with associated data sources, dashboards, and searches (continued)*

Domain of interest	Data sources that contribute to the insights	Associated dashboards	Associated predefined searches
Application program interfaces (APIs)	• "SMF 123 data" on page 171 (for IBM z/OS Connect Enterprise Edition)	• ***Log Analysis platform:*** No dashboards are provided for IBM z/OS Connect Enterprise Edition. • "Elastic Stack dashboards" on page 184 • "Splunk dashboards" on page 185	None
IBM CICS Transaction Server for z/OS	• z/OS SYSLOG • CICS Transaction Server for z/OS EYULOG and MSGUSR logs • "SMF 110 data" on page 161	• Log Analysis dashboards • "Elastic Stack dashboards" on page 184 • "Splunk dashboards" on page 185	"CICS Transaction Server for z/OS searches" on page 185
IBM Db2 for z/OS	• z/OS SYSLOG • ***Elastic Stack and Splunk platforms, and machine learning system:*** "SMF 100 data" on page 157 • ***Elastic Stack and Splunk platforms:*** "SMF 101 data" on page 160	• Log Analysis dashboards • "Elastic Stack dashboards" on page 184 • "Splunk dashboards" on page 185	"Db2 for z/OS searches" on page 187
IMS for z/OS	• z/OS SYSLOG	• Log Analysis dashboards • "Elastic Stack dashboards" on page 184 • "Splunk dashboards" on page 185	"IMS for z/OS searches" on page 187
IBM MQ for z/OS	• z/OS SYSLOG	• Log Analysis dashboards • "Elastic Stack dashboards" on page 184 • "Splunk dashboards" on page 185	"MQ for z/OS searches" on page 188

Table 24. Domains of interest with associated data sources, dashboards, and searches (continued)

Domain of interest	Data sources that contribute to the insights	Associated dashboards	Associated predefined searches
IBM WebSphere Application Server for z/OS	• ***Log Analysis platform:*** 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 • "SMF 120 data" on page 168	• Log Analysis dashboards • "Elastic Stack dashboards" on page 184 • "Splunk dashboards" on page 185	"WebSphere Application Server for z/OS searches" on page 191
Network	• z/OS SYSLOG • UNIX System Services system log (`syslogd`) • NetView for z/OS program	• Log Analysis dashboards • "Elastic Stack dashboards" on page 184 • "Splunk dashboards" on page 185	• "NetView for z/OS searches" on page 189 • "z/OS network searches" on page 192
Security	• z/OS SYSLOG • UNIX System Services system log (`syslogd`) • Access Monitor component of IBM Security zSecure Admin • "SMF 80 data" on page 148	• Log Analysis dashboards • "Elastic Stack dashboards" on page 184 • "Splunk dashboards" on page 185	• "Security searches: RACF" on page 189 • "Security searches: zSecure Access Monitor" on page 190

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 Z Common Data Provider

In the IBM Z Common Data Provider 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 Z Common Data Provider, 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” on page 146
- “NetView message data” on page 146
- “SMF 30 data” on page 147
- “SMF 80 data” on page 148
- “SMF 110 data” on page 161
- “SMF 120 data” on page 168
- “SYSLOG data” on page 178
- “syslogd data” on page 179
- “WebSphere SYSOUT data” on page 180
- “WebSphere SYSPRINT data” on page 181
- “zSecure data” on page 183
- For the Elastic Stack and Splunk platforms only:
  - “SMF 100 data” on page 157
  - “SMF 101 data” on page 160
  - “SMF110_1_SUMMARY record type” on page 166
- For the IBM Operations Analytics - Log Analysis platform only:
  - “WebSphere HPEL data” on page 180
- For the IBM Z Operations Analytics machine learning system only:
  - “SMF 100 machine learning data” on page 159
  - “SMF 110 machine learning data” on page 167

Data that is provided by IBM z Advanced Workload Analysis Reporter (IBM zAware)

“zAware interval anomaly data” on page 181 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 25. 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.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 10.

NetView message data

NetView message data includes network data from IBM Z NetView.

Data stream definition for NetView message data

Table 26 on page 146 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 26. 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.
Filter Transform	Not required

Table 26. Data stream definition for NetView message data (continued)

Type of node in the policy	Required configuration value
Subscriber	See "Subscribers for each type of source data" on page 10.

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" on page 147
- "Data stream definition for SMF 30 data" on page 147
- "Annotated fields for SMF 30 data" on page 147

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 Z Common Data Provider System Data Engine processes.

Data stream definition for SMF 30 data

For prerequisite requirements for defining SMF data streams, see "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

Table 27 on page 147 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 27. 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.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 10.

Annotated fields for SMF 30 data

Table 28. 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

Table 28. Annotated fields for SMF 30 data (continued)

Field	Description
`SystemID`	The MVS system ID, which is also the SMF system ID
`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" on page 148
- "Data stream definition for SMF 80 data" on page 148
- "Annotated fields for SMF 80 data" on page 149

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 Z Common Data Provider 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 151.

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 "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

Table 29 on page 149 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 29. 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.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 10.

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 30. Annotated fields for SMF 80 data

Field	Description	Corresponding SMF field
`AccessAllow`	Access authority allowed	SMF80DTA
`AccessReq`	Access authority requested	SMF80DTA
`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

Table 30. Annotated fields for SMF 80 data (continued)

Field	Description	Corresponding SMF field
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
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

Table 30. Annotated fields for SMF 80 data (continued)

Field	Description	Corresponding SMF field
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 Z Common Data Provider 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 MVS system ID, which is also the SMF system ID	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 Z Common Data Provider 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 31 on page 152 describes the event code qualifiers for events 8 - 25, which provide more information about why the command failed.

Table 31. 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 32 on page 152 describes the event code qualifiers for event 1, which provide more information about why the logon failed.

Table 32. 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
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

Table 32. SMF80_LOGON record type: event code qualifiers for event 1 (continued)

Event code qualifier	Description
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 33 on page 153 describes the event code qualifiers for events 28 - 30, which provide more information about the operation results.

Table 33. 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 34 on page 154, Table 35 on page 154, Table 36 on page 154, and Table 37 on page 154 describe the event code qualifiers for events 31 and 33 - 35, which provide more information about the operation results.

Table 34. 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 35. 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 36. 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
2	Security label failure

Table 37. 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 38 on page 155, Table 39 on page 155, Table 40 on page 155, Table 41 on page 156, Table 42 on page 156, and Table 43 on page 156 describe the event code qualifiers for events 2 - 7, which provide more information about the operation results.

Table 38. 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 39. 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 40. 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

Table 40. SMF80_OPERATION record type: event code qualifiers for event 4 (continued)

Event code qualifier	Description
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 41. 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 42. SMF80_OPERATION record type: event code qualifiers for event 6

Event code qualifier	Description
0	Successful deletion

Table 43. 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

Table 43. SMF80_OPERATION record type: event code qualifiers for event 7 (continued)

Event code qualifier	Description
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 44 on page 157 describes the event code qualifiers for event 2, which provide more information about the operation results.

Table 44. 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
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 Db2 for z/OS statistics data that is generated in a time interval. These records provide information about potential problems in Db2 subsystems.

- “SMF 100 data generation” on page 158

- “SMF 100_1 record type” on page 158
- “SMF 100 machine learning data” on page 159

SMF 100 data generation

When it sends data to IBM Z Operations Analytics, the IBM Z Common Data Provider System Data Engine collects only a subset of the SMF record type 100 data that is generated by Db2 for z/OS. 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)
```

Important:

For the IBM Z Operations Analytics machine learning system, SMF type 100 records must be generated in 1-minute intervals.

SMF 100_1 record type

SMF 100_1 records contain Db2 for z/OS statistics data for CICS Transaction Server for z/OS. These records provide information about potential problems in Db2 subsystems that are due to database deadlocks, database timeouts, and suspended database operations.

- “Data stream definition for SMF 100_1 data” on page 158
- “Annotated fields for SMF 100_1 data” on page 158

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 “Preparing the Configuration Tool to support SMF record types for Z Operations Analytics” on page 7.

Table 45 on page 158 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 45. 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.
Filter Transform	Not required
Subscriber	Elastic Stack platform or Splunk platform only. For the appropriate values, see “Subscribers for each type of source data” on page 10.

Annotated fields for SMF 100_1 data

Table 46. Annotated fields for SMF 100_1 data

Field	Description
DB2_SYSTEM_ID	Db2 system ID, which is also the SMF subsystem ID
DEADLOCKS	The number of times that deadlocks were detected

Table 46. Annotated fields for SMF 100_1 data (continued)

Field	Description
hostname	Host name
LATCH_SUSPENDS	The number of latch suspensions
LOCK_SUSPENDS	The number of times that a lock cannot be obtained, and the unit of work is suspended
LOCK_TIMEOUTS	The number of times that a unit of work was suspended for a period of time that exceeds the timeout value
MVS_SYSTEM_ID	MVS system ID, which is also the SMF system ID
OTHER_SUSPENDS	The number of suspensions that are caused by something other than lock or latch
sourcename	Source name
SSID	Subsystem ID
sysplex	Sysplex name
system	System name
TIMESTAMP	Time stamp
timezone	Time zone offset
TOTAL_SUSPENDS	The sum of the values of the following fields: • LOCK_SUSPENDS • LATCH_SUSPENDS • OTHER_SUSPENDS

SMF 100 machine learning data

The IBM Z Operations Analytics machine learning system uses SMF 100 data to provide insight into normal and abnormal resource usage in Db2 for z/OS subsystems, based on the CPU time, distributed data facility (DDF), locks, latches, and storage usage.

Data stream definition for SMF 100 machine learning data

Tip: These data streams can be defined only for the IBM Z Operations Analytics machine learning system.

For prerequisite requirements for defining SMF data streams, see “Preparing the Configuration Tool to support SMF record types for Z Operations Analytics” on page 7.

Table 47 on page 160 indicates the configuration values to use in defining these data streams in the IBM Z Common Data Provider Configuration Tool.

Table 47. Data stream definition for SMF 100 machine learning data

Type of node in the policy	Required configuration value
Data Stream	All of the following values: • **A_DB2_DB_SYS_I** • **A_DB2_DB_BIND_I** • **A_DB2_DB_I** • **A_DB2_BP_I** • **A_DB2_SHR_LOCK_I** • **A_DB2_SHR_BP_I** • **A_DB2_STORAGE_I** • **A_DB2_LATCH_I** **To select these data streams in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics – Machine Learning** > **Database** > **Db2**, and select the check box for each stream.
Filter Transform	Not required
Subscriber	Machine learning system only. For the appropriate values, see "Subscribers for each type of source data" on page 10.

SMF 101 data

System Management Facilities (SMF) record type 101 data is Db2 for z/OS accounting data about the resources that are used during a transaction.

- "SMF 101 data generation" on page 160
- "Data stream definition for SMF 101_SUMMARY data" on page 160

SMF 101 data generation

When it sends data to IBM Z Operations Analytics, the IBM Z Common Data Provider System Data Engine collects only a subset of the SMF record type 101 data that is generated by Db2 for z/OS. It collects data from SMF type 101 and summarizes these records before it sends them to the Elastic Stack or Splunk platforms. To enable the generation of this data, you must enable the following trace options in Db2 for z/OS:

```
START TRACE(ACCTG) DEST(SMF) CLASS(1,2,3)
```

Data stream definition for SMF 101_SUMMARY 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 "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

"Data stream definition for SMF 101_SUMMARY data" on page 160 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 48. Data stream definition for SMF 101_SUMMARY data

Type of node in the policy	Required configuration value
Data Stream	**SMF 101_SUMMARY** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **Database** > **Db2** > **Accounting**, and select the **SMF 101_SUMMARY** check box.
Filter Transform	Not required
Subscriber	Elastic Stack platform or Splunk platform only. For the appropriate values, see "Subscribers for each type of source data" on page 10.

SMF 110 data

System Management Facilities (SMF) record type 110 data is generated by CICS Transaction Server for z/OS for monitoring region-wide transaction statistics, transaction-level statistics, and storage exceptions. This data highlights potential problems in CICS system operation and can help you identify system constraints that affect the performance of your transactions.

- "SMF 110 data generation" on page 161
- "SMF110_E record type" on page 162
- "SMF110_S_10 record type" on page 164
- "SMF110_1_SUMMARY record type" on page 166
- "SMF 110 machine learning data" on page 167

SMF 110 data generation

When it sends data to IBM Z Operations Analytics, the IBM Z Common Data Provider System Data Engine collects only a subset of the SMF record type 110 data that is generated by CICS Transaction Server for z/OS. For example, the following data is some of the data that is collected 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`.

  The `SMF110_E` record type contains this data.
- 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`

  The `SMF110_S_10` record type contains this data.

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.

Important:

For the IBM Z Operations Analytics machine learning system, SMF type 110 records must be generated in 1-minute intervals.

SMF110_E record type

SMF110_E records (monitoring exceptions data) 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. CICS writes one exception record for each exception condition that occurs.

- "Data stream definition for SMF 110_E data" on page 162
- "Annotated fields for SMF 110_E data" on page 162

Data stream definition for SMF 110_E data

For prerequisite requirements for defining SMF data streams, see "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

Table 49 on page 162 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 49. 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.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 10.

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 50. 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

Table 50. Annotated fields for SMF 110_E data (continued)

Field	Description	Corresponding SMF field
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
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 MVS system ID, which is also the SMF system ID	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

Table 50. Annotated fields for SMF 110_E data (continued)

Field	Description	Corresponding SMF field
`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` record type

`SMF110_S_10` records (global transaction manager statistics data) 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 164
- "Annotated fields for SMF110_S_10 data" on page 164

Data stream definition for SMF110_S_10 data

For prerequisite requirements for defining SMF data streams, see "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

Table 51 on page 164 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 51. 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.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 10.

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 52. Annotated fields for SMF110_S_10 data

Field	Description	Corresponding SMF field
`ApplID`	The product name (Generic APPLID)	SMFSTPRN

Table 52. Annotated fields for SMF110_S_10 data (continued)

Field	Description	Corresponding SMF field
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
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 MVS system ID, which is also the SMF system ID	SMFMNSID
TransCount	The number of user and system transactions that are attached	XMGNUM

Table 52. Annotated fields for SMF110_S_10 data (continued)

Field	Description	Corresponding SMF field
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
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

SMF110_1_SUMMARY record type

SMF110_1_SUMMARY records contain information about CICS Transaction Server for z/OS transaction summary records for monitoring transaction counts, response times, and CPU usage.

Data stream definition for SMF110_1_SUMMARY 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 "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

Table 53 on page 167 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 53. Data stream definition for SMF110_1_SUMMARY data

Type of node in the policy	Required configuration value
Data Stream	**SMF110_1_SUMMARY** **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_1_SUMMARY** check box.
Filter Transform	Not required
Subscriber	Elastic Stack platform or Splunk platform only. For the appropriate values, see "Subscribers for each type of source data" on page 10.

SMF 110 machine learning data

The IBM Z Operations Analytics machine learning system uses SMF 110 data to provide insight into abnormal system constraints that affect the performance of your CICS regions, based on transactions counts and rates, resource manager interface (RMI) counts and rates, and CPU usage.

Data stream definition for SMF 110 machine learning data

Tip: These data streams can be defined only for the IBM Z Operations Analytics machine learning system.

For prerequisite requirements for defining SMF data streams, see "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

Table 54 on page 167 indicates the configuration values to use in defining these data streams in the IBM Z Common Data Provider Configuration Tool.

Table 54. Data stream definition for SMF 110 machine learning data

Type of node in the policy	Required configuration value
Data Stream	All of the following values: • **A_KC_S_STOR_DSA_T2** • **A_KC_S_STOR_G14_T** • **A_KC_S_STOR_G14_T2** • **A_KC_S_STOR_D14_T** • **A_KC_S_STOR_DSA_T1** • **A_KC_S_STOR_DSA_T0** • **A_KC_MON_TRAN_I** • **A_KC_T_TRAN_T** • **A_KC_RMI_PERF_T** **To select these data streams in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics – Machine Learning** > **CICS Transaction Server**, and select the check box for each stream.
Filter Transform	Not required

Table 54. Data stream definition for SMF 110 machine learning data (continued)

Type of node in the policy	Required configuration value
Subscriber	Machine learning system only. For the appropriate values, see "Subscribers for each type of source data" on page 10.

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" on page 168
- "Data stream definition for SMF 120 data" on page 169
- "Annotated fields for SMF 120 data" on page 169

SMF record type 120 data generation

When it sends data to IBM Z Operations Analytics, the IBM Z Common Data Provider 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 Z Common Data Provider 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 "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

"SMF 120 data" on page 168 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 55. 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.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 10.

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 56. 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

Table 56. Annotated fields for SMF 120 data (continued)

Field	Description	Corresponding SMF field
`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
`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 MVS system ID, which is also the SMF system ID	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 56. 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
`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

SMF 123 data

System Management Facilities (SMF) record type 123 data is generated by the audit interceptor that is included with IBM z/OS Connect Enterprise Edition. These records can be used for auditing, to gain insights into your API workloads for capacity planning, and for troubleshooting.

- “SMF 123 data generation” on page 172

- “SMF123_01_V2 record type” on page 172
- “SMF123_01_V2_SUMRY record type” on page 176

SMF 123 data generation

When it sends data to IBM Z Operations Analytics, the IBM Z Common Data Provider System Data Engine either can send individual transaction records (SMF123_01_V2 record type) or can generate summary records of the transaction records in a specified time interval (SMF123_01_V2_SUMRY record type). The IBM Z Operations Analytics dashboards can use either record type. However, for optimal performance, use the summary records (SMF123_01_V2_SUMRY record type) unless you must track each transaction, and you know that your environment has the capacity to process the additional volume of data.

Important: Do not send both types of record to your environment. Sending both types causes duplicate data to be shown in the IBM Z Operations Analytics dashboards.

To enable the generation of SMF 123 data, you must enable the IBM z/OS Connect Enterprise Edition audit interceptor for your API or service.

Restriction: IBM Z Operations Analytics processes only SMF type 123 subtype 1 version 2 records. Therefore, in your audit interceptor configuration, the value of the apiProviderSmfVersion attribute must be 2.

Important: To enable support for the SMF type 123 data on the IBM Operations Analytics - Log Analysis platform, an insight pack must be created and installed by using the **installCDPSourceTypes** script that is provided with IBM Z Operations Analytics.

For example, to install support for the SMF123_01_V2_SUMRY record, use the following command:

```
./installCDPSourceTypes.sh SMF123_01_V2_SUMRY
```

For more information about the installCDPSourceTypes.sh script, see Installing Insight Packs for SMF data source types that are provided by.

SMF123_01_V2 record type

SMF123_01_V2 records (SMF type 123 subtype 1 version 2 records) contain data about individual API provider requests and responses for IBM z/OS Connect Enterprise Edition.

- “Data stream definition for SMF123_01_V2 data” on page 172
- “Annotated fields for SMF123_01_V2 data” on page 173

Data stream definition for SMF123_01_V2 data

For prerequisite requirements for defining SMF data streams, see “Preparing the Configuration Tool to support SMF record types for Z Operations Analytics” on page 7.

Table 57 on page 172 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 57. Data stream definition for SMF123_01_V2

Type of node in the policy	Required configuration value
Data Stream	**SMF123_01_V2** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **z/OS** > **z/OS Connect EE**, and select the **SMF123_01_V2** check box.
Filter Transform	Not required

Table 57. Data stream definition for SMF123_01_V2 (continued)

Type of node in the policy	Required configuration value
Subscriber	For the appropriate values, see "Subscribers for each type of source data" on page 10.

Annotated fields for SMF123_01_V2 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 58. Annotated fields for SMF123_01_V2 data

Field	Description	Corresponding SMF field
API_NAME	API name	SMF123S1_API_NAME
ELAPSED_TIME	The time duration, in seconds, for the API request and response	SMF123S1_TIME_ZC_EXIT-F123S1_TIME_ZC_ENTRY
END_TIME	The time when the API response completed	TIMESTAMP
HTTP_RESP_CODE	HTTP response code	SMF123S1_HTTP_RESP_CODE
MAX_ELAPSED_TIME	The time duration, in seconds, for the API request and response	SMF123S1_TIME_ZC_EXIT-F123S1_TIME_ZC_ENTRY
MAX_SOR_ELAPSED_TIME	The time duration, in seconds, that is spent in the system of record	SMF123S1_TIME_SOR_RECV-SMF123S1_TIME_SOR_SENT
MAX_ZCONN_ELAPSED_TIME	The time duration, in seconds, that the API request and response is spent in IBM z/OS Connect Enterprise Edition	ZCONN_INTV
MVS_SYSTEM_ID	The MVS system ID, which is also the SMF system ID	SM123_SID
REC_COUNT	For this record type, this value is always 1.	Not applicable
REQ_HDR1	Request header *header1name*:*header1value*	SMF123S1_REQ_HDR1
REQ_HDR2	Request header *header2name*:*header2value*	SMF123S1_REQ_HDR2
REQ_HDR3	Request header *header3name*:*header3value*	SMF123S1_REQ_HDR3
REQ_HDR4	Request header *header4name*:*header4value*	SMF123S1_REQ_HDR4
REQ_ID	Request identifier that is unique within an IBM z/OS Connect Enterprise Edition server instance	SMF123S1_REQ_ID

Table 58. Annotated fields for SMF123_01_V2 data (continued)

Field	Description	Corresponding SMF field
REQ_METHOD	The HTTP request method for this API request, which is, for example, one of the following values: • GET • POST • PUT • DELETE	SMF123S1_REQ_METHOD
REQ_SIZE	The payload length, in bytes, for the API request	SMF123S1_REQ_PAYLOAD_LEN
REQ_TARGET_URI	The target URI for the API request	SMF123S1_REQ_TARGET_URI
RESP_HDR1	Response header *header1name*:*header1value*	SMF123S1_RESP_HDR1
RESP_HDR2	Response header *header2name*:*header2value*	SMF123S1_RESP_HDR2
RESP_HDR3	Response header *header3name*:*header3value*	SMF123S1_RESP_HDR3
RESP_HDR4	Response header *header4name*:*header4value*	SMF123S1_RESP_HDR4
RESP_SIZE	The payload length, in bytes, for the API response	SMF123S1_RESP_PAYLOAD_LEN
SERVER_JOBID	The job ID of the IBM z/OS Connect Enterprise Edition server (JSABJBID)	SMF123_SERVER_JOBID
SERVER_JOBNAME	The job name of the IBM z/OS Connect Enterprise Edition server (JSABJBNM)	SMF123_SERVER_JOBNAME
SERVER_STOKEN	The token that uniquely identifies the address space for the IBM z/OS Connect Enterprise Edition server (ASSBSTKN)	SMF123_SERVER_STOKEN
SERVER_SYSPLEX	The name of the sysplex (ECVTSPLX) where the IBM z/OS Connect Enterprise Edition server is running	SMF123_SERVER_SYSPLEX
SERVER_SYSTEM	The name of the system (CVTSNAME) where the IBM z/OS Connect Enterprise Edition server is running	SMF123_SERVER_SYSTEM
SERVICE_NAME	The name of the IBM z/OS Connect Enterprise Edition service that processes the API request	SMF123S1_SERVICE_NAME
SOR_ELAPSED_TIME	The time duration, in seconds, that is spent in the system of record	SMF123S1_TIME_SOR_RECV- SMF123S1_TIME_SOR_SENT

Table 58. Annotated fields for SMF123_01_V2 data (continued)

Field	Description	Corresponding SMF field
SOR_IDENTIFIER	The identifier for the system of record. This value is the value of the com.ibm.zosconnect.spi.Data.getData(SOR_IDENTIFIER) method.	SMF123S1_SOR_IDENTIFIER
SOR_REFERENCE	A reference to the element in the server.xml file that identifies the connection to the system of record. This value is the value of the com.ibm.zosconnect.spi.Data.getData(SOR_REFERENCE) method.	SMF123S1_SOR_REFERENCE
SOR_RESOURCE	The system of record resource. This value is the value of the com.ibm.zosconnect.spi.Data.getData(SOR_RESOURCE) method.	SMF123S1_SOR_RESOURCE
SP_NAME	The service provider name. This value is the value of the com.ibm.zosconnect.spi.Data.getData(SERVICE_PROVIDER_NAME) method.	SMF123S1_SP_NAME
START_TIME	The time when the API request was initiated	TIMESTAMP
SUB_SYSTEM_ID	The subsystem ID, which is set by the **SUBSYS** parameter that is specified in the SMF macros. The default value for the **SUBSYS** parameter is ZCON.	SM123_SSI
TIME_SOR_RECV	The time when the API response was received from the system of record	SMF123S1_TIME_SOR_RECV
TIME_SOR_SENT	The time when the API request was sent to the system of record	SMF123S1_TIME_SOR_SENT
TIME_ZC_ENTRY	The time when the API request was received by the IBM z/OS Connect Enterprise Edition server	SMF123S1_TIME_ZC_ENTRY
TIME_ZC_EXIT	The time when the response was ready to be returned to the HTTP client	SMF123S1_TIME_ZC_EXIT
TRACKING_TOKEN	Tracking token	SMF123S1_TRACKING_TOKEN
USER_NAME	The user name, which is a distributed or System Authorization Facility (SAF) identity	SMF123S1_USER_NAME
USER_NAME_MAPPED	SAF mapped user name	SMF123S1_USER_NAME_MAPPED
ZCONN_ELAPSED_TIME	The total time duration, in seconds, for the API request and response in IBM z/OS Connect Enterprise Edition	ZCONN_INTV

SMF123_01_V2_SUMRY record type

SMF123_01_V2_SUMRY records represent SMF type 123 subtype 1 version 2 records that are summarized over time. These records are generated in a specified time interval and contain data about individual API provider requests and responses for IBM z/OS Connect Enterprise Edition.

- "Data stream definition for SMF123_01_V2_SUMRY data" on page 176
- "Annotated fields for SMF123_01_V2_SUMRY data" on page 176

Data stream definition for SMF123_01_V2_SUMRY data

For prerequisite requirements for defining SMF data streams, see "Preparing the Configuration Tool to support SMF record types for Z Operations Analytics" on page 7.

Table 59 on page 176 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 59. Data stream definition for SMF123_01_V2_SUMRY data

Type of node in the policy	Required configuration value
Data Stream	**SMF123_01_V2_SUMRY** **To select this data stream in the Configuration Tool:** In the "Select data stream" window, click **IBM Z Operations Analytics** > **z/OS** > **z/OS Connect EE**, and select the **SMF123_01_V2_SUMRY** check box.
Filter Transform	Not required
Subscriber	For the appropriate values, see "Subscribers for each type of source data" on page 10.

Annotated fields for SMF123_01_V2_SUMRY 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 60. Annotated fields for SMF123_01_V2_SUMRY data

Field	Description	Corresponding SMF field
API_NAME	API name	SMF123S1_API_NAME
ELAPSED_TIME	The average time duration, in seconds, for all SMF records that are summarized in this event	SMF123S1_TIME_ZC_EXIT-F123S1_TIME_ZC_ENTRY
END_TIME	The time when the API response completed	TIMESTAMP
HTTP_RESP_CODE	HTTP response code	SMF123S1_HTTP_RESP_CODE
MAX_ELAPSED_TIME	The maximum time duration, in seconds, for all SMF records that are summarized in this event	SMF123S1_TIME_ZC_EXIT-F123S1_TIME_ZC_ENTRY

Table 60. Annotated fields for SMF123_01_V2_SUMRY data (continued)

Field	Description	Corresponding SMF field
MAX_SOR_ELAPSED_TIME	The maximum time, in seconds, that is spent in the system of record for all SMF records that are summarized in this event	SMF123S1_TIME_SOR_RECV-SMF123S1_TIME_SOR_SENT
MAX_ZCONN_ELAPSED_TIME	The maximum time duration, in seconds, that is spent in IBM z/OS Connect Enterprise Edition for all SMF records that are summarized in this event	ELAPSED_TIME-SOR_ELAPSED_TIME
MVS_SYSTEM_ID	The MVS system ID, which is also the SMF system ID	SM123_SID
REC_COUNT	The number of SMF 123 records that are summarized in this event	Not applicable
REQ_METHOD	The HTTP request method for this API request, which is, for example, one of the following values: • GET • POST • PUT • DELETE	SMF123S1_REQ_METHOD
REQ_SIZE	The total size, in bytes, of the API request payload length	SMF123S1_REQ_PAYLOAD_LEN
REQ_TARGET_URI	The target URI for the API request	SMF123S1_REQ_TARGET_URI
RESP_SIZE	The total size, in bytes, of the API response payload length	SMF123S1_RESP_PAYLOAD_LEN
SERVER_JOBID	The job ID of the IBM z/OS Connect Enterprise Edition server (JSABJBID)	SMF123_SERVER_JOBID
SERVER_JOBNAME	The job name of the IBM z/OS Connect Enterprise Edition server (JSABJBNM)	SMF123_SERVER_JOBNAME
SERVER_STOKEN	The token that uniquely identifies the address space for the IBM z/OS Connect Enterprise Edition server (ASSBSTKN)	SMF123_SERVER_STOKEN
SERVER_SYSPLEX	The name of the sysplex (ECVTSPLX) where the IBM z/OS Connect Enterprise Edition server is running	SMF123_SERVER_SYSPLEX
SERVER_SYSTEM	The name of the system (CVTSNAME) where the IBM z/OS Connect Enterprise Edition server is running	SMF123_SERVER_SYSTEM
SERVICE_NAME	The name of the IBM z/OS Connect Enterprise Edition service that processes the API request	SMF123S1_SERVICE_NAME

Table 60. Annotated fields for SMF123_01_V2_SUMRY data (continued)

Field	Description	Corresponding SMF field
SOR_ELAPSED_TIME	The average time duration, in seconds, that is spent in the system of record for all SMF records that are summarized in this event	SMF123S1_TIME_SOR_RECV-SMF123S1_TIME_SOR_SENT
SOR_IDENTIFIER	The identifier for the system of record. This value is the value of the com.ibm.zosconnect.spi.Data.getData(SOR_IDENTIFIER) method.	SMF123S1_SOR_IDENTIFIER
SOR_REFERENCE	A reference to the element in the server.xml file that identifies the connection to the system of record. This value is the value of the com.ibm.zosconnect.spi.Data.getData(SOR_REFERENCE) method.	SMF123S1_SOR_REFERENCE
SOR_RESOURCE	The system of record resource. This value is the value of the com.ibm.zosconnect.spi.Data.getData(SOR_RESOURCE) method.	SMF123S1_SOR_RESOURCE
SP_NAME	The service provider name. This value is the value of the com.ibm.zosconnect.spi.Data.getData(SERVICE_PROVIDER_NAME) method.	SMF123S1_SP_NAME
START_TIME	The time when the API request was initiated	TIMESTAMP
SUB_SYSTEM_ID	The subsystem ID, which is set by the **SUBSYS** parameter that is specified in the SMF macros. The default value for the **SUBSYS** parameter is ZCON.	SM123_SSI
USER_NAME	The user name, which is a distributed or System Authorization Facility (SAF) identity	SMF123S1_USER_NAME
USER_NAME_MAPPED	SAF mapped user name	SMF123S1_USER_NAME_MAPPED
ZCONN_ELAPSED_TIME	The average time duration, in seconds, that is spent in IBM z/OS Connect Enterprise Edition for all SMF records that are summarized in this event	ELAPSED_TIME-SOR_ELAPSED_TIME

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 61 on page 179 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 61. 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.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 10.

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 62 on page 179 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 62. 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.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 10.

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 63 on page 180 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 63. 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`.
Filter Transform	Not required
Subscriber	IBM Operations Analytics - Log Analysis platform only. For the appropriate values, see “Subscribers for each type of source data” on page 10.

WebSphere SYSOUT data

WebSphere Application Server for z/OS SYSOUT data is from the SYSOUT job log.

Data stream definition for SYSOUT data

Table 64 on page 180 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 64. 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.
Filter Transform	Not required
Subscriber	See “Subscribers for each type of source data” on page 10.

WebSphere SYSPRINT data

WebSphere Application Server for z/OS SYSPRINT data is from the SYSPRINT job log.

Data stream definition for SYSPRINT data

Table 65 on page 181 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 65. 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.
Filter Transform	Not required
Subscriber	See "Subscribers for each type of source data" on page 10.

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 66. 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
`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

Table 66. Annotated fields for zAware interval anomaly data (continued)

Field	Description	Data type
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
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” on page 183
- “Data stream definition for zSecure data” on page 183

Data generation

The Access Monitor component of IBM Security zSecure Admin generates security events that the IBM Z Common Data Provider 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 Z Common Data Provider 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 67 on page 183 indicates the configuration values to use in defining this data stream in the IBM Z Common Data Provider Configuration Tool.

Table 67. 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.
Filter Transform	Not applicable
Subscriber	See “Subscribers for each type of source data” on page 10.

Dashboards that represent the operational data

For each IBM Z Operations Analytics analytics platform, IBM Z Operations Analytics provides dashboards in the user interface to help you troubleshoot problems in your IT operations environment. This reference lists the available dashboards for each platform.

- “IBM Operations Analytics - Log Analysis dashboards” on page 184
- “Elastic Stack dashboards” on page 184
- “Splunk dashboards” on page 185

IBM Operations Analytics - Log Analysis dashboards

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 dashboards

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 Connect Enterprise Edition API Dashboard
- z/OS Connect Enterprise Edition Request URI Dashboard
- z/OS Connect Enterprise Edition Service Dashboard
- z/OS Job Dashboard
- z/OS Security Server RACF Dashboard
- zSecure Access Monitor Dashboard

Splunk dashboards

The following dashboards are provided for the Splunk platform.

- CICS Transaction Server for z/OS Enterprise Dashboard
- 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
- 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
- Systems Dashboard
- Welcome Dashboard
- z/OS Connect Enterprise Edition API Dashboard
- z/OS Connect Enterprise Edition Request URI Dashboard
- z/OS Connect Enterprise Edition Service 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 161.

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.

Chapter 11. Troubleshooting Z Operations Analytics

This reference lists known problems that you might experience in using IBM Z Operations Analytics and describes known solutions.

About this task

Table 68 on page 193 outlines where to find troubleshooting information for each IBM Z Operations Analytics component.

Table 68. Guide to the troubleshooting topics

Affected component in your environment	Where to find information
IBM Z Common Data Provider	IBM Z Common Data Provider documentation (troubleshooting information)
Problem Insights server	• "Log files for the Problem Insights server and the machine learning system" on page 194 • "Troubleshooting problems with the Problem Insights server" on page 195 • "Network communication problems are occurring" on page 193
Machine learning system	• "Log files for the Problem Insights server and the machine learning system" on page 194 • "Troubleshooting problems with the machine learning system" on page 197 • "Network communication problems are occurring" on page 193
Z Operations Analytics on your analytics platform	• "Troubleshooting Z Operations Analytics problems that can occur on multiple analytics platforms" on page 198 • Troubleshooting on the Log Analysis platform • Troubleshooting on the Elastic Stack platform • Troubleshooting on the Splunk platform

Network communication problems are occurring

In the configuration files for some IBM Z Operations Analytics components, the administrator must specify host names or IP addresses for services to which the component must connect. If a host name (rather than an IP address), is specified, some conditions in the IT environment might prevent the host name from being correctly resolved, which can cause communication problems between the IBM Z Operations Analytics component and the services to which it must connect.

You might notice this problem in both the machine learning system CLI and the Problem Insights server.

Solution

Depending on where the problem occurs, complete one of the following steps:

- If this problem occurs in the Problem Insights server on IBM z/OS UNIX System Services, complete the following steps:

1. Edit the started task (GLAPISRV).
2. Override the SYSTCPD DD statement to explicitly identify the data set to be used to obtain the parameters that are defined by the TCPIP.DATA statement when no GLOBALTCPIPDATA statement is configured.

   For information about the TCPIP.DATA search order, see the z/OS Communications Server IP Configuration Guide.

- For all other occurrences of the problem, the workaround for resolving network communication issues is to replace the host names with IP addresses in the configuration files.

Log files for the Problem Insights server and the machine learning system

To simplify the collection of log files for the Problem Insights server, the machine learning system, and related components, the **Collect logs** option is provided in the Z Operations Analytics Setup Menu when you run the setup script `izoa-setup.sh`.

Depending on the components that are installed in your environment, the **Collect logs** option collects the following logs:

- Logs for the Problem Insight server
- Logs that are stored in the installation directory for the machine learning system.

  The system determines the installation directory by using the value of the *IZOA_HOME* environment variable. If this variable is not set, you are prompted to provide the installation directory. For more information about the *IZOA_HOME* environment variable, see "Deploying the machine learning system on IBM z/OS UNIX System Services" on page 61.
- Logs that are stored in the machine learning system instance directory.

  The system determines the instance directory by using the value of the *IZOA_INSTANCE* environment variable. If this variable is not set, the value of the **`INSTANCEDIR`** property in the `$IZOA_HOME/config/izoaml.config` file is used. For more information about the *IZOA_INSTANCE* environment variable, see "Deploying the machine learning system on IBM z/OS UNIX System Services" on page 61.
- Logs for Apache Spark on z/OS.

  These logs are collected only if the *SPARK_CONF_DIR* environment variable is set. The system determines the log directory by using the value of the *SPARK_LOG_DIR* variable in the $*SPARK_CONF_DIR*/`spark-env.sh` script. For more information about the *SPARK_CONF_DIR* environment variable, see Creating the Apache Spark configuration directory in the IBM Open Data Analytics for z/OS documentation.
- Logs for IBM Watson Machine Learning for z/OS.

  These logs are collected only if the *IML_HOME* environment variable is set. The logs are collected from the `$`*`IML_HOME`*`/iml-logs` directory. The `$`*`IML_HOME`* directory is the user-defined directory in which IBM Watson Machine Learning for z/OS customizations are stored. It is created during the configuration of IBM Watson Machine Learning for z/OS.

All collected logs are stored in the archive file `izoa-support-logs.`*`TIMESTAMP`*`.zip`, where *TIMESTAMP* has the format YYYY-MM-DD_HH:MM:SS. The log archive is stored in the directory that is specified by the value of the *TMPDIR* environment variable. If this variable is not set, the `/tmp` directory is used. If the `/tmp` directory cannot be written to, the HOME directory for the user is used.

Troubleshooting problems with the Problem Insights server

This reference lists known problems that you might experience in using the IBM Z Operations Analytics Problem Insights server.

Before you begin

If network communication problems are occurring, see “Network communication problems are occurring” on page 193.

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 server, 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.

Problem Insights server task is not responsive, or Problem Insights server process terminates without a clear cause

If the IBM Z Operations Analytics Problem Insights server task is not responsive, or if the Problem Insights server process terminates without a clear problem cause in its logs, the reason might be that the Problem Insights server generated so many messages that the Job Entry Subsystem (JES) Spool SYSOUT limit is exceeded.

This problem occurs only if the Problem Insights server is run as a started task on IBM z/OS UNIX System Services. An indication of this problem is that you might see the following type of output in the z/OS system log:

```
$HASP375 GLAPISRV  ESTIMATED KBYTES EXCEEDED
$HASP395 GLAPISRV  ENDED - ABEND=S722
```

Cause

This problem is more likely to occur if the following properties have a value of `info` or `finest`:

- **`loglevel`** property (logging level) in the `cli.config` file
- **`traceSpecification`** property (tracing level) in the `server.xml` file

However, the main cause is based on whether the JES Spool SYSOUT limit is set, and what that limit value is. For more information about the JES Spool SYSOUT limit, see IEFUSO — SYSOUT Limit Exit.

The longer that the Problem Insights server runs, the more likely it is to exceed a specified SYSOUT limit.

Solution

Messages that are written to the STDOUT and STDERR DD names in the Problem Insights server task are also written to the `messages.log` and `trace.log` files in the Problem Insights server directory on the IBM z/OS UNIX System Services file system, for example:

```
IZOA_HOME/wlp/usr/servers/piFrameworkServer/logs
```

Therefore, you can manually suppress the writing of messages to the STDOUT and STDERR DD names without a risk of data loss. To prevent Problem Insights server availability problems, manually suppress the writing of these messages if your environment enforces a restrictive SYSOUT limit. To suppress the writing of these messages, update the Problem Insights server task to redirect the output to DUMMY.

Example

```
//STDOUT DD DUMMY
//STDERR DD DUMMY
```

Problem Insights server is not operational when first started

When the IBM Z Operations Analytics Problem Insights server is started for the first time, it is not operational.

Symptom

One of the following sets of symptoms is present:

Symptom 1

The following events occur:

- The message library load during the first Problem Insights server initialization returns `NULL`.
- The `messages.log` file in the *`INSTALL_DIR`*`/wlp/usr/server/piFrameworkServer/logs` directory contains messages that are similar to the following message:

```
[3/8/20 22:15:04:361 EDT] 00000030 com.ibm.ws.logging.internal.impl.IncidentImpl
I FFDC1015I: An FFDC Incident has been created:
"java.lang.IllegalArgumentException: Cannot support TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
with currently installed providers com.ibm.ws.channel.ssl.internal.SSLConnectionLink 238"
at ffdc_20.03.08_22.15.04.0.log
```

Symptom 2

The following events occur:

- The **`analysis.sh getlibrarylist`** command returns no output.
- The `messages.log` file in the *`INSTALL_DIR`*`/wlp/usr/server/piFrameworkServer/logs` directory contains messages that are similar to the following message:

```
[3/8/20 22:15:04:361 EDT] 00000030 com.ibm.ws.logging.internal.impl.IncidentImpl
I FFDC1015I: An FFDC Incident has been created:
"java.lang.IllegalArgumentException: Cannot support TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
with currently installed providers com.ibm.ws.channel.ssl.internal.SSLConnectionLink 238"
at ffdc_20.03.08_22.15.04.0.log
```

Cause

The Java Runtime Environment that is used to start the Problem Insights server does not support some of the encryption ciphers that are configured for use by the Problem Insights server.

Solution

Complete the following steps:

1. In the *INSTALL_DIR*/config/cli.config file, update the value of the **ciphers** property to list only ciphers that are supported by the Java Runtime Environment that is used to start the Problem Insights server.
2. Save and close the *INSTALL_DIR*/config/cli.config file.
3. Restart the Problem Insights server.

Problem Insights dashboard not showing message-based insights even though data is available

This problem applies only to the Elastic Stack and Splunk platforms. The IBM Z Operations Analytics Problem Insights dashboard does not show any message-based problem insights, even though appropriate data is available in the Splunk or Elasticsearch data store.

Symptom

The Problem Insights server log file (messages.log) includes entries that are similar to the following message:

```
[9/14/20 21:45:56:368 GMT] 00000197 com.ibm.zsystem.zmanaged.piFramework.SearchWorker
E Search runner task failed to complete. java.lang.RuntimeException: Remote host closed
connection
during handshake
```

Cause

The protocol that is used when the Problem Insights server polls the Splunk or Elasticsearch data store is incorrectly defined.

Example of an incorrect protocol definition for the Splunk platform

- Splunk is configured to use the HTTP protocol for communication.
- The value of the **splunk.scheme** property in the splunk.config file is https.

Solution

Ensure that the value of the **splunk.scheme** property (in the splunk.config file) or the **es.ssl.enabled** property (in the elk.config file) reflects your actual Splunk or Elasticsearch configuration.

Troubleshooting problems with the machine learning system

This reference lists known problems that you might experience in using the IBM Z Operations Analytics machine learning system.

Before you begin

If network communication problems are occurring, see "Network communication problems are occurring" on page 193.

Error occurs when you run the izoa-setup.sh script to remove a machine learning system instance

When you run the izoa-setup.sh script to remove a machine learning system instance, an error is shown in the console messages.

Symptom

The following message is shown in the console messages:

```
DB2 SQL Error: SQLCODE=-556, SQLSTATE=42504
```

Solution

Ignore the message because it does not affect the operation.

Troubleshooting Z Operations Analytics problems that can occur on multiple analytics platforms

This reference lists known problems that you might experience in using IBM Z Operations Analytics on multiple analytics platforms.

About this task

To review problems that are specific to only one platform, see the following topics:

- Troubleshooting on the Log Analysis platform
- Troubleshooting on the Elastic Stack platform
- Troubleshooting on the Splunk platform

APPLID values for CICS Transaction Server might not be correct in the user interface

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

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

In the IBM Z Common Data Provider 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 Z Common Data Provider 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 Z Common Data Provider Configuration Tool.

Depending on your environment, you must use only ***one*** of the following platform-dependent 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 platform-dependent configuration files is being used by the IBM Z Common Data Provider Configuration Tool.

Important: If you are sending SMF data to the IBM Z Operations Analytics machine learning system, you must also use the configuration files that are specific to the machine learning system.

For more information, see “Preparing the Configuration Tool to support SMF record types for Z Operations Analytics” on page 7.

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.

About this task

The IBM Operations Analytics - Log Analysis V1.3.5 documentation also contains troubleshooting information that might be helpful.

To review problems that can occur on multiple analytics platforms, see “Troubleshooting Z Operations Analytics problems that can occur on multiple analytics platforms” on page 198.

Log files for Log Analysis platform problems

Troubleshooting information that is related to the Log Analysis platform is available in log files from IBM Z Common Data Provider, IBM Operations Analytics - Log Analysis, Logstash, and the `ioaz` Logstash output plugin.

IBM Z Common Data Provider log files

The following log files provide information about the IBM Z Common Data Provider:

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 files

The log files are in the *LOGSTASH_INSTALL_DIR*/logs/logstash-ioaz directory.

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 Z Common Data Provider is running

Verify that IBM Z Common Data Provider 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 Z Common Data Provider documentation.

Step 2: Check the `logstash-ioaz.log` and `ioaz-logstash.`*`n`*`.log` files for error or warning messages

If no IBM Z Common Data Provider 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 Z Common Data Provider. 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 Z Common Data Provider 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 Z Common Data Provider 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 Z Common Data Provider the appropriate access.
- Because the IBM Z Common Data Provider 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.

About this task

To review problems that can occur on multiple analytics platforms, see "Troubleshooting Z Operations Analytics problems that can occur on multiple analytics platforms" on page 198.

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 Z Common Data Provider

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.

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.

About this task

To review problems that can occur on multiple analytics platforms, see "Troubleshooting Z Operations Analytics problems that can occur on multiple analytics platforms" on page 198.

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 Z Common Data Provider 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 Z Common Data Provider

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 Z Common Data Provider documentation.

If data is being written to disk

1. Verify that both the IBM Z Common Data Provider 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 Z Common Data Provider Buffered Splunk Ingestion App.

2. Verify that the IBM Z Common Data Provider 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 Z Common Data Provider 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.

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.

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.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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.

IBM®